

Oracle® Solaris Cluster Geographic Edition System Administration Guide

ORACLE®

Part No: E39667
July 2014, E39667-01

Copyright © 2004, 2014, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2004, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Using This Documentation	17
1 Introduction to Administering the Geographic Edition Software	19
Geographic Edition Administration	19
Geographic Edition Administration Tools	20
Graphical User Interface	20
Command-Line Interface	21
2 Overview of Administering a Geographic Edition Configuration	23
Overview of Oracle Solaris Cluster Administration Concepts	23
Configuring Resources and Resource Groups	23
Configuring Logical Hostnames	24
Managing Device Groups	25
Geographic Edition Configuration and Administration Tasks	25
Example Geographic Edition Cluster Configuration	26
3 Administering the Geographic Edition Infrastructure	29
Geographic Edition Infrastructure Resource Groups	29
Enabling the Geographic Edition Software	30
▼ How to Enable Geographic Edition Software	30
Disabling the Geographic Edition Software	33
▼ How to Disable the Geographic Edition Software	33
Checking the Status of the Geographic Edition Infrastructure	35
Booting a Cluster	36
4 Administering RBAC	37
Modifying a User's RBAC Properties	37
5 Administering Cluster Partnerships	39

Modifying Partnership Properties	39
▼ How to Modify Partnership Properties	39
Adding a New Cluster Node	40
▼ How to Add a New Node to a Cluster in a Partnership	41
Renaming a Cluster Node	42
Renaming a Cluster That Is in a Partnership	42
▼ How to Rename a Cluster That Is in a Partnership	42
Removing Trust Between Partner Clusters	46
▼ How to Remove Trust Between Two Clusters	46
Leaving or Deleting a Partnership	48
▼ How to Leave or Delete a Partnership	48
Resynchronizing a Partnership	49
▼ How to Resynchronize a Partnership	50
6 Administering Heartbeats	51
Introduction to Heartbeats	51
Creating a Heartbeat	52
▼ How to Create a Heartbeat	52
Creating a Heartbeat Plug-in	53
▼ How to Create Heartbeat Plug-in	54
Modifying a Heartbeat Plug-in Property	54
▼ How to Modify the Properties of a Heartbeat Plug-in	55
Deleting Heartbeats and Heartbeat Plug-ins	56
▼ How to Delete a Heartbeat	56
▼ How to Delete a Plug-in From a Heartbeat	56
Displaying Heartbeat Configuration Information	57
▼ How to Display Heartbeat Configuration Information	57
Tuning the Heartbeat Properties	58
▼ How to Modify the Heartbeat Properties	59
Creating a Heartbeat That Uses a Custom Heartbeat Plug-in	60
Creating a Custom Heartbeat Plug-in	60
▼ How to Add a Custom Heartbeat Plug-in to an Existing Default Heartbeat	61
▼ How to Create a Custom Heartbeat Plug-in and Add It to a Custom Heartbeat	62
Configuring Heartbeat-Loss Notification	64
Introduction to Configuring Heartbeat-Loss Notification	64
Configuring the Heartbeat-Loss Notification Properties	65
Creating an Action Shell Script for Heartbeat-Loss	65

7 Administering Protection Groups	69
Activating and Deactivating a Protection Group	69
Guidelines for Activating and Deactivating a Protection Group	69
▼ How to Activate a Protection Group	71
▼ How to Deactivate a Protection Group	75
Resynchronizing a Protection Group	78
▼ How to Resynchronize a Protection Group	78
Modifying Protection Groups and Data Replication Components	79
▼ How to Modify a Protection Group	79
▼ How to Modify a Data Replication Component	80
Deleting Protection Groups and Data Replication Components	81
▼ How to Delete a Protection Group	81
▼ How to Remove a Data Replication Component From a Protection Group	83
8 Administering Sites	85
Adding a Cluster to a Site	85
▼ How to Add a Cluster to a Site	85
▼ How to Accept an Invitation to Join a Site	87
Changing the Site Role of a Site Cluster	89
▼ How to Change a Site Membership Role	89
▼ How to Forcibly Change a Site Member to a Site Controller	90
Synchronizing Site Configuration Information	92
▼ How to Synchronize Site Configuration Information	92
Removing a Cluster From a Site	93
▼ How to Remove a Cluster From a Site	93
▼ How to Remove an Unreachable Cluster From a Site a Cluster	94
Deleting a Site	95
▼ How to Delete a Site	95
9 Administering Multigroups	97
Modifying Protection Groups in a Multigroup	97
▼ How to Add a Protection Group to a Multigroup	97
▼ How to Change the Protection Groups in a Multigroup Dependency Chain	99
▼ How to Remove a Protection Group From a Multigroup	100
Starting and Stopping Multigroups	101
▼ How to Start All Protection Groups in a Multigroup	101
▼ How to Stop All Protection Groups in a Multigroup	102

Synchronizing Multigroup Configuration Information	103
▼ How to Synchronize Multigroup Configuration Information	103
Deleting a Multigroup	104
▼ How to Delete a Multigroup	104
10 Monitoring and Validating the Geographic Edition Software	107
Monitoring the Runtime Status of the Geographic Edition Software	107
Viewing the Geographic Edition Log Messages	114
Displaying Configuration Information for Partnerships and Protection Groups	114
▼ How to Display Configuration Information About Partnerships	115
▼ How to Display Configuration Information About Protection Groups	115
11 Migrating Services	117
Detecting Cluster Failure	117
Detecting Primary Cluster Failure	117
Detecting Secondary Cluster Failure	118
Migrating Replication Services by Switching Over Protection Groups	118
Actions Performed by the Geographic Edition Software During a Switchover	119
▼ How to Switch Over Replication From the Primary Cluster to the Secondary Cluster	120
▼ How to Switch Over a Multigroup	122
Forcing a Takeover of a Protection Group	123
Actions Performed by the Geographic Edition Software After a Takeover	124
▼ How to Force Immediate Takeover of a Protection Group by a Secondary Cluster	124
▼ How to Take Over a Multigroup	126
Recovering a Protection Group to a Cluster	127
12 Customizing Switchover and Takeover Actions	129
Creating a Role-Change Action Script	129
Configuring a Protection Group to Run a Script at Switchover or Takeover	132
▼ How to Configure a Protection Group to Run a Script at Switchover or Takeover	132
13 Script-Based Plug-Ins	133
Introduction to Geographic Edition Script-Based Plug-Ins	133
Advantages and Disadvantages of Using Script-Based Plug-Ins	134
Script-Based Plug-In Architecture	134

