



Sun Cluster Geographic Edition System Administration Guide



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

Part No: 820-2996-10
February 2008, Revision A

Copyright 2008 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Sun Fire, Sun StorageTek, Sun StorEdge, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. ORACLE is a registered trademark of Oracle Corporation.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2008 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Sun Fire, Sun StorageTek, Sun StorEdge, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. ORACLE est une marque déposée registre de Oracle Corporation.

L'interface d'utilisation graphique OPEN LOOK et Sun a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

Contents

Preface	13
1 Introduction to Administering the Sun Cluster Geographic Edition Software	17
Sun Cluster Geographic Edition Administration	17
Sun Cluster Geographic Edition Administration Tools	18
Graphical User Interface	18
Command-Line Interface	19
Overview of Disaster Recovery Administration	19
Analyzing the Application for Suitability	21
2 Before You Begin	23
Overview of Sun Cluster Administration Concepts	23
Configuring Resources and Resource Groups	23
Configuring Logical Hostnames	24
Managing Device Groups	25
Overview of Sun Cluster Geographic Edition Administration Tasks	26
Prerequisite Administration Tasks	26
Sun Cluster Geographic Edition Administration Tasks	27
Example Sun Cluster Geographic Edition Cluster Configuration	31
3 Administering the Sun Cluster Geographic Edition Infrastructure	33
Sun Cluster Geographic Edition Infrastructure Resource Groups	33
Enabling the Sun Cluster Geographic Edition Software	34
▼ How to Enable Sun Cluster Geographic Edition Software	34
Disabling the Sun Cluster Geographic Edition Software	37
▼ How to Disable the Sun Cluster Geographic Edition Software	37
Checking the Status of the Sun Cluster Geographic Edition Infrastructure	39

Booting a Cluster	40
Applying Patches to a Sun Cluster Geographic Edition System	40
▼ How to Prepare a Sun Cluster Geographic Edition System for Patches	41
▼ How to Install Patches on a Sun Cluster Geographic Edition System	42
4 Administering Access and Security	45
Sun Cluster Geographic Edition Software and RBAC	45
Setting Up and Using RBAC	45
RBAC Rights Profiles	46
Modifying a User's RBAC Properties	46
Configuring Secure Cluster Communication Using Security Certificates	47
Configuring Secure Cluster Communication Using IPsec	47
▼ How to Configure IPsec for Secure Cluster Communication	48
5 Administering Cluster Partnerships	51
Configuring Trust Between Partner Clusters	51
▼ How to Configure Trust Between Two Clusters	51
▼ How to Remove Trust Between Two Clusters	53
Creating and Modifying a Partnership	54
▼ How to Create a Partnership	54
▼ How to Modify Partnership Properties	56
Joining an Existing Partnership	57
▼ How to Join a Partnership	57
Leaving or Deleting a Partnership	60
▼ How to Leave a Partnership	60
Resynchronizing a Partnership	61
▼ How to Resynchronize a Partnership	62
6 Administering Heartbeats	63
Introduction to Heartbeats	63
Creating a Heartbeat	64
▼ How to Create a Heartbeat	64
Creating a Heartbeat Plug-in	65
▼ How to Create Heartbeat Plug-in	65

Modifying a Heartbeat Plug-in Property	66
▼ How to Modify the Properties of a Heartbeat Plug-in	66
Deleting Heartbeats and Heartbeat Plug-ins	67
▼ How to Delete a Heartbeat	67
▼ How to Delete a Plug-in From a Heartbeat	68
Displaying Heartbeat Configuration Information	68
▼ How to Display Heartbeat Configuration Information	69
Tuning the Heartbeat Properties	69
▼ How to Modify the Heartbeat Properties	70
Creating a Heartbeat That Uses a Custom Heartbeat Plug-in	71
Creating a Custom Heartbeat Plug-in	71
▼ How to Add a Custom Heartbeat Plug-in to an Existing Default Heartbeat	72
▼ How to Create a Custom Heartbeat Plug-in and Add It to a Custom Heartbeat	73
Configuring Heartbeat-Loss Notification	75
Configuring the Heartbeat-Loss Notification Properties	75
Creating an Action Shell Script for Heartbeat-Loss	76
7 Administering Protection Groups	79
Introduction to Protection Groups	79
Creating a Protection Group That Does Not Require Data Replication	80
▼ How to Create a Protection Group That Is Configured Not to Use Data Replication	80
8 Monitoring and Validating the Sun Cluster Geographic Edition Software	83
Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software	83
Viewing the Sun Cluster Geographic Edition Log Messages	89
Displaying Configuration Information for Partnerships and Protection Groups	89
▼ How to Display Configuration Information About Partnerships	90
▼ How to Display Configuration Information About Protection Groups	90
9 Customizing Switchover and Takeover Actions	93
Creating a Role-Change Action Script	93
Configuring a Protection Group to Run a Script at Switchover or Takeover	95
▼ How to Configure a Protection Group to Run a Script at Switchover or Takeover	96

A	Standard Sun Cluster Geographic Edition Properties	97
	General Heartbeat Properties	97
	General Heartbeat Plug-in Properties	98
	Partnership Properties	99
	General Properties of a Protection Group	100
B	Legal Names and Values of Sun Cluster Geographic Edition Entities	103
	Legal Names for Sun Cluster Geographic Edition Entities	103
	Legal Values for Sun Cluster Geographic Edition Entities	104
C	Takeover Postconditions	105
	Results of a Takeover When the Partner Cluster Can Be Reached	105
	Results of a Takeover When the Partner Cluster Cannot Be Reached	106
D	Troubleshooting Sun Cluster Geographic Edition Software	109
	Troubleshooting Monitoring and Logging	109
	Configuring the Logger File to Avoid Too Many Traces	109
	Configuring the Log File to Avoid Detailed Messages From the gcr Agent	110
	Configuring the Log File to Avoid jmx Remote Traces	110
	Troubleshooting Migration Problems	110
	Resolving Problems With Application Resource Group Failover When Communication Lost With the Storage Device	110
	Troubleshooting Cluster Start and Restart	111
	Validating Protection Groups in an Error State	111
	Restarting the Common Agent Container	111
	Index	113

Figures

FIGURE 2-1	Example Cluster Configuration	31
------------	-------------------------------------	----

Tables

TABLE 1-1	Sun Cluster Geographic Edition CLI	19
TABLE 2-1	IP Addresses Required by Sun Cluster Geographic Edition Software	24
TABLE 2-2	Sun Cluster Geographic Edition Prerequisite Tasks	26
TABLE 2-3	Sun Cluster Geographic Edition Administration Tasks	27
TABLE 4-1	Sun Cluster Geographic Edition RBAC Rights Profiles	46
TABLE 8-1	Status Value Descriptions	85
TABLE A-1	General Heartbeat Properties	97
TABLE A-2	General Heartbeat Plug-in Properties	98
TABLE A-3	Partnership Properties	99
TABLE A-4	General Properties of a Protection Group	100
TABLE C-1	Takeover Results of Running the geopg takeover Command on the Secondary Cluster	105
TABLE C-2	Takeover Results of Running the geopg takeover Command on the Primary Cluster	106
TABLE C-3	Takeover Results of Running the geopg takeover Command on the Secondary Cluster When the Primary Cluster Cannot Be Reached	107
TABLE C-4	Takeover Results of Running the geopg takeover Command on the Primary Cluster When the Secondary Cluster Cannot Be Reached	107

Examples

EXAMPLE 3-1	Enabling the Sun Cluster Geographic Edition Infrastructure a Cluster	36
EXAMPLE 3-2	Disabling a Cluster	38
EXAMPLE 3-3	Displaying Whether the Sun Cluster Geographic Edition Infrastructure Has Been Enabled	39
EXAMPLE 5-1	Creating a Partnership	56
EXAMPLE 5-2	Modifying the Properties of a Partnership	57
EXAMPLE 5-3	Joining a Partnership	59
EXAMPLE 5-4	Creating and Joining a Partnership With a Remote Cluster in a Different Domain	59
EXAMPLE 5-5	Leaving a Partnership	61
EXAMPLE 5-6	Deleting a Partnership	61
EXAMPLE 5-7	Resynchronizing a Partnership	62
EXAMPLE 6-1	Creating a Heartbeat	65
EXAMPLE 6-2	Creating a Heartbeat Plug-in	66
EXAMPLE 6-3	Modifying the Properties of the Heartbeat Plug-in	67
EXAMPLE 6-4	Deleting a Heartbeat	68
EXAMPLE 6-5	Deleting a Plug-in From a Heartbeat	68
EXAMPLE 6-6	Displaying Heartbeat Configuration Information	69
EXAMPLE 6-7	Modifying the Properties of the Default Heartbeat	71
EXAMPLE 6-8	Adding a Custom Heartbeat Plug-in to the Default Heartbeat	73
EXAMPLE 6-9	Adding a Custom Heartbeat Plug-in to a New Custom Heartbeat	74
EXAMPLE 6-10	Configuring Heartbeat-Loss Notification for an Existing Partnership	76
EXAMPLE 6-11	How a Notification Action Script Parses the Command-Line Information Provided by the Sun Cluster Geographic Edition Software	77
EXAMPLE 7-1	Creating and Configuring a Protection Group That Is Configured to Not Use Data Replication	82
EXAMPLE 8-1	Displaying Partnership Configuration Information	90
EXAMPLE 8-2	Displaying Configuration Information About a Protection Group	91
EXAMPLE 9-1	Switchover Action Script for Updating the DNS	94

EXAMPLE 9-2	Configuring a Protection Group to Run a Command at Cluster Switchover or Takeover	96
-------------	---	----

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), of Sun Cluster, and expertise with the volume manager software that is used with Sun Cluster software.

Note –

This release of Sun Cluster Geographic Edition supports the following releases of Sun's availability suite software:

- Sun StorageTek Availability Suite 4
- Sun StorEdge Availability Suite 3.2.1

In this manual, references to Sun StorageTek™ Availability Suite software also apply to Sun StorEdge™ Availability Suite software unless specifically stated otherwise.

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	<i>Sun Cluster Geographic Edition Overview</i>
Glossary	<i>Sun Java Enterprise System Glossary</i>
Hardware administration	Individual hardware administration guides

Topic	Documentation
Software installation	<i>Sun Cluster Geographic Edition Installation Guide</i>
System administration	<i>Sun Cluster Geographic Edition System Administration Guide</i> <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>
Command and function references	<i>Sun Cluster Geographic Edition Reference Manual</i>

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

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

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

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, 3.2 2/08)
- The contents of the `/var/cacao/instances/default/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
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>geoadm -V</code>	Displays the Sun Cluster Geographic Edition software release information
<code>cluster status</code>	Provides a snapshot of the cluster status
<code>cluster show</code>	Lists cluster configuration information
<code>geoadm status</code>	Displays 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 conventions that are used in this book.

TABLE P-1 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>
Bourne shell and Korn shell	<code>\$</code>
Bourne shell and Korn shell for superuser	<code>#</code>

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 sections:

- “Sun Cluster Geographic Edition Administration” on page 17
- “Sun Cluster Geographic Edition Administration Tools” on page 18
- “Overview of Disaster Recovery Administration” on page 19
- “Analyzing the Application for Suitability” on page 21

Sun Cluster Geographic Edition Administration

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. You can install, configure, start, use, stop, and uninstall the Sun Cluster Geographic Edition software 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 by using the CLI.