Restrictions of Script-Based Plug-Ins	136
Creating Script-Based Plug-Ins	136
Property Descriptions for Script-Based Plug-Ins	137
Protection Group Properties - Overview	137
Replicated Component Properties - Overview	138
Protection Group Property Descriptions	139
Internals for Script-Based Plug-Ins	148
Plug-In Script Functional Requirements	148
Script-Based Plug-In Replication Resource Groups and Resources	150
Protection Group Status Mapped from Replication Resource Status	150
How Geographic Edition Handles Password Properties	151
A Standard Geographic Edition Properties	153
Heartbeat Properties	153
Heartbeat Plug-in Properties	154
Partnership Properties	155
Protection Group Properties	156
Site Properties	157
Multigroup Properties	158
B Legal Names and Values of Geographic Edition Entities	161
Legal Names for Geographic Edition Entities	161
Legal Values for Geographic Edition Entities	162
C Disaster Recovery Administration Example	163
D Takeover Postconditions	165
Results of a Takeover When the Partner Cluster Can Be Reached	165
Results of a Takeover When the Partner Cluster Cannot Be Reached	166
E Troubleshooting Geographic Edition Software	169
Troubleshooting Monitoring and Logging	169
Configuring the Logger File to Avoid Too Many Traces	169
Configuring the Logger File to Avoid Detailed Messages From the gcr Agent	170
Configuring the Logger File to Avoid jmx Remote Traces	170
Troubleshooting Migration Problems	170

Resolving Problems With Application Resource Group Failover When Communication Lost With the Storage Device	170
Troubleshooting Cluster Start and Restart	171
Validating Protection Groups in an Error State	171
Restarting the Common Agent Container	172
Matching the <code>NodeList</code> Property of an Availability Suite Protection Group to Those of Its Device Group and Resource Group	172
F Error Return Codes for Script-Based Plug-Ins	173
Error Codes That Are Returned in Script-Based Plug-Ins	173
Index	177

Figures

FIGURE 2-1	Example Cluster Configuration	27
FIGURE 13-1	Script-Based Plug-Ins Framework	135
FIGURE 13-2	Script-Based Plug-In Replication Resource Group	136

Tables

TABLE 1-1	Geographic Edition CLI	21
TABLE 2-1	IP Addresses Required by Geographic Edition Software	24
TABLE 2-2	Geographic Edition Administration Tasks	25
TABLE 7-1	Commands Used to Stop EMC SRDF Data Replication	71
TABLE 10-1	Status Value Descriptions	109
TABLE 13-1	Protection Group Global Policies	137
TABLE 13-2	Optional Replicated Component Properties	139
TABLE A-1	Geographic Edition Heartbeat Properties	153
TABLE A-2	Geographic Edition Heartbeat Plug-in Properties	154
TABLE A-3	Geographic Edition Partnership Properties	155
TABLE A-4	Geographic Edition Protection Group Properties	156
TABLE A-5	Geographic Edition Site Properties	157
TABLE A-6	Geographic Edition Multigroup Properties	159
TABLE D-1	Takeover Results of Running the <code>geopg takeover</code> Command on the Secondary Cluster	165
TABLE D-2	Takeover Results of Running the <code>geopg takeover</code> Command on the Primary Cluster	166
TABLE D-3	Takeover Results of Running the <code>geopg takeover</code> Command on the Secondary Cluster When the Primary Cluster Cannot Be Reached	166
TABLE D-4	Takeover Results of Running the <code>geopg takeover</code> Command on the Primary Cluster When the Secondary Cluster Cannot Be Reached	167

Examples

EXAMPLE 3-1	Enabling the Geographic Edition Infrastructure in a Cluster	32
EXAMPLE 3-2	Disabling a Cluster	34
EXAMPLE 3-3	Displaying Whether the Geographic Edition Infrastructure Has Been Enabled	36
EXAMPLE 5-1	Modifying the Properties of a Partnership	40
EXAMPLE 5-2	Renaming a Cluster in a Partnership	45
EXAMPLE 5-3	Leaving a Partnership	49
EXAMPLE 5-4	Deleting a Partnership	49
EXAMPLE 5-5	Resynchronizing a Partnership	50
EXAMPLE 6-1	Creating a Heartbeat	53
EXAMPLE 6-2	Creating a Heartbeat Plug-in	54
EXAMPLE 6-3	Modifying the Properties of the Heartbeat Plug-in	55
EXAMPLE 6-4	Deleting a Heartbeat	56
EXAMPLE 6-5	Deleting a Plug-in From a Heartbeat	57
EXAMPLE 6-6	Displaying Heartbeat Configuration Information	58
EXAMPLE 6-7	Modifying the Properties of the Default Heartbeat	60
EXAMPLE 6-8	Adding a Custom Heartbeat Plug-in to the Default Heartbeat	62
EXAMPLE 6-9	Adding a Custom Heartbeat Plug-in to a New Custom Heartbeat	64
EXAMPLE 6-10	Configuring Heartbeat-Loss Notification for an Existing Partnership	65
EXAMPLE 6-11	How a Notification Action Script Parses the Command-Line Information Provided by the Geographic Edition Software	66
EXAMPLE 7-1	Globally Activating a Protection Group	74
EXAMPLE 7-2	Locally Activating a Protection Group	75
EXAMPLE 7-3	Deactivating a Protection Group on All Clusters	77
EXAMPLE 7-4	Deactivating a Protection Group on the Local Cluster	77
EXAMPLE 7-5	Stopping Remote Replication While Leaving the Protection Group Online	77
EXAMPLE 7-6	Deactivating a Protection Group While Keeping Application Resource Groups Online	77
EXAMPLE 7-7	Modifying the Properties of an Oracle ZFS Storage Appliance Remote Replication Component	81