Graphical User Interface

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

Note – To administer Sun Cluster Geographic Edition software by using the Sun Cluster Manager – Geographic Edition GUI, ensure that the root passwords are 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 by using the `geoadm start` command. Use a shell to run the `geoadm start` and `geoadm stop` commands. For information about enabling and disabling the Sun Cluster Geographic Edition infrastructure, see [Chapter 3, “Administering the Sun Cluster Geographic Edition Infrastructure.”](#)

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 run the `geops join-partnership` command.

To start the GUI, go to the following URL from any Java-enabled and Javascript-enabled browser, and log in to the Sun Administration Console as root.

Note – RBAC is not supported in the GUI.

`https://clustername:6789`

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 displays 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. An administrator can expect the following events:

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 54](#).
2. The administrator receives the notification email and follows the company procedure to verify that the disconnect occurred because of a situation that requires a takeover by the secondary cluster. Because a takeover might take a long time, depending on the requirements of the applications being protected, 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, see one of following data replication guides:

- “Detecting Cluster Failure on a System That Uses Sun StorageTek Availability Suite Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*
 - “Detecting Cluster Failure on a System That Uses Hitachi TrueCopy Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
 - “Detecting Cluster Failure on a System That Uses EMC Symmetrix Remote Data Facility Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*
3. Because the `cluster-paris` cluster cannot be brought online again for at least another day, the administrator runs a `geopg takeover` command on a node in the cluster in New York. This command starts the protection group on the secondary cluster `cluster-newyork` in New York.

For information about performing a takeover on a system, see one of the following data replication guides:

- “Forcing a Takeover on Systems That Use Sun StorageTek Availability Suite” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*
 - “Forcing a Takeover on a System That Uses Hitachi TrueCopy Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
 - “Forcing a Takeover on a System That Uses EMC Symmetrix Remote Data Facility Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*
4. After the takeover, the secondary cluster `cluster-newyork` becomes the new primary cluster. The failed cluster in Paris is still configured to be the primary cluster. Therefore, when the `cluster-paris` cluster restarts, the cluster detects that the primary cluster was down and lost contact with the partner cluster. Then, the `cluster-paris` cluster enters an error state that requires administrative action to clear. You might also be required to recover and resynchronize data on the cluster.

For information about recovering data after a takeover, see one of the following data replication guides:

- “Recovering Sun StorageTek Availability Suite Data After a Takeover” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*
- “Recovering From a Hitachi TrueCopy Data Replication Error” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- “Recovering From an EMC Symmetrix Remote Data Facility Data Replication Error” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

Analyzing the Application for Suitability

This section describes the guidelines you must follow in creating applications to be managed by Sun Cluster Geographic Edition software.

Before you create an application to be managed by Sun Cluster Geographic Edition software, determine whether the application satisfies the following requirements for being made highly available or scalable.

Note – If the application fails to meet all requirements, modify the application source code to make it highly available or scalable.

- Both network-aware (client-server model) and network-unaware (client-less) applications are potential candidates for being made highly available or scalable in the Sun Cluster Geographic Edition environment. However, Sun Cluster Geographic Edition cannot provide enhanced availability in timesharing environments in which applications are run on a server that is accessed through `telnet` or `rlogin`.
- The application must be crash tolerant. That is, it must recover disk data (if necessary) when it is started after an unexpected node death. Furthermore, the recovery time after a crash must be bounded. Crash tolerance is a prerequisite for making an application highly available because the ability to recover the disk and restart the application is a data integrity issue. The data service is not required to be able to recover connections.
- The application must not depend on the physical host name of the node on which it is running.
- The application must operate correctly in environments in which multiple IP addresses are configured to go up. Examples include environments with multihomed hosts, in which the node is located on more than one public network, and environments with nodes on which multiple, logical interfaces are configured to go up on one hardware interface.
- Application binaries and libraries can be located locally on each node or in the cluster file system. The advantage of being located in the cluster file system is that a single installation is sufficient. The disadvantage is that when you use rolling upgrade for Sun Cluster software, the binaries are in use while the application is running under the control of the Resource Group Manager (RGM).
- The client must have capacity to retry a query automatically if the first attempt times out. If the application and the protocol already handle the case of a single server crashing and rebooting, they also can handle the containing resource group failing over or switching over.
- The application must not have UNIX® domain sockets or named pipes in the cluster file system.

A scalable service must meet all the preceding conditions for high availability as well as the following additional requirements.

- The application must have the ability to run multiple instances, all operating on the same application data in the cluster file system.
- The application must provide data consistency for simultaneous access from multiple nodes.
- The application must implement sufficient locking with a globally visible mechanism, such as the cluster file system.

For a scalable service, application characteristics also determine the load-balancing policy. For example, the load-balancing policy `Lb_weighted`, which allows any instance to respond to client requests, does not work for an application that makes use of an in-memory cache on the server for client connections. In this case, you should specify a load-balancing policy that restricts a given client's traffic to one instance of the application. The load-balancing policies `Lb_sticky` and `Lb_sticky_wild` repeatedly send all requests by a client to the same application instance, where they can make use of an in-memory cache. If multiple client requests come in from different clients, the RGM distributes the requests among the instances of the service.

See Chapter 2, “Developing a Data Service,” in *Sun Cluster Data Services Developer's Guide for Solaris OS* for more information about setting the load-balancing policy for scalable data services.

The application must be able to meet the following data replication requirements:

- Information replicated must not be host- or cluster-specific.
When the application fails over to the remote site, the application might run on a host with a different IP address. To allow client nodes to find the remote site, use a Sun Cluster Geographic Edition action script to update the DNS/NIS mapping.
- If you don't want your application to tolerate any data loss, the application should use synchronous replication.

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 contains the following sections:

- [“Overview of Sun Cluster Administration Concepts” on page 23](#)
- [“Overview of Sun Cluster Geographic Edition Administration Tasks” on page 26](#)
- [“Example Sun Cluster Geographic Edition Cluster Configuration” on page 31](#)

Overview of Sun Cluster Administration Concepts

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

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

Configuring Resources and Resource Groups

You use either Sun Cluster commands or the Sun Cluster 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 that are described in [Appendix B, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

For more information about using the `geoadm start` command, see [“Enabling the Sun Cluster Geographic Edition Software” on page 34](#).

Note – If you are using Sun StorageTek Availability Suite for data replication, a logical hostname is created for each device group to be replicated. For more information, see Chapter 1, “Replicating Data With Sun StorageTek Availability Suite Software,” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*.

The following table lists the Sun Cluster and Sun Cluster Geographic Edition components that require 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 the Solaris OS 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	<ul style="list-style-type: none"> ▪ 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
Console-access device	1
Logical addresses	1 per logical host resource, per subnet

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

Component	Number of IP Addresses Needed
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 StorageTek Availability Suite 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 Software on the Cluster,” 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 5, “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 1, “Replicating Data With Sun StorageTek Availability Suite Software,” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*, Chapter 1, “Replicating Data With Hitachi TrueCopy Software,” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*, and Chapter 1, “Replicating Data With EMC Symmetrix Remote Data Facility Software,” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*.

Overview of Sun Cluster Geographic Edition Administration Tasks

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

- [“Prerequisite Administration Tasks” on page 26](#)
- [“Sun Cluster Geographic Edition Administration Tasks” on page 27](#)

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 <code>SC-clustername</code> to the cluster name you want to use with the Sun Cluster Geographic Edition software.	Use the <code>cluster(1CL)</code> command. For more information, see “How to Enable Sun Cluster Geographic Edition Software” on page 34 .
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 Software on the Cluster,” in <i>Sun Cluster Software Installation Guide for Solaris OS</i> .
Install and configure your data replication product.	See the Sun StorageTek Availability Suite, Hitachi TrueCopy, or EMC Symmetrix Remote Data Facility documentation. This step is required before you can create protection groups with the <code>geopg create</code> command.
Port and configure application configuration and corresponding resource groups on clusters that are candidates for partnership.	You can use the Sun Cluster <code>scsnapshot</code> tool to facilitate porting of application resource groups. See “Creating and Modifying a Partnership” on page 54 for more information.
Enable the common agent container on all nodes of both clusters.	See “Enabling the Sun Cluster Geographic Edition Software” on page 34 .

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 <i>Sun Cluster Geographic Edition Installation Guide</i> .
Set up security between the candidate partner clusters.	<ul style="list-style-type: none"> ■ Exchange certificates, as described in “Configuring Secure Cluster Communication Using Security Certificates” on page 47. ■ (Optional) Configure a secure logical hostname that uses IP Security Architecture (IPsec), as described in “Configuring Secure Cluster Communication Using IPsec” on page 47.
Enable the Sun Cluster Geographic Edition software.	<p>Use the <code>geoadm start</code> command.</p> <p>For more information, see “Enabling the Sun Cluster Geographic Edition Software” on page 34.</p>
Create partnerships.	<p>See “How to Create a Partnership” on page 54. This procedure includes the following:</p> <ul style="list-style-type: none"> ■ Modifying the default heartbeat. For more information, see Chapter 6, “Administering Heartbeats.” ■ Configuring loss of heartbeat notification. For more information, see “Configuring Heartbeat-Loss Notification” on page 75.
Configure data replication.	<p>For information about replicating data by using Sun StorageTek Availability Suite, see Chapter 1, “Replicating Data With Sun StorageTek Availability Suite Software,” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i>.</p> <p>For information about replicating data by using Hitachi TrueCopy, see Chapter 1, “Replicating Data With Hitachi TrueCopy Software,” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i>.</p> <p>For information about replicating data by using EMC Symmetrix Remote Data Facility, see Chapter 1, “Replicating Data With EMC Symmetrix Remote Data Facility Software,” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>.</p>

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

Task	Description and Documentation
<p>Create protection groups.</p>	<ul style="list-style-type: none"> ■ Create a protection group. See one of the following data replication guides: <ul style="list-style-type: none"> ■ “How to Create and Configure a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “How to Create and Configure an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> ■ Add data replication device groups. See one of the following data replication guides: <ul style="list-style-type: none"> ■ “How to Add a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “How to Add a Data Replication Device Group to an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> ■ Add application resource groups to the protection group. See one of the following data replication guides: <ul style="list-style-type: none"> ■ “How to Add an Application Resource Group to a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “How to Add an Application Resource Group to a Hitachi TrueCopy Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “How to Add an Application Resource Group to an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> ■ Create a protection group that does not require data replication. See “Creating a Protection Group That Does Not Require Data Replication” on page 80.