EXAMPLE 7-8	Deleting a Protection Group	82
EXAMPLE 7-9	Deleting an EMC SRDF Protection Group While Keeping Application Resource Groups Online	83
EXAMPLE 7-10	Removing a Remote Replication Component From an Oracle ZFS Storage Appliance Protection Group	84
EXAMPLE 10-1	Displaying Partnership Configuration Information	115
EXAMPLE 10-2	Displaying Configuration Information About a Protection Group	116
EXAMPLE 11-1	Forcing a Switchover From the Primary Cluster to the Secondary Cluster	122
EXAMPLE 11-2	Forcing a Takeover by a Secondary Cluster	126
EXAMPLE 12-1	Switchover Action Script for Updating the DNS	131
EXAMPLE 12-2	Configuring a Protection Group to Run a Command at Cluster Switchover or Takeover	132

Using This Documentation

- **Overview** – Provides procedures for administering Oracle Solaris Cluster Geographic Edition (Geographic Edition) software.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

This document is not to be used as a planning or presales guide.

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/pls/topic/lookup?ctx=E39579>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program web site at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Introduction to Administering the Geographic Edition Software

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

This chapter contains the following sections:

- [“Geographic Edition Administration” on page 19](#)
- [“Geographic Edition Administration Tools” on page 20](#)

Geographic Edition Administration

Familiarize yourself with the planning information in the [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#) and the [“Oracle Solaris Cluster Geographic Edition Overview”](#) before beginning administration tasks. This guide contains the standard tasks that are used to administer and maintain the Geographic Edition configurations.

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

You can perform all administration tasks on a cluster that is running the Geographic Edition software without causing any nodes or the cluster to fail. You can install, configure, start, use, stop, and uninstall the 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 Oracle Solaris Cluster administrative tasks. Refer to the appropriate product documentation for administration restrictions.

Geographic Edition Administration Tools

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

- [“Graphical User Interface” on page 20](#)
- [“Command-Line Interface” on page 21](#)

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

Graphical User Interface

Geographic Edition software supports Oracle Solaris Cluster Manager, a graphical user interface (GUI) tool that you can use to monitor Geographic Edition status and perform various administrative tasks on your cluster. For specific information about how to use Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#) and the Oracle Solaris Cluster Manager online help.

Note - To administer Oracle Solaris Cluster software by using the Oracle Solaris Cluster Manager GUI, ensure that the root passwords are the same on all nodes of both clusters in the partnership.

You can use the GUI to administer Geographic Edition software only after the Geographic Edition software infrastructure has been enabled by using the `geoadm start` command. For information about enabling and disabling the Geographic Edition infrastructure, see [Chapter 3, “Administering the 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 access the GUI, go to the following URL from any Java-enabled and Javascript-enabled browser, and log in as root.

`https://node:8998/scm`

Note - RBAC is not currently supported in the GUI.

Command-Line Interface

Table 1-1 lists the commands that you can use to administer the Geographic Edition software. For more information about each command, refer to the “[Oracle Solaris Cluster Geographic Edition Reference Manual](#)”.

TABLE 1-1 Geographic Edition CLI

Command	Description
geoadm	Enables or disables the 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 Geographic Edition software
geomg	Configures and manages multigroups of protection groups
geops	Creates and manages the partnerships between clusters
geopg	Configures and manages protection groups
geosite	Configures and manages cluster sites

Overview of Administering a Geographic Edition Configuration

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

This chapter contains the following sections:

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

Overview of Oracle Solaris Cluster Administration Concepts

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

This section describes the following Oracle Solaris Cluster administration topics that you need to understand before you administer the Geographic Edition framework:

- [“Configuring Resources and Resource Groups” on page 23](#)
- [“Configuring Logical Hostnames” on page 24](#)
- [“Managing Device Groups” on page 25](#)

Configuring Resources and Resource Groups

You use either Oracle Solaris Cluster commands or the Oracle Solaris Cluster Manager GUI to create failover and scalable resource groups.

For more information about administering resources and resource groups in Oracle Solaris Cluster software, see the [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

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 Geographic Edition Entities”](#).

For more information about using the `geoadm start` command, see [“Enabling the Geographic Edition Software” on page 30](#).

Note - If you are using 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 the Availability Suite Feature of Oracle Solaris,”](#) in [“Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Solaris Availability Suite”](#).

The following table lists the Oracle Solaris Cluster and 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 Oracle Solaris OS software

TABLE 2-1 IP Addresses Required by Geographic Edition Software

Component	Number of IP Addresses Needed
Oracle Solaris Cluster administrative console	1 per subnet
IP Network Multipathing groups	<ul style="list-style-type: none"> ■ Single-adapter groups – 1 primary IP address. ■ 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	1 per domain
Console-access device	1
Logical addresses	1 per logical host resource, per subnet
Geographic Edition infrastructure hostname	1 logical IP address per cluster infrastructure. For example, if you have two clusters in your Geographic Edition infrastructure, you need two IP addresses.
Replication with Availability Suite from Oracle	1 dedicated logical IP address on the local cluster for each device group to be replicated. For example, if

Component	Number of IP Addresses Needed
	you have two clusters in your Geographic Edition infrastructure, you need two IP addresses.

For more information about configuring the IP address and host maps during the installation of Oracle Solaris Cluster software, refer to [Chapter 2, “Installing Software on Global-Cluster Nodes,”](#) in [“Oracle Solaris Cluster Software Installation Guide ”](#).