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

Task	Description and Documentation
Bring the protection groups online.	<p>See one of the following data replication guides:</p> <ul style="list-style-type: none"> ■ “How to Activate a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “How to Activate a Hitachi TrueCopy Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “How to Activate an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>
Test the configured partnership and protection groups to validate the setup.	<p>Perform a trial switchover or takeover and test some simple failure scenarios. See one of the following data replication guides:</p> <ul style="list-style-type: none"> ■ Chapter 3, “Migrating Services That Use Sun StorageTek Availability Suite Data Replication,” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ Chapter 3, “Migrating Services That Use Hitachi TrueCopy Data Replication,” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ Chapter 3, “Migrating Services That Use EMC Symmetrix Remote Data Facility Data Replication,” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> <p>Note – You cannot perform personality swaps if you are running EMC Symmetrix Remote Data Facility/Asynchronous data replication.</p>
Migrate services to the partner cluster.	<p>See one of the following data replication guides:</p> <ul style="list-style-type: none"> ■ “How to Switch Over a Sun StorageTek Availability Suite Protection Group From Primary to Secondary” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “How to Switch Over a Hitachi TrueCopy Protection Group From Primary to Secondary” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “How to Switch Over an EMC Symmetrix Remote Data Facility Protection Group From Primary to Secondary” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> <p>Note – You cannot perform personality swaps if you are running EMC Symmetrix Remote Data Facility/Asynchronous data replication.</p>

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

Task	Description and Documentation
Take over services from primary to secondary during a disaster.	<p>See one of the following data replication guides:</p> <ul style="list-style-type: none"> ■ “How to Force Immediate Takeover of Sun StorageTek Availability Suite Services by a Secondary Cluster” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “How to Force Immediate Takeover of Hitachi TrueCopy Services by a Secondary Cluster” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “How to Force Immediate Takeover of EMC Symmetrix Remote Data Facility Services by a Secondary Cluster” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>
Recover from a takeover.	<ul style="list-style-type: none"> ■ Data recovery and error repair outside of the Sun Cluster Geographic Edition infrastructure. See the Sun StorageTek Availability Suite, Hitachi TrueCopy, or EMC Symmetrix Remote Data Facility documentation. ■ Resynchronize the partner clusters. See “Recovering Sun StorageTek Availability Suite Data After a Takeover” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i>, “Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy Replication” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i>, or “Recovering Services to a Cluster on a System That Uses EMC Symmetrix Remote Data Facility Replication” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>.
Take a protection group offline.	<p>See “How to Deactivate a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i>, “How to Deactivate a Hitachi TrueCopy Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i>, or “How to Deactivate an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>.</p>
Delete a protection group.	<p>See “How to Delete a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i>, “How to Delete a Hitachi TrueCopy Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i>, or “How to Delete an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>.</p>
Delete a partnership.	<p>See “Leaving or Deleting a Partnership” on page 60.</p>
Disable the Sun Cluster Geographic Edition software.	<p>See “How to Disable the Sun Cluster Geographic Edition Software” on page 37.</p>
Uninstall the Sun Cluster Geographic Edition software.	<p>See the <i>Sun Cluster Geographic Edition Installation Guide</i>.</p>

Example Sun Cluster Geographic Edition Cluster Configuration

The following figure describes a 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`.

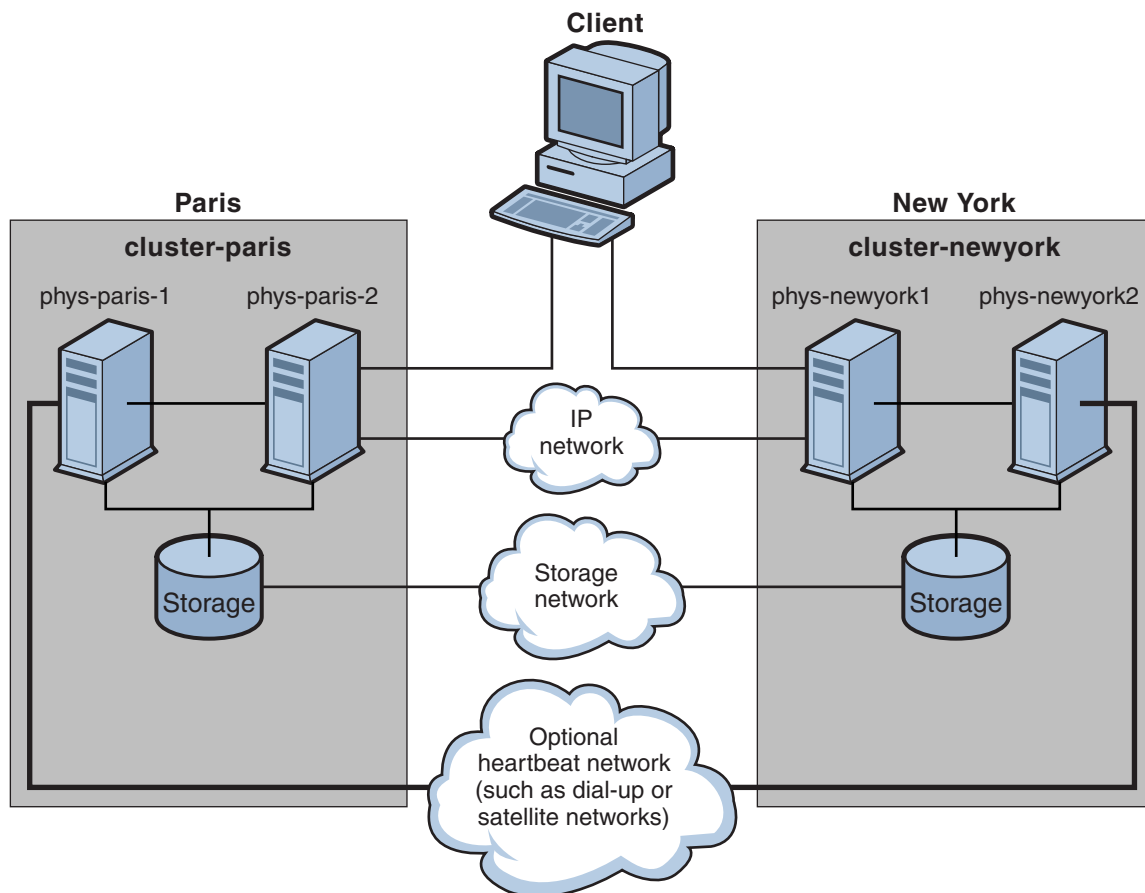


FIGURE 2-1 Example Cluster Configuration

Administering the Sun Cluster Geographic Edition Infrastructure

This chapter contains information about enabling your cluster for participation in a partnership. 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 sections:

- “Sun Cluster Geographic Edition Infrastructure Resource Groups” on page 33
- “Enabling the Sun Cluster Geographic Edition Software” on page 34
- “Disabling the Sun Cluster Geographic Edition Software” on page 37
- “Checking the Status of the Sun Cluster Geographic Edition Infrastructure” on page 39
- “Booting a Cluster” on page 40
- “Applying Patches to a Sun Cluster Geographic Edition System” on page 40

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` – 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. An entry in the naming services 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 `clresource status` command. For more information about this command, see the `clresource(1CL)` 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

This procedure enables the Sun Cluster Geographic Edition infrastructure on the local cluster only. Repeat this procedure on all the clusters of your geographically separated cluster.

Before You Begin 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 Sun Cluster Manager is running.
- The Sun Cluster Geographic Edition software is installed.
- The cluster has been configured for secure cluster communication by using security certificates, that is, nodes within the same cluster must share the same security certificates. This is done during Sun Cluster installation.

When you upgrade to Sun Cluster 3.2 software, the security certificates must be identical on all nodes of the cluster. Therefore, you must copy the security certificates manually from one node of the cluster to the other nodes of the cluster. For more information on copying the security files for the common agent container, see the procedures in “How to Finish Upgrade to Sun Cluster 3.2 12/07 Software” in *Sun Cluster Upgrade Guide for Solaris OS*.

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 45](#).

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

```
# cluster list
```

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

```
# cluster rename -c newclustername oldclustername
```

For more information, see the `cluster(1CL)` 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.

4 Confirm that the naming service and the local hosts files contain a host entry that matches the cluster name.

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

5 On a node of the cluster, start the Sun Cluster Geographic Edition infrastructure.

```
# geoadm start
```

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

6 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 [“Sun Cluster Geographic Edition Infrastructure Resource Groups” on page 33](#).

```
# geoadm show
# clresourcegroup status
# clresource status
```

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

The output for the `clresourcegroup sstatus` and `clresource status` commands displays 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 `clresourcegroup(1CL)` and `clresource(1CL)` man pages.

Example 3-1 Enabling the Sun Cluster Geographic Edition Infrastructure a Cluster

This example enables the Sun Cluster Geographic Edition software on the `cluster-paris` cluster.

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

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

2. Ensure that 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# clresourcegroup status
```

```
=== Cluster Resource Groups ===
```

Group Name	Node Name	Suspended	Status
-----	-----	-----	-----
geo-clusterstate	phys-paris-1	No	Online
	phys-paris-2	No	Online
geo-infrastructure	phys-paris-1	No	Online
	phys-paris-2	No	Offline

```
# clresource status
```

```
=== Cluster Resources ===
```

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

Next Steps For information about creating protection groups, see the Sun Cluster Geographic Edition Data Replication Guide that corresponds to the type of data replication software you are using.

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 Ensure that all protection groups on the local cluster are offline.

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 45.](#)

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 83.](#)



Caution – If you want to keep the application resource groups online while deactivating a protection group, follow the procedure described in the following data replication guides:

- “How to Deactivate a Sun StorageTek Availability Suite Protection Group” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*
- “How to Deactivate a Hitachi TrueCopy Protection Group” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- “How to Deactivate an EMC Symmetrix Remote Data Facility Protection Group” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

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 – Disabling the Sun Cluster Geographic Edition software removes only the infrastructure resource groups. Resource groups that have been created to support data replication are not removed unless you remove the protection group that the resource groups are supporting by using the `geopg delete` command.

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# clresourcegroup status
```

For more information, see the `clresourcegroup(1CL)` man page.

Example 3-2 Disabling a Cluster

This example disables the `cluster-paris` cluster.

1. Confirm that all protection groups are offline.

```
phys-paris-1# geoadm status

Cluster: cluster-paris

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

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

Protection group "tcpg"      :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. Confirm that 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.

```

phys-paris-1# clresourcegroup status
phys-paris-1#

```

Checking the Status of the Sun Cluster Geographic Edition Infrastructure

Use the `geoadm show` command to determine whether the Sun Cluster Geographic Edition infrastructure is enabled on the local cluster and 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`.

EXAMPLE 3-3 Displaying Whether the Sun Cluster Geographic Edition Infrastructure Has Been Enabled

This example displays information on the `phys-paris-1` node of the `cluster-paris` cluster.

```

phys-paris-1# geoadm show

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

```

EXAMPLE 3-3 Displaying Whether the Sun Cluster Geographic Edition Infrastructure Has Been Enabled (Continued)

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

Booting a Cluster

The following events take place when you boot a cluster:

1. After the Sun Cluster infrastructure is enabled, the Sun Cluster Geographic Edition software starts automatically. Verify that the software started successfully 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 83.

Applying Patches to a Sun Cluster Geographic Edition System

Observe the following guidelines and requirements to patch Sun Cluster Geographic Edition software:

- You must run the same patch levels for Sun Cluster software and the common agent container software on all nodes of both clusters.
- The patch level for each node on which you have installed Sun Cluster Geographic Edition software must meet the Sun Cluster patch-level requirements.
- All nodes in one cluster must have the same version of Sun Cluster Geographic Edition software and the same patch level. However, primary and secondary clusters can run different versions of Sun Cluster Geographic Edition software, provided that each version is correctly patched and the versions are no more than one release different. For example, if one cluster is running Sun Cluster Geographic Edition 3.2 software that has been fully patched, and the partner cluster is running Sun Cluster Geographic Edition 3.2 2/08 software that has been fully patched, then both clusters should be brought to the same patch level as soon as possible. Also, if both partner clusters are running Sun Cluster Geographic Edition 3.2 2/08, then both partner clusters should be brought to the same patch level as soon as possible.
- To ensure that the patches have been installed properly, install the patches on your secondary cluster before you install the patches on the primary cluster.
- For additional information about Sun Cluster Geographic Edition patches, see the patch README file.

- See “Required Patches” in *Sun Cluster Geographic Edition 3.2 2/08 Release Notes* for a list of required patches.

▼ How to Prepare a Sun Cluster Geographic Edition System for Patches

1 Ensure that the cluster is functioning properly.

To view the current status of the cluster, run the following command from any node:

```
% cluster status
```

See the `cluster(1CL)` man page for more information.

Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.

2 Become superuser on a node of the cluster.

3 Remove all application resource groups from protection groups.

This step ensures that resource groups are not stopped when you later stop the protection groups.

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

See the `geopg(1M)` man page for more information.

4 Perform the preceding steps on all clusters that have a partnership with this cluster.

5 Stop all protection groups that are active on the cluster.

```
# geopg stop protectiongroup -e local
```

See the `geopg(1M)` man page for more information.

6 Stop the Sun Cluster Geographic Edition infrastructure.

```
# geoadm stop
```

Shutting down the infrastructure ensures that a patch installation on one cluster does not affect the other cluster in the partnership.

See the `geoadm(1M)` man page for more information.

7 On each node, stop the common agent container.

```
# /usr/sbin/cacaoadm stop
```

Note – You must use common agent container 2, which is located in the `/usr/sbin` directory. Use the `/usr/sbin/cacaoadm -V` command to check which version of common agent container you are using.

Next Steps Install the required patches for the Sun Cluster Geographic Edition software. Go to [“How to Install Patches on a Sun Cluster Geographic Edition System”](#) on page 42.

▼ How to Install Patches on a Sun Cluster Geographic Edition System

Perform this procedure on all nodes of the cluster.

Patch the secondary cluster before you patch the primary cluster, to permit testing.

Before You Begin Perform the following tasks:

- Ensure that the Solaris OS is installed to support Sun Cluster Geographic Edition software. If Solaris software is already installed on the node, you must ensure that the Solaris installation meets the requirements for Sun Cluster Geographic Edition software and any other software that you intend to install on the cluster.
- Ensure that Sun Cluster Geographic Edition software packages are installed on the node.
- Ensure that you completed all steps in [“How to Prepare a Sun Cluster Geographic Edition System for Patches”](#) on page 41.

1 Ensure that all the nodes are online and part of the cluster.

To view the current status of the cluster, run the following command from any node:

```
% cluster status
```

See the `cluster(1CL)` man page for more information.

Search the `/var/adm/messages` log on the same node for unresolved error messages or warning messages.

2 Become superuser on each node.

3 On each node, install any necessary patches to support Sun Cluster Geographic Edition software by using the `patchadd` command.

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

- 4 **After you have installed all required patches on all nodes of the cluster, on each node start the common agent container.**

```
# /usr/sbin/cacaoadm start
```

Note – You must use common agent container 2, which is located in the `/usr/sbin` directory. Use the `/usr/sbin/cacaoadm -V` command to check which version of common agent container you are using.

- 5 **On one node, enable Sun Cluster Geographic Edition software.**

```
# geoadm start
```

- 6 **Add all application resource groups that you removed while you were preparing the cluster for a patch installation back to the protection group.**

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

See the `geopg(1M)` man page for more information.

- 7 **Start all the protection groups that you have added.**

```
# geopg start protectiongroup -e local [-n]
```

See the `geopg(1M)` man page for more information.

Next Steps After you patch the secondary cluster, perform a sanity test on the Sun Cluster Geographic Edition software, and then repeat this procedure on the primary cluster.

Administering Access and Security

This chapter describes how to administer access and security. It contains the following sections:

- “Sun Cluster Geographic Edition Software and RBAC” on page 45
- “Configuring Secure Cluster Communication Using Security Certificates” on page 47
- “Configuring Secure Cluster Communication Using IPsec” on page 47

Sun Cluster Geographic Edition Software and RBAC

This section describes role-based access control (RBAC) in Sun Cluster Geographic Edition software. It contains the following sections:

- “Setting Up and Using RBAC” on page 45
- “RBAC Rights Profiles” on page 46
- “Modifying a User’s RBAC Properties” on page 46

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 software, 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 in the search list. Including `files` in the search list 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 [Table 4-1](#).

TABLE 4-1 Sun Cluster Geographic Edition RBAC Rights Profiles

Rights Profile	Included Authorizations	Role Identity Permission
Geo Management	<code>solaris.cluster.geo.read</code>	Read information about the Sun Cluster Geographic Edition entities
	<code>solaris.cluster.geo.admin</code>	Perform administrative tasks with the Sun Cluster Geographic Edition software
	<code>solaris.cluster.geo.modify</code>	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
	<code>solaris.cluster.geo.read</code>	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, so cluster `cluster-paris` must be configured to trust its partner cluster `cluster-newyork` and cluster `cluster-newyork` must be configured to trust its partner cluster `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 Trust Between Partner Clusters](#)” on page 51.

For information about the example cluster configuration, see “[Example Sun Cluster Geographic Edition Cluster Configuration](#)” on page 31.

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 StorageTek Availability Suite communications, if you use the Sun StorageTek Availability Suite software 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 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 are different on the two clusters of a partnership.
- **IPsec keys file**, `/etc/init/secret/ipseckeys`. Contains keys files for specific authentication and encryption algorithms. The contents of this file are identical on both clusters of a partnership.

The following procedure configures a cluster, `cluster-paris`, for IPsec secure communication with another 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`.

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 31.

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 `ipseconf(1M)` man page.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see [Appendix B, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

a. Configure the communication policy.

The default port for the `tcp_udp` plug-in is 2084. You can specify this value in `theetc/cacao/instances/default/modules/com.sun.cluster.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`, you must reverse the policies.

```
# {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.

```
# ipseconf -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/ipseckey`. Permissions on the file should be `600`.

Add keys:

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

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-algorithm \
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 `ipseconf(1M)` man page.

Administering Cluster Partnerships

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

This chapter contains the following sections:

- “Configuring Trust Between Partner Clusters” on page 51
- “Creating and Modifying a Partnership” on page 54
- “Joining an Existing Partnership” on page 57
- “Leaving or Deleting a Partnership” on page 60
- “Resynchronizing a Partnership” on page 61

Configuring Trust Between Partner Clusters

Before you create a partnership between two clusters, you must configure the Sun Cluster Geographic Edition software for secure communication between the two clusters. The configuration must be reciprocal. For example, you must configure the cluster `cluster-paris` to trust the cluster `cluster-newyork`, and you must also configure the cluster `cluster-newyork` to trust the cluster `cluster-paris`.

▼ How to Configure Trust Between Two Clusters

Before You Begin Ensure that the following conditions are met:

- The cluster on which you want to create the partnership is running.
- The `geoadm start` command must have already been run on 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 34](#).
- The cluster name of the partner cluster is known.

- The host information of the partner cluster must be defined in the local host file. The local cluster needs to know how to reach the partner cluster by name.

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 45](#).

2 Import the public keys from the remote cluster to the local cluster.

Running this command on one node of the local cluster imports the keys from the remote cluster to one node of the cluster.

```
# geops add-trust -c remotepartnerclustername
```

```
-c remotepartnerclustername
```

Specifies the logical hostname of the cluster with which to form a partnership. The logical hostname is used by the Sun Cluster Geographic Edition software and maps to the name of the remote partner cluster. For example, a remote partner cluster name might resemble the following:

```
cluster-paris
```

When you use this option with the `add-trust` or `remote-trust` subcommand, the option specifies the alias where the public keys on the remote cluster are stored. An alias for certificates on the remote cluster has the following pattern:

```
remotepartnercluster.certificate[0-9]*
```

Keys and only keys that belong to the remote cluster should have their alias match this pattern.

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

3 Repeat the preceding steps on a node of the remote partner cluster.

If you choose to use the Sun Cluster Manager, skip this step. The Sun Cluster Manager handles all nodes in a single operation.

4 Verify trust from one node of each cluster.

```
# geops verify-trust -c remotepartnerclustername
```

This command verifies the trust from the node on which you run the command to all nodes of the partner cluster. If you choose to use the Sun Cluster Manager, it verifies the trust from all nodes of the local cluster to all nodes of the partner cluster.

See Also For a complete example of how to configure and join a partnership, see [Example 5–4](#).

▼ How to Remove Trust Between Two Clusters

Before You Begin Ensure that the following conditions are met:

- The cluster on which you want to remove trust is running.
- The cluster name of the partner cluster is known.
- The host information of the partner cluster must be defined in the local host file. The local cluster needs to know how to reach the partner cluster by name.

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 45](#).

2 If there is a partnership configured between the two clusters, dissolve that partnership.

Run the following command on both clusters:

```
# geops leave
```

3 On all nodes of both clusters, remove all keys for the remote cluster from the truststore file on the local node.

```
# geops remove-trust -c remotepartnerclustername
```

Perform this step on all the nodes of the local cluster, and then repeat this step on all nodes of the partner cluster.

-cremotepartnerclustername

Specifies the logical hostname of the cluster from which you want to remove the keys. The name for the remote cluster must be identical to the cluster name you specified when adding trust with the `geops add-trust` command. You do not need to specify the fully qualified name if the remote cluster is reachable by partial name.

When you use this option with the `add-trust` or `remote-trust` subcommand, the option specifies the alias where the public keys on the remote cluster are stored. An alias for certificates on the remote cluster has the following pattern:

```
remotepartnercluster.certificate[0-9]*
```

Keys and only keys that belong to the remote cluster should have their alias match this pattern.

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

If you choose to use the Sun Cluster Manager, it handles all nodes of a cluster in a single operation.

4 Repeat the preceding steps on a node of the remote partner cluster.

Creating and Modifying a Partnership

The Sun Cluster Geographic Edition software enables clusters to form partnerships between clusters to provide mutual protection against disasters. The clusters in a partnership monitor each other by sending heartbeat messages to each other in the same way that nodes of a single cluster do. Unlike local clusters, the clusters in a partnership use the public network for these messages, but support additional, plug-in mechanisms as well.

You create only one partnership between two specific clusters by using the `geops(1M)` command. After you have created a partnership, you can use this command to modify the properties of this partnership.

When creating partnerships, ensure that the name of all the clusters in the partnership are unique. For example, if you have a cluster wholly within the domain `.france`, you can use hostnames like `paris` and `grenoble`. However, if you have a cross-domain cluster, you must specify the hostnames with enough qualification to identify the host on the network. You can link `paris` and `munich` with hostnames `paris.france` and `munich.germany`, and the cluster names remain `paris` and `munich`.

You cannot create a partnership between clusters `paris.france` and `paris.texas` because of a collision on the cluster name `paris`.

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, you might have to change the IP address and host names in the script. 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 define only one partnership between two specific clusters. A single cluster can participate in other partnerships with different clusters.

▼ How to Create a Partnership

Before You Begin 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 34.

- The cluster name of the partner cluster is known.
- The host information of the partner cluster must be defined in the local host file. The local cluster needs to know how to reach the partner cluster by name.
- Security has been configured on the two clusters by installing the appropriate certificates. See [“Configuring Trust Between Partner Clusters”](#) on page 51 for more information.

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 45.

2 Create the partnership.

```
# geops create -c remotepartnerclustername [-h heartbeatname] \
[-p propertysetting [-p...]] partnershipname
```

-c remotepartnerclustername

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 heartbeatname

Specifies a custom heartbeat to use in the partnership to monitor the availability of the partner cluster.

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 6, “Administering Heartbeats.”](#)

If you create a custom heartbeat, you must add at least one plug-in to prevent the partnership from remaining in degraded mode.

You must configure the custom heartbeat that you provide in this option before you run the `geops create` command.

Note – 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 you run the `geops create` command.

-p propertysetting

Specifies 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_emailaddr` and `Notification_actioncmd` properties. For more information about configuring heartbeat-loss notification, see “[Configuring Heartbeat-Loss Notification](#)” on page 75.

For more information about the properties you can set, see [Appendix A, “Standard Sun Cluster Geographic Edition Properties”](#).

partnershipname

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, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

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

3 Verify that the partnership was created and the status of the partnership.