Managing Device Groups

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

For information about configuring device groups in Oracle Solaris Cluster software, refer to [Chapter 5, “Administering Global Devices, Disk-Path Monitoring, and Cluster File Systems,”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#).

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

For more information about configuring data replication in Geographic Edition software, see the Geographic Edition documentation for the data replication product you use.

Geographic Edition Configuration and Administration Tasks

The following table lists administration tasks for your Geographic Edition configuration.

Note - For procedures to create partnerships, heartbeats, protection groups, sites, and multigroups, see [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide ”](#) and the Geographic Edition manual for your data replication product.

TABLE 2-2 Geographic Edition Administration Tasks


Task	Description and Documentation
Switch over services to the partner cluster.	See the procedures to switch over a protection group from primary to secondary in the Geographic Edition manual for your data replication product. You can also use the Oracle Solaris Cluster Manager GUI.

Task	Description and Documentation
You can switch over individual protection groups or multigroups.	To switch over multiple protection groups that are configured in a multigroup, see “How to Switch Over a Multigroup” on page 122.
Take over services from primary to secondary during a disaster. You can take over services for an individual cluster or for a multigroup.	See the procedures to force immediate takeover of services by a secondary cluster in the Geographic Edition manual for your data replication product. You can also use the Oracle Solaris Cluster Manager GUI. To take over multiple protection groups that are configured in a multigroup, see “How to Take Over a Multigroup” on page 126.
Recover from a takeover.	<ul style="list-style-type: none"> ■ Perform data recovery and error repair outside of the Geographic Edition infrastructure. See the Geographic Edition documentation for the data replication product you are using. ■ Resynchronize the partner clusters. See “Resynchronizing a Protection Group” on page 78. You can also use the Oracle Solaris Cluster Manager GUI.
Start an individual protection group or a multigroup.	To start an individual protection group, see “How to Activate a Protection Group” on page 71 or use the Oracle Solaris Cluster Manager GUI. To start a multigroup, see “How to Start All Protection Groups in a Multigroup” on page 101.
Stop an individual protection group or a multigroup.	To stop an individual protection group, see “How to Deactivate a Protection Group” on page 75. You can also use the Oracle Solaris Cluster Manager GUI. To stop a multigroup, see “How to Stop All Protection Groups in a Multigroup” on page 102.
Delete a multigroup	See “Deleting a Multigroup” on page 104.
Delete a site.	See “Deleting a Site” on page 95.
Delete a protection group or a data replication component.	See “Deleting Protection Groups and Data Replication Components” on page 81. You can also use the Oracle Solaris Cluster Manager GUI.
Delete a partnership.	See “Leaving or Deleting a Partnership” on page 48. You can also use the Oracle Solaris Cluster Manager GUI.
Disable the Geographic Edition software.	See “How to Disable the Geographic Edition Software” on page 33.
Uninstall the Geographic Edition software.	See the “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide” .

Example Geographic Edition Cluster Configuration

The following figure describes an Geographic Edition cluster configuration that is used throughout this guide to illustrate the 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 Geographic Edition Infrastructure

This chapter contains information about enabling your cluster for participation in a partnership. It also contains information for disabling the Geographic Edition software so that your cluster no longer can participate in partnerships.

This chapter contains the following sections:

- [“Geographic Edition Infrastructure Resource Groups” on page 29](#)
- [“Enabling the Geographic Edition Software” on page 30](#)
- [“Disabling the Geographic Edition Software” on page 33](#)
- [“Checking the Status of the Geographic Edition Infrastructure” on page 35](#)
- [“Booting a Cluster” on page 36](#)

Geographic Edition Infrastructure Resource Groups

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

- `geo-clusterstate` – A scalable resource group that the Geographic Edition software uses to distinguish between node failover and cluster reboot scenarios. This resource group does not contain any resources. The resource group contains the following resources:
 - `geo-servicetag` - A scalable resource that is started on all nodes of a cluster when Geographic Edition is present. When the Geographic Edition software is started, this resource checks for the existence of an Oracle Solaris service tag for the running version of Geographic Edition on each node, and creates a service tag if necessary. The service tag indicates that Geographic Edition has been used on the cluster. This service tag is removed from the node when the Geographic Edition packages are removed.
 - `geo-zc-sysevent` - (For zone clusters only) Runs resource methods in the global zone when Geographic Edition is started in a non-global zone. It manages the mechanism which transfers cluster events to subscribers in the zone cluster.

- `geo-infrastructure` – A failover resource group that encapsulates the Geographic Edition infrastructure. The resource group contains the following resources:
 - `geo-cluster-name` – The logical hostname for the Geographic Edition software. The Geographic Edition software uses the logical hostname of a cluster for intercluster 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 Geographic Edition software.
 - `geo-failovercontrol` – Encapsulates the Geographic Edition software itself. The 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 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 Geographic Edition Software

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

For more information about setting up and installing the Geographic Edition software, see the [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

▼ How to Enable Geographic Edition Software

This procedure enables the 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 Oracle Solaris Operating System and the Oracle Solaris Cluster software.
- The 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 Oracle Solaris Cluster installation.

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

```
# cluster list
```

3. If the cluster name is not the name you want to use, change the cluster name.

If you must change the name of a cluster that is configured in a partnership, do not perform this step. Instead, follow instructions in [“Renaming a Cluster That Is in a Partnership”](#) on page 42.

Follow cluster naming guidelines as described in [“Planning Required IP Addresses and Hostnames”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#). Cluster names must follow the same requirements as for host names. You cannot rename a cluster that is part of a zone cluster, but you can rename a cluster in a global zone.

```
# cluster rename -c new-cluster-name old-cluster-name
```

For more information, see the [cluster\(1CL\)](#) man page.

Note - After you have enabled the 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 Geographic Edition infrastructure.

```
# geoadm start
```

The `geoadm start` command enables the 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 Geographic Edition resource groups are online.