```
# geoadm status
```

Example 5-1 Creating a Partnership

This example creates the `paris-newyork-ps` partnership on the `cluster-paris` cluster.

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

See Also For a complete example of how to configure and join a partnership, see [Example 5-4](#).

▼ How to Modify Partnership Properties

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 45.

2 Modify partnership properties.

```
# geops set-prop -p propertysetting [-p...] partnershipname
```

`-p propertysetting` Specifies 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_emailaddr`s and `Notification_actioncmd` properties. For more information about configuring heartbeat-loss notification, see “[Configuring Heartbeat-Loss Notification](#)” on page 75.

For more information about the properties you can set, see [Appendix A](#), “[Standard Sun Cluster Geographic Edition Properties](#)”.

partnershipname 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](#), “[Legal Names and Values of Sun Cluster Geographic Edition Entities](#)”.

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

This example modifies the notification email address for the `cluster-paris` cluster.

```
# geops set-prop -p Notification_emailaddr=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 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 47 for more information.

1 Log in to a node of the cluster that is joining 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 45.

2 Confirm that the remote cluster that originally created the partnership, `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 34.

3 Join the partnership.

```
# geops join-partnership [-h heartbeatname] remoteclustername partnershipname
```

`-h heartbeatname` Specifies a custom heartbeat to use in the partnership to monitor the availability of the partner cluster.

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 6, “Administering Heartbeats.”](#)

If you create a custom heartbeat, you must add at least one plug-in to prevent the partnership from remaining in degraded mode.

You must configure the custom heartbeat that you provide in this option before you run the `geops` command.

`remoteclustername` Specifies the name of a cluster that is currently a member of the partnership that is being joined. This cluster is used to retrieve the partnership configuration information.

`partnershipname` 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, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

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

This example joins the `cluster-newyork` cluster in the partnership that was created on `cluster-paris` in [Example 5-1](#).

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

Example 5-4 Creating and Joining a Partnership With a Remote Cluster in a Different Domain

This example creates and configures the `paris-newyork-ps` partnership between clusters `cluster-paris.france` and `cluster-newyork.usa`.

1. On one node of `cluster-paris.france`, configure trust for the partnership.

```
phys-paris-1# geops add-trust -c cluster-newyork.usa
```

2. On one node of `cluster-newyork.usa`, configure trust for the partnership.

```
phys-newyork-1# geops add-trust -c cluster-paris.france
```

3. On each node of both clusters, verify that trust has been set up properly, both between the local cluster and partner cluster and among nodes of the local cluster.

```
phys-newyork-1# geops verify-trust -c cluster-paris.france
phys-newyork-2# geops verify-trust -c cluster-paris.france
phys-newyork-1# geops verify-trust
phys-newyork-2# geops verify-trust
phys-paris-1# geops verify-trust -c cluster-newyork.usa
phys-paris-2# geops verify-trust -c cluster-newyork.usa
phys-paris-1# geops verify-trust
phys-paris-2# geops verify-trust
```

4. On `cluster-paris.france`, create the partnership `paris-newyork-ps`.

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

5. On `cluster-newyork.usa`, join the partnership `paris-newyork-ps`.

```
cluster-newyork# geops join-partnership cluster-paris.france paris-newyork-ps
```

6. Verify that the partnership has been created successfully.

```
# 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 the resources that are associated with the partnership.

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 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.

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 45](#).

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 protection groups, see one of the following data replication guides:

- “How to Delete a Sun StorageTek Availability Suite Protection Group” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*
- “How to Delete a Hitachi TrueCopy Protection Group” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- “How to Delete an EMC Symmetrix Remote Data Facility Protection Group” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

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

```
# geops leave-partnership partnershipname
```

partnershipname 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 5-5 Leaving a Partnership

In this example, the `cluster-paris` cluster leaves the `paris-newyork-ps` partnership.

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

Example 5-6 Deleting a Partnership

After the `cluster-paris` cluster leaves the `paris-newyork-ps` partnership, as described in the previous example, the only remaining member of the partnership is the `cluster-newyork` cluster. You can delete the `paris-newyork-ps` partnership by forcing the `cluster-newyork` cluster to leave the partnership.

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

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

Resynchronizing a Partnership

Partner clusters that become disconnected during a disaster situation might force the administrator to perform a takeover for a protection group that the partners share. When both clusters are brought online again, both partner clusters might 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 cluster.

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 do not need to resynchronize the configuration information in the following situations if the original secondary cluster goes down and resumes operation later.

Use the `geoadm status` command to check whether you need to resynchronize a partnership. If the Configuration status is Synchronization Status Error, you need to synchronize the partnership. If the Local status is Partnership Error, do not resynchronize the partnership. Instead, wait until a heartbeat exchange occurs.

▼ How to Resynchronize a Partnership

Before You Begin 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.

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 45.

2 Resynchronize the partnership.

```
# geops update partnershipname  
partnershipname Specifies the name of the partnership
```

Example 5-7 Resynchronizing a Partnership

This example resynchronizes a partnership.

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

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 contains the following sections:

- “Introduction to Heartbeats” on page 63
- “Creating a Heartbeat” on page 64
- “Creating a Heartbeat Plug-in” on page 65
- “Modifying a Heartbeat Plug-in Property” on page 66
- “Deleting Heartbeats and Heartbeat Plug-ins” on page 67
- “Displaying Heartbeat Configuration Information” on page 68
- “Tuning the Heartbeat Properties” on page 69
- “Creating a Heartbeat That Uses a Custom Heartbeat Plug-in” on page 71
- “Configuring Heartbeat-Loss Notification” on page 75

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.

Note – The Internet Assigned Numbers Authority (IANA) has officially assigned port number 2084 for use by the Sun Cluster Geographic Edition heartbeats.

- `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_localclustername-remoteclustername`. 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.

If you create a custom heartbeat, you must add at least one plug-in to prevent the partnership from remaining in degraded mode.

Creating a Heartbeat

This section describes procedures for creating heartbeats.

▼ How to Create a Heartbeat

Use this procedure to create a new heartbeat. To use 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.

If you create a custom heartbeat, you must add at least one plug-in to prevent the partnership from remaining in degraded mode.

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.

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 45](#).

2 Create the heartbeat.

```
# geohb create -r remoteclustername \  
[-p propertysetting [-p...]] heartbeatname
```

-r *remoteclustername* Specifies the name of the remote, secondary partner cluster.

-p *propertysetting* Specifies a heartbeat property that 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, “Standard Sun Cluster Geographic Edition Properties”](#).

heartbeatname Specifies an identifier for the heartbeat.

If you create a custom heartbeat, you must add at least one plug-in to prevent the partnership from remaining in degraded mode.



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 6-1 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

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 45](#).

2 Add the heartbeat plug-in to an existing heartbeat.

```
# geohb add-plugin heartbeatname pluginname \
[-p propertysetting [-p...]]
```

heartbeatname Specifies the identifier for heartbeat on the local cluster.

pluginname Specifies the name of the heartbeat plug-in.

-ppropertysetting Specifies a heartbeat plug-in property that 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, “Standard Sun Cluster Geographic Edition Properties”](#).

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

Example 6-2 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

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 45](#).

2 Modify the heartbeat plug-in properties.

```
# geohb modify-plugin -p propertysetting \
[-p...] pluginname heartbeatname
```

heartbeatname Specifies an identifier for the heartbeat.

pluginname Specifies the name of the heartbeat plug-in.

`-p propertysetting` Specifies a heartbeat plug-in property that 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, “Standard Sun Cluster Geographic Edition Properties”](#).

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, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

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

Example 6-3 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

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 45](#).

2 Delete the heartbeat.

```
# geohb delete heartbeatname
```

heartbeatname Specifies an identifier for the heartbeat settings.

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

Example 6-4 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

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 45](#).

2 Remove the plug-in from the heartbeat.

```
# geohb remove-plugin pluginname heartbeatname
```



Caution – Do not delete the default heartbeat plug-ins `tcp_upd_plugin` and `ping_plugin`.

pluginname Specifies the name of the custom heartbeat plug-in

heartbeatname Specifies an identifier for the heartbeat that contains this plug-in

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see [Appendix B, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

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

Example 6-5 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
```

Displaying Heartbeat Configuration Information

This section describes procedures for displaying heartbeat configuration information.

▼ How to Display Heartbeat Configuration Information

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 45](#).

2 Display the current configuration information for a specific heartbeat or the whole heartbeat subsystem.

```
# geohb list [heartbeatnamelist]
```

heartbeatnamelist Specifies the names of the heartbeats on the local cluster for which configuration information should be displayed.

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 6-6 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.geocontrol.xml` file.

If you adjust the delay setting of the `Query_interval` property, ensure that the following condition is met:

```
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 change the `heartbeat.retries` value. If you want to change the default value of the `heartbeat.retries` property, contact a Sun service representative.

▼ How to Modify the Heartbeat Properties

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 45](#).

2 Modify the heartbeat properties.

```
# geohb set-prop -p propertysetting \  
[-p...] heartbeatname
```

`-p propertysetting` Specifies 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, “Standard Sun Cluster Geographic Edition Properties”](#).

heartbeatname 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, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

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

Example 6-7 Modifying the Properties of the Default Heartbeat

This example modifies 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:

<i>queryinterval</i>	The value of the <code>Query-interval</code> property, which defines the delay in seconds after which a heartbeat status request is declared a failure.
<i>mode</i>	The mode for the plug-in startup, either <code>Normal</code> or <code>Emergency</code> .
<i>pluginpropertyvalues</i>	The value of the <code>Plugin-properties</code> property that is configured for the heartbeat plug-in, if any.

For more information about the properties you can set, see [Appendix A, “Standard Sun Cluster Geographic Edition Properties”](#).

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 — Indicates that the secondary cluster is alive
- Nonzero, on failure — Indicates that the secondary cluster did not respond to the heartbeat check

▼ How to Add a Custom Heartbeat Plug-in to an Existing Default Heartbeat

1 Log in to 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 45](#).

2 Add the custom heartbeat plug-in to the default heartbeat.

```
# geohb add-plugin -p propertysetting [-p...] \  
pluginname hb_localclustername-remotecclustername
```

-p propertysetting

Specifies 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, “Standard Sun Cluster Geographic Edition Properties”](#).

pluginname

Specifies the name of the custom heartbeat plug-in.

hb_localclustername-remotecclustername

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 a node of the secondary cluster.

Example 6-8 Adding a Custom Heartbeat Plug-in to the Default Heartbeat

This example adds 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

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 45](#).

2 Create the new custom heartbeat.

```
# geohb create -r remoteclassname \
[-p propertysetting [-p...]] heartbeatname
```

`-r remoteclassname` Specifies the name of the remote, secondary partner cluster.

`-p propertysetting` Specifies 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, “Standard Sun Cluster Geographic Edition Properties”](#).

heartbeatname 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 uniquely identifies the heartbeat, 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 propertysetting [-p...] \
pluginname heartbeatname
```

- p *propertysetting* Specifies 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, “Standard Sun Cluster Geographic Edition Properties”](#).
- pluginname* Specifies the name of the custom heartbeat plug-in.
- heartbeatname* Specifies an identifier for the heartbeat.

4 Create the partnership that will use the heartbeat that you created in the previous step.

```
# geops create -c remoteclustername -h heartbeatname \  
[-p propertysetting [-p...]] partnershipname
```

- c *remoteclustername* 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 *heartbeatname* Specifies the custom heartbeat to be used in the partnership to monitor the availability of the partner cluster.