For a list of the Geographic Edition resource groups, see [“Geographic Edition Infrastructure Resource Groups”](#) on page 29.

```
# geoadm show
# clresource-group status
# clresource status
```

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

The output for the `clresource-group status` and `clresource status` commands displays that the `geo-failovercontrol`, `geo-hbmonitor`, and `geo-cluster-name` 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 Geographic Edition Infrastructure in a Cluster

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

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

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

2. Ensure that the Geographic Edition infrastructure was successfully enabled.

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

```
--- CLUSTER LEVEL INFORMATION ---
Oracle Solaris 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 Geographic Edition resource groups and resources.

```
phys-paris-1# clresource-group 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-cluster-name	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 - Service is online
phys-paris-2      Offline           Offline

geo-servicetag      phys-paris-1  Online_not_monitored Online_not_monitored
phys-paris-1      Offline           Offline

```

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

Disabling the Geographic Edition Software

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

▼ How to Disable the Geographic Edition Software

Before You Begin Ensure that all protection groups on the local cluster are offline. For more information, see the following procedures:

- To stop an individual protection group, see [“How to Deactivate a Protection Group” on page 75](#).
- To stop a set of protection groups that are configured in a multigroup, see [“How to Stop All Protection Groups in a Multigroup” on page 102](#).



Caution - If you want to keep the application resource groups online while deactivating a protection group, follow [Example 7-6](#) in [“How to Deactivate a Protection Group” on page 75](#).

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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 Geographic Edition Software” on page 107](#).

3. Disable the Geographic Edition software.

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

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

For more information about this command, see the [geoadm\(1M\)](#) man page.

Note - Disabling the 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 Geographic Edition resource groups are no longer displayed.

```
phys-paris-1# geoadm show
phys-paris-1# clresource-group 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 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 Geographic Edition infrastructure was successfully disabled.

```
phys-paris-1# geoadm show

--- CLUSTER LEVEL INFORMATION ---
Oracle Solaris 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 Geographic Edition resource groups and resources have been removed.

```
phys-paris-1# clresource-group status
phys-paris-1#
```

Checking the Status of the Geographic Edition Infrastructure

Use the `geoadm show` command to determine whether the Geographic Edition infrastructure is enabled on the local cluster and on which node the infrastructure is active. The 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 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 ---
Oracle Solaris Cluster Geographic Edition is active on:
node phys-paris-2, cluster cluster-paris

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

Booting a Cluster

The following events take place when you boot a cluster:

1. After the Oracle Solaris Cluster infrastructure is enabled, the 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 Geographic Edition Software” on page 107](#).

Administering RBAC

This chapter describes how to administer role-based access control (RBAC) in Geographic Edition software. For information about RBAC rights for Geographic Edition software, see [“Planning Security” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Modifying a User's RBAC Properties

Note - The RBAC rights for Oracle Solaris Cluster software, including the Geographic Edition framework, are intended to simplify the assignment of management roles and protect against accidental errors. However, a malicious user would be able to abuse RBAC-based cluster administration privileges to gain wider system privileges. Therefore, RBAC rights should be assigned with care.

When you grant authorization to users other than the root role, you must do so on all nodes of both partner clusters. Otherwise, some operations that have a global scope might fail, due to insufficient user rights on one or more nodes in the partnership.

To modify the RBAC rights for a user, you must be logged in as the root role or assume a role that is assigned the System 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, “Oracle Solaris Cluster and RBAC,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Administering Cluster Partnerships

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

This chapter contains the following sections:

- “Modifying Partnership Properties” on page 39
- “Adding a New Cluster Node” on page 40
- “Renaming a Cluster Node” on page 42
- “Renaming a Cluster That Is in a Partnership” on page 42
- “Removing Trust Between Partner Clusters” on page 46
- “Leaving or Deleting a Partnership” on page 48
- “Resynchronizing a Partnership” on page 49

Modifying Partnership Properties

This section provides procedures to modify a property of a partnership.

▼ How to Modify Partnership Properties

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the name of the partnership to go to its page, click the Properties tab, and click Edit to access the listed properties for editing. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in “[Oracle Solaris Cluster System Administration Guide](#)”.

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 “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

2. Modify partnership properties.

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

-p property-setting

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 64](#).

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

partnership-name

Specifies the name of the partnership.

For information about the names and values that are supported by Geographic Edition software, see [Appendix B, “Legal Names and Values of 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 5-1 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@example.com \  
paris-newyork-ps  
# geops list
```

Adding a New Cluster Node

When you add a new node to a cluster that is in a partnership, you must perform additional tasks on that node to make it an active participant in the Geographic Edition configuration.

▼ How to Add a New Node to a Cluster in a Partnership

Perform all steps from the new node.

1. Add the new node to the cluster.

Follow procedures in [Chapter 8, “Administering Cluster Nodes,”](#) in “Oracle Solaris Cluster System Administration Guide”.

2. Install Geographic Edition, data replication, and application software on the new node.

- To install Geographic Edition software, see “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”.
- To install data replication and application software, see the appropriate manual for the software that you use.

3. If the cluster with the new node is the primary for any activated protection groups, remove application resource groups from those protection groups.

This step is necessary to avoid application downtime.

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

4. Deactivate all protection groups that are active on this cluster locally.

```
# geopg stop -e local protection-group
```

5. Stop the Geographic Edition infrastructure.

```
# geoadm stop
```

6. Re-enable the Geographic Edition infrastructure.

This action recreates each Geographic Edition resource group and adds all nodes in the cluster, including the new node, to the node list.

```
# geoadm start
```

7. Reactivate the protection groups that you deactivated in [Step 4](#).

```
# geopg start -e local protection-group
```

8. Restore any application resource groups that you removed in [Step 3](#).

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

Renaming a Cluster Node

You can rename a node in a Geographic Edition cluster that is in a partnership of an Oracle Solaris Cluster configuration.

If the cluster where you are performing the rename procedure is primary for the protection group, and you want to have the application in the protection group online, you can switch the primary group to a secondary during the rename procedure.