- p *propertysetting* 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, “Standard Sun Cluster Geographic Edition Properties”](#).

- partnershipname* Specifies the name of the partnership.

For more information about using `geops create` command to create a partnership, see [“How to Create a Partnership” on page 54](#).

5 Verify that your changes were made correctly.

```
# geoadm status
```

Example 6–9 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 associates 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 Heartbeat-Loss Notification

You can configure the Sun Cluster Geographic Edition software to send email notification and to run an action script when a heartbeat is lost. You configure heartbeat-loss notification by using the optional `Notification_emailaddr`s 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 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)
```

Delays can occur between the generation of the heartbeat-loss event and the triggering of the heartbeat-loss 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 runs after a heartbeat-loss event.

Configuring the Heartbeat-Loss Notification Properties

You can configure heartbeat-loss notification by using two partnership properties, `Notification_emailaddr`s and `Notification_actioncmd`. You specify these properties by using the `geops` command.

You can specify these properties on the default heartbeat during partnership creation. For more information, see [“How to Create a Partnership” on page 54](#). You can also modify these properties by using the procedure that is described in [“How to Modify the Heartbeat Properties” on page 70](#).

If you want to be notified of heartbeat loss by email, set the `Notification_emailaddr` 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 to run a command in response to heartbeat loss, set the `Notification_actioncmd` property.

EXAMPLE 6-10 Configuring Heartbeat-Loss Notification for an Existing Partnership

This example specifies a notification email address and a custom notification script for the partnership, `paris-newyork-ps`.

```
phys-paris-1# geops set-prop \  
-p Notification_emailaddr=ops@paris.com,ops@newyork.com \  
-p Notification_actioncmd=/opt/hb_action.sh paris-newyork-ps
```

Creating an Action Shell Script for Heartbeat-Loss

You can create an action shell script that runs when the local cluster detects a heartbeat-loss in the partner cluster. The script runs with root permissions. The file must have root ownership and execution permissions, but the script should not have write permissions.

If you have configured the `Notification_actioncmd` property, the action command runs with arguments that provide information about the event in the following command line:

```
# customactioncommandpath -c localclustername -r remoteclustername -e 1 \  
-n nodename -t time
```

<code>customactioncommandpath</code>	Specifies a path to the action command you have created.
<code>-c localclustername</code>	Specifies the name of the local cluster.
<code>-p remoteclustername</code>	Specifies the name of the remote partner cluster.
<code>-e 1</code>	Specifies that <code>HBLOST=1</code> , which indicates that a heartbeat-loss event has occurred. The Sun Cluster Geographic Edition software only supports heartbeat-loss notification, so <code>-e 1</code> is the only value that can be passed to the action shell script.

<code>-n <i>nodename</i></code>	Specifies the name of the cluster node that sent the heartbeat-loss event notification.
<code>-t <i>timestamp</i></code>	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 6-11 How a Notification Action Script Parses the Command-Line Information Provided by the Sun Cluster Geographic Edition Software

This example displays 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
```


Administering Protection Groups

This chapter contains the procedures for creating and configuring protection groups that do not require data replication. The chapter contains the following sections:

- “Introduction to Protection Groups” on page 79
- “Creating a Protection Group That Does Not Require Data Replication” on page 80

Introduction to Protection Groups

Protection groups enable a set of clusters to tolerate and recover from disaster by managing the resource groups for services. Protection groups can exist only in a partnership. You must create a partnership before you can create a protection group for that partnership. A protection group contains application resource groups and properties for managing data replication for those application resource groups.

You can duplicate the application resource group configuration on partner clusters. The configuration for a protection group is identical on partner clusters, so partner clusters must have the application resource groups of the protection group defined in their configuration. The Sun Cluster Geographic Edition software propagates protection group configurations between partners.

You can specify a data replication type in the protection group to indicate the mechanism that is used for data replication between partner clusters. When a service is protected from disaster by data replication, the protection group also contains replication resource groups. Protection groups link an application in a resource group with the application data that should be replicated. This linkage and replication enable the application to fail over seamlessly from one cluster to another cluster.

For information about how to create a protection group that requires data replication, see the following data replication guides:

- Chapter 2, “Administering Sun StorageTek Availability Suite Protection Groups,” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*

- Chapter 2, “Administering Hitachi TrueCopy Protection Groups,” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- Chapter 2, “Administering EMC Symmetrix Remote Data Facility Protection Groups,” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

Creating a Protection Group That Does Not Require Data Replication

Some protection groups do not require data replication. If you are using the Sun Cluster Geographic Edition software to manage only resource groups, 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 not to use data replication.

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 Not to 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.

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 45.

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 partnershipname -o localrole \  
[-p property [-p...]] \  
protectiongroupname
```


- s *partnershipname* Specifies the name of the partnership.
- o *localrole* Specifies the role of this protection group on the local cluster as either Primary or Secondary.
- p *propertysetting* Specifies the properties of the protection group.

You can specify the following properties:

- `Description` – Describes the protection group.
- `External_Dependency_Allowed` – Specifies whether to allow any dependencies between resource groups and resources that belong to this protection group and resource groups and resources that do not belong to this protection group.
- `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 path to an executable command. This script is invoked during a switchover or takeover on the new primary cluster when the protection group is started on the new primary cluster. The script is invoked on the new primary cluster after the data replication role changes from secondary to primary and before the application resource groups are brought online. If the data replication role change does not succeed, then the script is not called.

This path should be valid on all nodes of all partner clusters that can host the protection group.

- `Timeout` – Specifies the timeout period for the protection group in seconds. You can change the timeout period from the default value depending on the complexity of your data replication configuration. For more information on setting the timeout period, see [Table A-4](#).

For more information about the properties you can set, see [Appendix A, “Standard Sun Cluster Geographic Edition Properties”](#).

- protectiongroupname* Specifies the name of the protection group.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see [Appendix B, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

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

Example 7-1 Creating and Configuring a Protection Group That Is Configured to Not Use Data Replication

This example creates a protection group that is configured to not use data replication.

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

Next Steps See one of the following guides for information about adding resource groups to a protection group.

- “Administering Sun StorageTek Availability Suite Application Resource Groups” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*
- “Administering Hitachi TrueCopy Application Resource Groups” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- “Administering EMC Symmetrix Remote Data Facility Application Resource Groups” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

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 contains the following sections:

- [“Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software” on page 83](#)
- [“Viewing the Sun Cluster Geographic Edition Log Messages” on page 89](#)
- [“Displaying Configuration Information for Partnerships and Protection Groups” on page 89](#)

Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software

You can display 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 – Provides the name of the local cluster
- Partnership – Provides information all partnership, including the name of the partner cluster, the synchronization state, the local heartbeats, and the local heartbeat plug-in
- Protection group – Provides information about the status of protection groups, including information about the local cluster and the remote cluster
- Pending operations – Provides status information about any 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 45](#).

For example, an administrator runs the `geoadm status` command on `cluster-paris` and the following information is displayed:

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

```
Cluster: cluster-paris
```

```
Partnership "paris-newyork-ps": OK
```

```
  Partner clusters      : cluster-newyork
```

```
  Synchronization      : OK
```

```
  ICRM Connection      : OK
```

```
Heartbeat "paris-to-newyork" monitoring "cluster-newyork": OK
```

```
  Heartbeat plug-in "ping_plugin"      : Inactive
```

```
  Heartbeat plug-in "tcp_udp_plugin"    : OK
```

```
Protection group "tcpg"      : 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 displayed 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 8-1 Status Value Descriptions

Field	Value Descriptions
Partnership	<p>OK – The partners are connected.</p> <p>Error – The connection between the partner clusters is lost.</p> <p>Degraded – The partnership has been successfully created but a connection with the partner cluster has not yet been established. This status value occurs when the partnership has been created and the partner cluster has not been configured.</p>
Synchronization	<p>OK – The configuration information is synchronized between partner clusters.</p> <p>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.</p> <p>For information about resynchronizing a partnership, see “Resynchronizing a Partnership” on page 61.</p> <p>For information about resynchronizing a protection group, see one of the following data replication guides:</p> <ul style="list-style-type: none"> ■ “Resynchronizing a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “Resynchronizing a Hitachi TrueCopy Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “Resynchronizing an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> <p>Mismatch – Configuration information has been created separately on the clusters. The configuration information must be replaced by a copy of the configuration information from the partner cluster. You can synchronize the protection group configuration by using the <code>geopg get</code> command.</p> <p>Unknown – Information is not accessible because the partners are disconnected or because some components of the protection group cannot be reached.</p>
ICRM Connection	<p>OK – The Intercluster Resource Management (ICRM) module is running properly.</p> <p>Error – The ICRM module on the local cluster is unable to communicate with the ICRM module on the remote cluster.</p>

TABLE 8-1 Status Value Descriptions (Continued)

Field	Value Descriptions
Heartbeat	<p>OK – Heartbeat checks are running and the partner cluster responds within the specified timeout and retry periods.</p> <p>Offline – Heartbeat checks are not running.</p> <p>Error – Heartbeat checks are running but the partner is not responding and retries have timed out.</p> <p>Degraded – Heartbeat checks are running but one of the primary plug-ins is degraded or not running.</p>
Heartbeat plug-in	<p>OK – Responses are being received from the partner.</p> <p>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.</p> <p>No-Response – Partner cluster is not responding.</p>
Protection group (overall protection group state)	<p>OK – The synchronization state is OK and the state of the protection group on each cluster is OK.</p> <p>Degraded – The synchronization state is OK. The state of the protection group is Degraded on either one or both clusters in the partnership.</p> <p>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.</p> <p>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.</p>
Protection group > Cluster (state of protection group on each cluster)	<p>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.</p> <p>Degraded – The state of one or more of the protection group components is in the Degraded state on the cluster.</p> <p>Unknown – The state of some components of the protection group, such as configuration data, data replication, or resource groups, is unavailable.</p> <p>Error – The state of some components of the protection group, such as configuration data, data replication, or resource groups, is in Error.</p>
Protection group > Cluster > Role	<p>Primary – The cluster is the Primary for this protection group.</p> <p>Secondary – The cluster is the Secondary for this protection group.</p> <p>Unknown – Information is not accessible because the partners are disconnected or because some components of the protection group cannot be reached.</p>

TABLE 8-1 Status Value Descriptions (Continued)

Field	Value Descriptions
Protection group > Cluster > PG activation state	<p>Activated – The protection group is activated.</p> <p>Deactivated – The protection group is deactivated.</p> <p>Unknown – Information is not accessible because the partners are disconnected or because some components of the protection group cannot be reached.</p>
Protection group > Cluster > Configuration	<p>OK – Protection group configuration has been validated without errors on the cluster.</p> <p>Error – Protection group configuration validation resulted in errors on the cluster. You need to revalidate the protection group. For information about validating a protection group, see one of the following data replication guides:</p> <ul style="list-style-type: none"> ■ “How to Validate a Sun StorageTek Availability Suite Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> ■ “Validating a Hitachi TrueCopy Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> ■ “Validating an EMC Symmetrix Remote Data Facility Protection Group” in <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> <p>Unknown – Information is not accessible because the partners are disconnected or because some components of the protection group cannot be reached.</p>

TABLE 8-1 Status Value Descriptions (Continued)

Field	Value Descriptions
Protection group > Cluster > Data replication	<p>None – Data replication is not configured.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>Unknown – Information is not accessible because the partners are disconnected or because some components of the protection group cannot be reached.</p> <p>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 StorageTek Availability Suite data replication.</p>
Protection group > Cluster > Resource groups	<p>None – No resource group is protected by this protection group.</p> <p>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 has the Secondary role, all resource groups are unmanaged.</p> <p>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 has the Secondary role, not all resource groups are unmanaged.</p> <p>Unknown – Information is not accessible because the partners are disconnected or because some components of the protection group cannot be reached.</p>

For more specific information about checking the runtime status of replication, see one of the following data replication guides:

- “Checking the Runtime Status of Sun StorageTek Availability Suite Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite*
- “Checking the Runtime Status of Hitachi TrueCopy Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- “Checking the Runtime Status of EMC Symmetrix Remote Data Facility Data Replication” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

Viewing the Sun Cluster Geographic Edition Log Messages

All 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 in 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/cacao/instances/default/logs/cacao.0`
- `/var/cacao/instances/default/logs/cacao.1`
- `/var/cacao/instances/default/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 log messages of the Sun Cluster Geographic Edition modules 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.