For instructions on renaming a node in a Geographic Edition cluster, see [“How to Rename a Node”](#) in [“Oracle Solaris Cluster System Administration Guide”](#). The procedure includes instructions that pertain to a node in a Geographic Edition configuration.

Renaming a Cluster That Is in a Partnership

When you rename a cluster that is in a partnership, the partnership becomes invalid. You must fully unconfigure the existing partnership and create a new one that uses the new cluster name.

▼ How to Rename a Cluster That Is in a Partnership

This procedure demonstrates how to rename one of the global clusters that is in a partnership. You can rename more than one of the clusters at the same time.

Note - You cannot use this procedure to rename a zone cluster in a partnership.

If the cluster that you rename belongs to more than one partnership, perform each step on all clusters that share a partnership with the cluster to rename, before you proceed to the next step in the procedure.

- 1. From one node of the cluster that you are renaming, remove resource groups from each protection group that the cluster belongs to.**

This task avoids production application downtime.

```
# geopg remove-resource-group application-resource-group protection-group
```

- 2. From one node of each cluster in a protection group, confirm that application resource groups have been removed.**

```
# geopg list protection-group
```

3. **From one node of the cluster that you are renaming, stop each protection group globally.**

This task stops data replication.

```
# geopg stop protection-group -e global
```

4. **From one node of each cluster in a protection group, delete the protection group.**

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

5. **From one node of each cluster in a partnership, leave the partnership.**

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

6. **From one node of each cluster, confirm that the protection group and the partnership have been removed.**

```
# geoadm status
```

7. **From one node of each cluster, disable Geographic Edition software.**

```
# geoadm stop
```

8. **From one node of each cluster, confirm that Geographic Edition software was disabled.**

Verify that the geo-infrastructure, geo-clusterstate, and data-replication resource groups are deleted.

```
# clrg list  
# geoadm status
```

9. **From one node of the cluster that you are renaming, change the cluster name.**

Follow cluster naming guidelines as described in [“Planning Required IP Addresses and Hostnames”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

```
# cluster rename -c new-cluster-name
```

Note - The name of the cluster must not include the domain. If a partnership contains clusters that are in different domains, you specify the domain to administrative commands, when necessary, by appending the domain name to the cluster name as *cluster.domain*. Only certain Geographic Edition administrative commands require this fully qualified name when clusters in a partnership are not in the same domain.

10. **Confirm that the cluster name is changed.**

```
# cluster list
```

11. **On each node of both clusters, ensure that hostname entries that match the new cluster name are free and are added to the local `/etc/inet/hosts` files.**

If clusters in the partnership are in different domains, include the domain in the `/etc/hosts` entry for each cluster.

```
# ping new-cluster
    There should be no response
# echo "IP-address new-cluster" >> /etc/inet/hosts
```

12. **From one node of each cluster, start Geographic Edition software.**

```
# geoadm start
```

If Geographic Edition software fails to start, and the failure is not due to problems with the new logical host, restart the common agent container on all nodes by using the `cacoadm restart` command, then start Geographic Edition software.

13. **From one node of each cluster, verify that Geographic Edition software is successfully started.**

```
# geoadm status
```

14. **From one node of each cluster, add trust between the clusters.**

```
# geops add-trust -c remote-partner-cluster[.domain-name]
```

15. **From one node of each cluster, confirm that trust is added successfully.**

Note - Do not specify a domain name to the `verify-trust` subcommand.

```
# geops verify-trust -c remote-partner-cluster
```

16. **Create and join a new partnership between the clusters.**

- a. **From the primary cluster, create the partnership.**

```
# geops create -c remote-partner-cluster[.domain-name] partnership
```

- b. **From the secondary cluster, join the partnership.**

```
# geops join-partnership remote-partner-cluster[.domain-name] partnership
```

17. **On each cluster, confirm that the new partnership is successfully created and joined.**

```
# geoadm status
```

18. **If you did not reboot the nodes of the cluster that you renamed, restart the heartbeats on each node of the renamed cluster.**

Restarting the heartbeat initiates the heartbeat to read and store the new cluster name.

```
# svcadm disable svc:/system/cluster/gchb_resd:default
# svcadm enable svc:/system/cluster/gchb_resd:default
```

Example 5-2 Renaming a Cluster in a Partnership

This example renames the cluster `newyork`, in the `paris-newyork-ps` partnership, to `chicago`. The names of the nodes in this cluster are not changed, so `phys-newyork-1` becomes a node in the newly named `chicago` cluster. The `paris-newyork-ps` partnership is first unconfigured. After the cluster is renamed, a new `paris-chicago-ps` partnership is created with the `chicago` cluster as primary and the `paris` cluster as secondary. The two clusters belong to the same domain, so the domain name is not specified to the commands.

```
phys-newyork-1# geopg remove-resource-group app-rg

phys-newyork-1# geopg list examplepg
phys-paris-1# geopg list examplepg

phys-newyork-1# geopg stop examplepg -e global

phys-newyork-1# geopg delete examplepg
phys-paris-1# geopg delete examplepg

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

phys-newyork-1# geoadm stop
phys-paris-1# geoadm stop

phys-newyork-1# clrg list
phys-newyork-1# geoadm status
phys-paris-1# clrg list
phys-paris-1# geoadm status

phys-newyork-1# cluster rename -c chicago
phys-newyork-1# cluster list

phys-newyork-1# ping chicago
phys-newyork-1# echo "192.168.10.1 chicago" >> /etc/hosts
    repeat on each node of the chicago cluster

phys-paris-1# ping chicago
phys-paris-1# echo "192.168.20.1 chicago" >> /etc/hosts
    repeat on each node of the paris cluster

phys-newyork-1# geoadm start
phys-paris-1# geoadm start

phys-newyork-1# geoadm status
```

```
phys-paris-1# geoadm status

phys-newyork-1# geops add-trust -c paris
phys-paris-1# geops add-trust -c chicago

phys-newyork-1# geops verify-trust -c paris
phys-paris-1# geops verify-trust -c chicago

phys-newyork-1# geops create -c paris paris-chicago-ps
phys-paris-1# geops join-partnership chicago paris-chicago-ps

phys-newyork-1# geoadm status
phys-paris-1# geoadm status

phys-newyork-1# /etc/init.d/initgchb_resd stop
phys-newyork-1# /etc/init.d/initgchb_resd start
    repeat on each node of the chicago cluster

phys-paris-1# svcadm disable svc:/system/cluster/gchb_resd:default
phys-paris-1# svcadm enable svc:/system/cluster/gchb_resd:default
    repeat on each node of the paris cluster
```

Next Steps Perform the following tasks:

- Create a new protection group and replicate it to partner.
- Add device groups.
- Start globally.
- Add resource groups to the protection group and verify the configuration.

Note - When you create the new protection group, pay close attention to which cluster is the primary and which is the secondary, to ensure that data replication is started in the desired direction.

Follow procedures in the appropriate Geographic Edition data-replication guide.

Removing Trust Between Partner Clusters

This section provides procedures to remove trust between partner clusters.

▼ How to Remove Trust Between Two Clusters

Perform this procedure to remove trust from between two clusters.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships and then click Remove Partner Trust. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,” in “Oracle Solaris Cluster System Administration Guide ”](#).

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 [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide ”](#).

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 remote-partner-cluster-name
```

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

```
-c remote-partner-cluster-name
```

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 `remove-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:

```
remote-partner-cluster-name.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.

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

Leaving or Deleting a Partnership

This section provides procedures to leave or delete a partnership. Because the `geops leave-partnership` command destroys the local partnership configuration information, when the last member leaves a partnership, the partnership no longer exists.

You can alternatively use the `geops` command to remove a cluster from a partnership and release all the resources that are associated with the partnership.

▼ How to Leave or Delete a Partnership

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, highlight the partnership to leave, and click Leave Partnership. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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 [“Security Certificates”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. **Verify that the partnership does not contain 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 [“Deleting Protection Groups and Data Replication Components”](#) on page 81.

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

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


partnership-name

Specifies the name of the partnership

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

For more information, refer to the [geops\(1M\)](#) man page.

Example 5-3 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 5-4 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

This section provides information about resynchronizing partnership information between its partner clusters.

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 or Oracle Solaris Cluster Manager 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

Perform this procedure to resynchronize partnership configuration information.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, then highlight the partnership name and click `Update Partnership`. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Before You Begin Ensure that the following conditions are met:

- The local cluster is 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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Resynchronize the partnership.

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

```
partnership-name
```

Specifies the name of the partnership

Example 5-5 Resynchronizing a Partnership

This example resynchronizes a partnership.

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

Administering Heartbeats

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 51
- “Creating a Heartbeat” on page 52
- “Creating a Heartbeat Plug-in” on page 53
- “Modifying a Heartbeat Plug-in Property” on page 54
- “Deleting Heartbeats and Heartbeat Plug-ins” on page 56
- “Displaying Heartbeat Configuration Information” on page 57
- “Tuning the Heartbeat Properties” on page 58
- “Creating a Heartbeat That Uses a Custom Heartbeat Plug-in” on page 60
- “Configuring Heartbeat-Loss Notification” on page 64

Introduction to Heartbeats

A heartbeat in 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 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 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_local-cluster-name-remote-partner-cluster-name`. For more information about the `geops` command, refer to the [geops\(1M\)](#) man page.

You can create custom heartbeat plug-ins and associate them with existing default heartbeats or with new custom heartbeats.

Note - Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Oracle 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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Create the heartbeat.

```
# geohb create -r remote-partner-cluster-name \
[-p property-setting [-p...]] heartbeat-name
```

-r remote-partner-cluster-name

Specifies the name of the remote, secondary partner cluster.

-p property-setting

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 Geographic Edition Properties”](#).

heartbeat-name

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

```
# geohb add-plugin heartbeat-name plug-in-name [-p property-setting [-p..]]
```

heartbeat-name

Specifies the identifier for heartbeat on the local cluster.

plug-in-name

Specifies the name of the heartbeat plug-in.

-p property-setting

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 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

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, then click the partnership name. In the Heartbeats section of the Partnership panel, click the heartbeat name. In the Heartbeat panel, click the plug-in name. In the Heartbeat Plug-Ins panel, click the Properties tab, then click Edit. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Modify the heartbeat plug-in properties.

```
# geohb modify-plugin -p property-setting [-p...] plug-in-name heartbeat-name
```

heartbeat-name

Specifies an identifier for the heartbeat.

plug-in-name

Specifies the name of the heartbeat plug-in.

-p *property-setting*

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 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 Geographic Edition software, see [Appendix B, “Legal Names and Values of 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 provides the following procedures:

- [“How to Delete a Heartbeat” on page 56](#)
- [“How to Delete a Plug-in From a Heartbeat” on page 56](#)

▼ 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 [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Delete the heartbeat.

```
# geohb delete heartbeat-name
```

heartbeat-name

Specifies an identifier for the heartbeat settings.

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

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Remove the plug-in from the heartbeat.


```
# geohb remove-plugin plug-in-name heartbeat-name
```



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

plug-in-name

Specifies the name of the custom heartbeat plug-in

heartbeat-name

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

For information about the names and values that are supported by Geographic Edition software, see [Appendix B, “Legal Names and Values of 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

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, then click the partnership name. In the Heartbeats section click the heartbeat name, then click the Properties tab. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

```
# geohb list [heartbeat-name-list]
```

heartbeat-name-list

Specifies the names of the 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 Oracle 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 an Oracle service representative.

▼ How to Modify the Heartbeat Properties

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, then click the partnership name. In the Heartbeats section click the heartbeat name. In the Heartbeat panel, click the Properties tab then click Edit. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in “Oracle Solaris Cluster System Administration Guide”.