Displaying Configuration Information for Partnerships and Protection Groups

You can display the current local cluster partnership configuration, including a list of all partnerships that are defined between the local cluster and remote clusters.

You can also display 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

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 45](#).

2 Display information about the partnership.

```
# geops list partnershipname
```

partnershipname Specifies the name of the partnership. If you do not specify a partnership, then the `geops list` command displays information on all partnerships.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see [Appendix B, “Legal Names and Values of Sun Cluster Geographic Edition Entities”](#).

Example 8-1 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

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 45](#).

2 Display information about a protection group.

```
# geopg list [protectiongroupname]
```

protectiongroupname Specifies the name of a protection group.

If you do not specify a protection group, then the command lists information about all the protection groups that are configured on your system.

Example 8-2 Displaying Configuration Information About a Protection Group

This example displays configuration information for avspg, which is configured on cluster-paris.

```
# geopg list avspg
```


Customizing Switchover and Takeover Actions

This chapter describes how to create a script that runs when the role of a protection group changes from secondary to primary. The chapter contains the following sections:

- “Creating a Role-Change Action Script” on page 93
- “Configuring a Protection Group to Run a Script at Switchover or Takeover” on page 95

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 runs during a switchover or takeover on the new primary cluster when the protection group is started on the new primary cluster. The script is invoked on the new primary cluster after the data replication role changes from secondary to primary and before the application resource groups are brought online. If the data replication role change does not succeed, then the script is not called.

The path to this script should be valid on all nodes of all partner clusters that can host the protection group.

The following command-line runs the script:

```
# custom-action-command-path -o primary -c clustername \  
-s partnershipname protectiongroupname userarguments
```

<i>customactioncommandpath</i>	Specifies a path to the action command you have created.
<i>-o primary</i>	Specifies that the role being assumed by the cluster is primary.
<i>-c clustername</i>	Specifies the name of the secondary cluster that is assuming the new role of primary cluster.

<i>-s partnershipname</i>	Specifies the name of the partnership that hosts the protection group.
<i>protectiongroupname</i>	Specifies the name of the protection group that is undergoing the role change.
<i>userarguments</i>	Specifies static arguments that are passed after all 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 <code>key=value</code> pairs, such as <code>name=sun.com, ip=10.1.2.3</code> . You could also specify a sequence of options, such as <code>-n sun.com -a 10.1.2.3.4</code> . 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 started 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.

The Sun Cluster Geographic Edition software waits for the script to return before the software processes operations such as bringing online application resource groups. Therefore, you must know in advance the amount of time required to run the script when you create the action script so that you can set the timeout period for the protection group accordingly. Setting the timeout period to include enough time for the script to complete to avoid switchovers or takeovers timing out and leaving the application resource group offline on the new primary.

EXAMPLE 9-1 Switchover Action Script for Updating the DNS

This sample script 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
# sample 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"
```

EXAMPLE 9-1 Switchover Action Script for Updating the DNS (Continued)

```

# 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"
#
# 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 Run a Script at Switchover or Takeover

After you have created a script, you must configure the protection group to run the script when a switchover or takeover occurs. If a switchover or takeover occurs, the script runs on the cluster that is becoming the new primary cluster.

▼ How to Configure a Protection Group to Run a Script at Switchover or Takeover

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 45](#).

2 Configure the `RoleChange_ActionCmd` and `RoleChange_ActionArgs` properties of the protection group.

```
# geopg set-prop -p RoleChange_ActionCmd=fullyqualifiedscript -p RoleChange_ActionArgs=scriptarguments
```

`-p propertysetting` Specifies 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, “Standard Sun Cluster Geographic Edition Properties”](#).

`protectiongroupname` Specifies the name of the protection group.

Example 9-2 Configuring a Protection Group to Run a Command at Cluster Switchover or Takeover

This example configures a protection group to run 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
```


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 sections:

- “General Heartbeat Properties” on page 97
- “General Heartbeat Plug-in Properties” on page 98
- “Partnership Properties” on page 99
- “General Properties of a Protection Group” on page 100

Note – The property names and values, such as `Query_interval`, `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
<code>Query_interval</code> (integer)	<p>Specifies the delay in seconds between heartbeat status requests.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime.</p> <p>Category: Optional</p> <p>Default: 120 seconds</p>

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)	<p>Specifies a property string specific to the plug-in.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime.</p> <p>Category: Optional</p> <p>Default: None except for heartbeats that use the default heartbeat plug-ins, <code>tcp_udp_plugin</code> and <code>ping_plugin</code>.</p> <p>For the <code>tcp_udp_plugin</code> plug-in, the format of this string is predefined as <code>remoteIPaddress/UDP/2084/ipsec</code>, <code>remoteIPaddress/TCP/2084/ipsec</code>. The <code>remote_IP_address</code> argument specifies the IP address of the partner cluster. The optional <code>ipsec</code> argument specifies if the plug-in uses IPsec with a Boolean value of <code>true</code> or <code>false</code>.</p> <p>For the <code>ping_plugin</code>, the format of this string is predefined as <code>remote_IP_address</code>, where <code>remote_IP_address</code> specifies the IP address of the partner cluster.</p>
Query_cmd (string)	<p>Specifies the path to the heartbeat status request command.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime.</p> <p>Category: Required property if the plug-in does not specify a predefined plug-in.</p> <p>Default: None</p>
Requester_agent (string)	<p>Specifies the absolute path to the requester agent.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. However, the <code>Requester_agent</code> property of the default plug-in should never need to be tuned except for testing purposes.</p> <p>Category: Optional</p> <p>Default: None</p>

TABLE A-2 General Heartbeat Plug-in Properties (Continued)

Property	Description
Responder_agent (string)	<p>Specifies the absolute path to the responder agent.</p> <p>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.</p> <p>Category: Optional</p> <p>Default: None</p>
Type (enum)	<p>Designates the type of plug-in. Set to either primary or backup.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime.</p> <p>Category: Required</p> <p>Default: None, except for the default heartbeat that is named ping_plugin. If using this plug-in, the default value is backup.</p>

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)	<p>Describes the partnership.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime.</p> <p>Category: Optional</p> <p>Default: Empty string</p>
Notification_ActionCmd (string)	<p>Provides the path to the action script that is triggered when heartbeat-loss notification is issued.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime.</p> <p>Category: Optional</p> <p>Default: Empty string</p>

TABLE A-3 Partnership Properties (Continued)

Property	Description
Notification_EmailAdrrs (string array)	<p>Lists the email addresses that are sent email when heartbeat-loss notification is issued. The list is comma delimited.</p> <p>Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime.</p> <p>Category: Optional</p> <p>Default: Empty string</p>

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)	<p>Describes the protection group.</p> <p>Tuning recommendations: This property can be tuned at any time.</p> <p>Category: Optional</p> <p>Default: Empty string</p>
	<p>Allow dependencies between resource groups and resources that belong to this protection group and resource groups and resources that do not belong to this protection group when set to true.</p> <p>Tuning recommendations: This property can be tuned at any time.</p> <p>Category: Optional</p> <p>Default: false</p>
RoleChange_ActionArgs (string)	<p>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.</p> <p>Tuning recommendations: This property can be tuned at any time.</p> <p>Category: Optional</p> <p>Default: Empty string</p>

TABLE A-4 General Properties of a Protection Group (Continued)

Property	Description
RoleChange_ActionCmd (string)	<p data-bbox="684 230 1328 430">Specifies the path to an executable command. This script is invoked during a switchover or takeover on the new primary cluster when the protection group is started on the new primary cluster. The script is invoked on the new primary cluster after the data replication role changes from secondary to primary and before the application resource groups are brought online. If the data replication role change does not succeed, then the script is not called.</p> <p data-bbox="684 447 1328 499">This path should be valid on all nodes of all partner clusters that can host the protection group.</p> <p data-bbox="684 517 1278 543">Tuning recommendations: This property can be tuned at any time.</p> <p data-bbox="684 560 856 586">Category: Optional</p> <p data-bbox="684 604 878 630">Default: Empty string</p>

TABLE A-4 General Properties of a Protection Group (Continued)

Property	Description
<p>Timeout (integer)</p>	<p>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 you run a <code>geogg</code> command, such as <code>geogg start</code>, <code>geogg stop</code>, <code>geogg switchover</code>, and <code>geogg takeover</code>. If the command does not respond within the timeout period, the Sun Cluster Geographic Edition software reports the operation as timed out, even if the underlying command eventually completes successfully.</p> <p>You should identify the amount of time required to perform a role-reversal of the data replication, and set the timeout value to 150% to 200% of that value to ensure enough time for the role-reversal to complete.</p> <p>To ensure that an operation has finished on the remote cluster, check system status after a timeout before attempting the operation again. For more information, see “Troubleshooting Migration Problems” on page 110.</p> <p>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.</p> <p>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 that an operation with a global scope times out during one of the following conditions:</p> <ul style="list-style-type: none"> ▪ The local operation does not complete after the specified timeout period. ▪ The remote operation does not complete after the specified timeout period. <p>Tuning recommendations: This property can be tuned only when the protection group is offline.</p> <p>Category: Optional</p> <p>Range: 20-1000000 seconds</p> <p>Default: 200</p>

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 sections:

- [“Legal Names for Sun Cluster Geographic Edition Entities” on page 103](#)
- [“Legal Values for Sun Cluster Geographic Edition Entities” on page 104](#)

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 Sun Cluster Geographic Edition entity values fall into two categories: property values and description values. Both types of values share the following rules:

- 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 a newline or a semicolon

Takeover Postconditions

This appendix provides details about the state of the primary and secondary clusters after you run the `geopg takeover` command.

This appendix contains the following sections:

- [“Results of a Takeover When the Partner Cluster Can Be Reached” on page 105](#)
- [“Results of a Takeover When the Partner Cluster Cannot Be Reached” on page 106](#)

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 you run the `geopg takeover` command. The results described in this section assume that the partner cluster can be reached.

The following table describes the states of the clusters when you run the `geopg takeover` command on the secondary cluster, `cluster-newyork`.

TABLE C-1 Takeover Results of Running the `geopg takeover` Command on the Secondary Cluster

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
<code>cluster-paris: primary, deactivated</code> <code>cluster-newyork: secondary, deactivated</code>	<code>cluster-paris: secondary, deactivated</code> <code>cluster-newyork: primary, deactivated</code>
<code>cluster-paris: primary, activated</code> <code>cluster-newyork: secondary, deactivated</code>	<code>cluster-paris: secondary, deactivated</code> <code>cluster-newyork: primary, deactivated</code>

TABLE C-1 Takeover Results of Running the `geopg takeover` Command on the Secondary Cluster
(Continued)

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
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 you run the `geopg takeover` command on the primary cluster, `cluster-paris`.

TABLE C-2 Takeover Results of Running the `geopg takeover` Command 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 you run a `geopg takeover` command 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 you run the `geopg takeover` command 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 available only when the partner cluster can be reached again.

TABLE C-3 Takeover Results of Running the `geopg takeover` Command on the Secondary Cluster When 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 you run the `geopg takeover` command on the primary cluster, `cluster-paris`, and the secondary cluster cannot be reached or the protection group on the secondary cluster is busy.

TABLE C-4 Takeover Results of Running the `geopg takeover` Command on the Primary Cluster When 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

TABLE C-4 Takeover Results of Running the `geopg takeover` Command on the Primary Cluster When the Secondary Cluster Cannot Be Reached (Continued)

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
<p><code>cluster-paris</code>: primary, activated, synchronization status Unknown</p> <p><code>cluster-newyork</code>: secondary, deactivated, synchronization status Unknown</p>	<p><code>cluster-paris</code>: primary, activated, with data replication stopped, synchronization status OK, Error, or Mismatch</p> <p><code>cluster-newyork</code>: secondary, deactivated, synchronization status OK, Error, or Mismatch</p>
<p><code>cluster-paris</code>: primary, deactivated, synchronization status Unknown</p> <p><code>cluster-newyork</code>: secondary, activated, synchronization status Unknown</p>	<p><code>cluster-paris</code>: primary, deactivated, synchronization status OK, Error, or Mismatch</p> <p><code>cluster-newyork</code>: secondary, activated, synchronization status OK, Error, or Mismatch</p>
<p><code>cluster-paris</code>: primary, activated, synchronization status Unknown</p> <p><code>cluster-newyork</code>: secondary, activated, synchronization status Unknown</p>	<p><code>cluster-paris</code>: primary, activated, with data replication stopped, synchronization status OK, Error, or Mismatch</p> <p><code>cluster-newyork</code>: secondary, activated, synchronization status OK, Error, or Mismatch</p>

Troubleshooting Sun Cluster Geographic Edition Software

This appendix describes procedures for troubleshooting your application of the Sun Cluster Geographic Edition software.

This appendix contains the following sections:

- “[Troubleshooting Monitoring and Logging](#)” on page 109
- “[Troubleshooting Migration Problems](#)” on page 110
- “[Troubleshooting Cluster Start and Restart](#)” on page 111

Troubleshooting Monitoring and Logging

This section provides information about setting up logging and problems that you might encounter with monitoring the Sun Cluster Geographic Edition software.

Configuring the Logger File to Avoid Too Many Traces

Configure the logger file, `/etc/cacao/instances/default/private/logger.properties`, as following depending on the `cmass` messages you want logged:

- To select only `WARNING` and `SEVERE` messages, the first line of the file should read as follows:

```
com.sun.cluster.level=WARNING
```

- To enable all `geocontrol` messages, the second line of the file should read as follows:

```
com.sun.cluster.agent.geocontrol.level=ALL
```

The enabled traces are copied to the `/var/cacao/instances/default/logs/cacao.0` file.

Configuring the Log File to Avoid Detailed Messages From the gcr Agent

If you want to avoid too detailed messages in your log file from the gcr agent, use entries similar to the following in your log file

/etc/cacao/instances/default/private/logger.properties:

```
com.sun.cluster.level=WARNING
com.sun.cluster.agent.geocontrol.gcr.level=INFO
com.sun.cluster.agent.geocontrol.level=ALL
```

This property file is updated each time you reinstall the SUNWs cmsa package.

Configuring the Log File to Avoid jmx Remote Traces

To avoid jmx remote traces add the following lines to the beginning of your logger.properties file:

```
javax.management.remote.level=OFF
com.sun.jmx.remote.level=OFF
java.io.level=OFF
```

Troubleshooting Migration Problems

This section provides information about problems that you might encounter when services are migrated by using Sun Cluster Geographic Edition software.

Resolving Problems With Application Resource Group Failover When Communication Lost With the Storage Device

When a loss of communication occurs between a node on which the application is online and the storage device, some application resource groups might not failover gracefully to the nodes from which the storage is accessible. The application resource group might result in a `ERROR_STOP_FAILED` state.

Solution or Workaround

The Sun Cluster infrastructure does not initiate a switchover when I/O errors occur in a volume or its underlying devices. Because no switchover or failover occurs, the device service remains online on this node despite the fact that storage has been rendered inaccessible.

If this problem occurs, restart the application resource group on the correct nodes by using the standard Sun Cluster procedures. Refer to “Clearing the STOP_FAILED Error Flag on Resources” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* about recovering from the ERROR_STOP_FAILED state and restarting the application.

The Sun Cluster Geographic Edition software detects state changes in the application resource group and displays the states in the output of the `geoadm status` command. For more information about using this command, see [“Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software”](#) on page 83.

Troubleshooting Cluster Start and Restart

This section provides information about troubleshooting problems that you might encounter with starting and restarting the Sun Cluster Geographic Edition software.

Validating Protection Groups in an Error State

After a cluster reboot the protection group configuration might be in an error state. This problem might be caused by the common agent container process not being available on one of the nodes of the cluster when the protection group is initialized after the reboot.

Solution or Workaround

To fix the configuration error, use the `geopg validate` command on the protection group that is in an error state.

Restarting the Common Agent Container

The Sun Cluster software enables the common agent container only during the Sun Cluster software installation. Therefore, if you disable the common agent container at any time after the installation, the common agent container remains disabled.

Solution or Workaround

To enable the common agent container after a node reboot, use the `/usr/lib/cacao/bin/cacaoadm enable` command.

Index

A

- activating Sun Cluster Geographic Edition
 - software, 34-36
- adding, resource groups to protection groups, 43
- administering
 - access, 45-49
 - heartbeats, 63-77
 - protection groups, 79-82
 - security, 45-49
- administration tasks, 26-31
 - prerequisite, 26-27
 - Sun Cluster, 23-25
 - Sun Cluster Geographic Edition, 27-31
- application suitability, 21-22

B

- booting cluster, 40

C

- certificates, configuring, 47
- cluster
 - administration concepts, 23-25
 - booting, 40
 - example configuration, 31
 - status of, 83-89
- command-line interface, overview of, 19
- common agent container
 - required version, 42

common agent container (*Continued*)

- starting, 43
- stopping, 41
- configuring
 - IPsec, 47-49
 - logical hostname, 24-25
 - protection groups
 - unreplicated, 80-82
 - RBAC, 45-46
 - role-change action script, 95-96
 - security certificates, 47
- creating
 - heartbeats, 64-65
 - partnerships, 54-56
 - protection groups
 - unreplicated, 80-82
 - role-change action script, 93-95
- custom heartbeat action script, 76-77
- custom heartbeat plug-in
 - adding to custom heartbeat, 73-75
 - adding to default heartbeat, 72-73
 - creating heartbeat for, 71-75

D

- deleting
 - heartbeats, 67-68
 - partnerships, 60-61
 - plug-in from heartbeat, 68
- device groups, overview, 25

- disabling Sun Cluster Geographic Edition
 - software, 37-39
- disaster recovery overview, 19-20
- displaying
 - heartbeat configuration, 69
 - partnership configuration, 89-91

E

- enabling Sun Cluster Geographic Edition
 - software, 34-36
 - after adding patches, 43
 - /etc/inet/ipsecinit.conf, 48-49
 - /etc/init/secret/ipseckeys, 48-49
- example cluster configuration, 31

G

- geo-clustername, 33-34
- geo-clusterstate, 33-34
- geo-failovercontrol, 33-34
- geo-hbmonitor, 33-34
- geo-infrastructure, 33-34
- geoadm command, enabling Sun Cluster Geographic Edition software, 43
- geoadm show, 39-40
- geoadm status, 83-89
- geopg command
 - adding resource groups to protection groups, 43
 - removing resource groups from protection groups, 41
 - starting protection groups, 43
 - stopping protection groups, 41
- graphical user interface (GUI), overview of, 18

H

- heartbeat-loss notification, 75-77
- heartbeat plug-in
 - deleting from a heartbeat, 68
 - modifying properties of, 66-67

- heartbeats
 - administering, 63-77
 - creating, 64-65
 - custom action script, 76-77
 - deleting, 67-68
 - deleting plug-in from, 68
 - displaying configuration of, 69
 - general heartbeat plug-in properties, 98-99
 - general properties of, 97-98
 - introduction to, 63-64
 - IPsec security with, 47-49
 - loss notification, 75-77
 - tuning the properties of, 69-71

I

- IPsec, 47-49
 - keys file, 48-49
 - policy file, 48-49

J

- joining, partnerships, 57-60

L

- leaving, partnerships, 60-61
- logging, 89
- logical hostname, configuring, 24-25
- loss of heartbeat notification, 75-77
 - creating action shell script, 76-77
 - properties, 75-76

M

- modifying
 - heartbeat plug-in properties, 66-67
 - heartbeat properties, 70-71
 - partnerships, 56-57
 - RBAC rights, 46-47
- monitoring, infrastructure resource groups, 33-34

monitoring Sun Cluster Geographic Edition, 83-89

N

notification_actioncmd, 75-77
notification_emailaddr, 75-77

O

operations, status of, 83-89

P

partnerships
 creating, 54-56
 deleting, 60-61
 displaying configuration information, 89-91
 joining, 57-60
 leaving, 60-61
 modifying, 56-57
 properties of, 99-100
 resynchronizing, 61-62
 status of, 83-89
patches, applying, 40-43
properties
 general heartbeat, 97-98
 general heartbeat plug-in, 98-99
 general protection group, 100-102
 partnership, 99-100
 tuning heartbeat, 69-71
protection groups
 adding resource groups, 43
 administering, 79-82
 configuring
 role-change action, 95-96
 general properties of, 100-102
 removing resource groups, 41
 starting, 43
 status of, 83-89
 stopping, 41
 unreplicated
 creating, 80-82

R

RBAC, 45-47
 modifying rights, 46-47
 rights profiles, 46
 setting up and using, 45-46
resource groups
 configuring, 23
 removing from protection groups, 41
 Sun Cluster Geographic Edition
 infrastructure, 33-34
resources, configuring, 23
resynchronizing, partnerships, 61-62
role-based access control, *See* RBAC
role-change action script, 93-96
 configuring protection group for, 95-96
 creating, 93-95
runtime status, Sun Cluster Geographic Edition, 83-89

S

scripts
 custom loss of heartbeat action, 76-77
 switchover and takeover action, 93-96
security
 administering, 45-49
 configuring certificates, 47
 IPsec, 47-49
solaris.cluster.geo.admin, 46
solaris.cluster.geo.modify, 46
solaris.cluster.geo.read, 46
starting
 See also enabling
 common agent container, 43
 protection groups, 43
status, Sun Cluster Geographic Edition, 39-40
status descriptions, 83-89
stopping
 common agent container, 41
 Sun Cluster Geographic Edition software, 41
Sun Cluster
 administration concepts, 23-25
 resources, 23
Sun Cluster Geographic Edition software
 disabling, 37-39

Sun Cluster Geographic Edition software (*Continued*)

enabling, 34-36

enabling after adding patches, 43

stopping, 41

Sun Cluster Manager, 18

Sun StorageTek Availability Suite, IPsec, 47-49

switchover, custom action script, 93-96

T

takeover, custom action script, 93-96

timeout, description of, 100-102

tuning, heartbeat properties, 69-71