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 “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

2. Modify the heartbeat properties.

```
# geohb set-prop -p property-setting [-p...] heartbeat-name
```

`-p property-setting`

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 Geographic Edition Properties”](#).

`heartbeat-name`

Specifies an identifier for the heartbeat settings.

For information about the names and values that are supported by Geographic Edition software, see [Appendix B, “Legal Names and Values of 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 Oracle 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.

This section contains the following information:

- [“Creating a Custom Heartbeat Plug-in” on page 60](#)
- [“How to Add a Custom Heartbeat Plug-in to an Existing Default Heartbeat” on page 61](#)
- [“How to Create a Custom Heartbeat Plug-in and Add It to a Custom Heartbeat” on page 62](#)

Creating a Custom Heartbeat Plug-in

When a heartbeat is created, your custom heartbeat plug-in is passed the following arguments by the Geographic Edition software:

query-interval

The value of the `Query-interval` property, which defines the delay in seconds after which a heartbeat status request is declared a failure.

mode

The mode for the plug-in startup, either `Normal` or `Emergency`.

plug-in-property-values

The value of the `Plugin-properties` property that is configured for the heartbeat plug-in, if any.

For more information about the properties you can set, see [Appendix A, “Standard 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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

```
# geohb add-plugin -p property-setting [-p...] \  
plug-in-name hb_local-cluster-name-remote-partner-cluster-name
```

-p property-setting

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 Geographic Edition Properties”](#).

plug-in-name

Specifies the name of the custom heartbeat plug-in.

hb_local-cluster-name-remote-partner-cluster-name

Specifies the name of the default heartbeat to which you want to add the custom heartbeat plug-in.

3. Verify that your changes were made correctly.

```
# geoadm status
```

4. **Repeat the previous steps on 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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. **Create the new custom heartbeat.**

```
# geohb create -r remote-partner-cluster-name [-p property-setting [-p...]] heartbeat-name
```

-r remote-partner-cluster-name

Specifies the name of the remote, secondary partner cluster.

-p property-setting

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 Geographic Edition Properties”](#).

heartbeat-name

Specifies an identifier for the heartbeat settings.



Caution - The name of the custom heartbeat on each cluster in the same partnership must be different. Choose a name that 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 property-setting [-p...] plug-in-name heartbeat-name
```

-p property-setting

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 Geographic Edition Properties”](#).

plug-in-name

Specifies the name of the custom heartbeat plug-in.

heartbeat-name

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

-c remote-partner-cluster-name

Specifies the name of remote cluster that will participate in the partnership. This name matches the logical hostname used by the Geographic Edition infrastructure on the remote cluster.

-h heartbeat-name

Specifies the custom heartbeat to be used in the partnership to monitor the availability of the partner cluster.

-p property-setting

Sets the value of partnership properties with a string of *name=value* pair statements. For more information about the properties you can set, see [Appendix A, “Standard Geographic Edition Properties”](#).

partnership-name

Specifies the name of the partnership.

For more information about using `geops create` command to create a partnership, see [“How to Create a Partnership”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

This section contains the following information:

- [“Introduction to Configuring Heartbeat-Loss Notification” on page 64](#)
- [“Configuring the Heartbeat-Loss Notification Properties” on page 65](#)
- [“Creating an Action Shell Script for Heartbeat-Loss” on page 65](#)

Introduction to Configuring Heartbeat-Loss Notification

You can configure the 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_emailaddrs` and `Notification_actioncmd` properties.

Heartbeat-loss notification occurs if the heartbeat still fails after the interval you configure with the `Query_interval` property of the heartbeat. The heartbeat monitor sends 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 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_emailaddrs` 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”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#). You can also modify these properties by using the procedure that is described in [“How to Modify the Heartbeat Properties”](#) on page 59.

If you want to be notified of heartbeat loss by email, set the `Notification_emailaddrs` property. You can specify a list of email addresses, separated by commas. If you want to use email notification, the cluster nodes must be configured as email clients. For more information about configuring mail services, see [Chapter 2, “Administering Mail Services,”](#) in [“Managing sendmail Services in Oracle Solaris 11.2”](#).

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_emailaddrs=ops@paris.example.com,ops@newyork.example.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:

```
# custom-action-command-path -c local-cluster-name -r remote-partner-cluster-name -e 1 \
-n nodename -t time
```

custom-action-command-path

Specifies a path to the action command you have created.

-c local-cluster-name

Specifies the name of the local cluster.

-p remote-partner-cluster-name

Specifies the name of the remote partner cluster.

-e1

Specifies that `HBLOST=1`, which indicates that a heartbeat-loss event has occurred. The Geographic Edition software only supports heartbeat-loss notification, so `-e 1` is the only value that can be passed to the action shell script.

-nnodename

Specifies the name of the cluster node that sent the heartbeat-loss event notification.

-t timestamp

Specifies the time of the heartbeat-loss event as the number of milliseconds since January 1, 1970, 00:00:00 GMT.



Caution - You can use this script to perform an automatic takeover on the secondary cluster. However, such an automated action is risky. If the heartbeat-loss notification is caused by a total loss of all heartbeat connectivity on both the primary and secondary clusters, such an automated action could lead to a situation where two primary clusters exist.

EXAMPLE 6-11 How a Notification Action Script Parses the Command-Line Information Provided by the 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 following information:

- [“Activating and Deactivating a Protection Group”](#) on page 69
- [“Resynchronizing a Protection Group”](#) on page 78
- [“Modifying Protection Groups and Data Replication Components”](#) on page 79
- [“Deleting Protection Groups and Data Replication Components”](#) on page 81

Activating and Deactivating a Protection Group

This section provides the following information:

- [“Guidelines for Activating and Deactivating a Protection Group”](#) on page 69
- [“How to Activate a Protection Group”](#) on page 71
- [“How to Deactivate a Protection Group”](#) on page 75

Guidelines for Activating and Deactivating a Protection Group

When you activate a protection group, the protection group assumes the role that you assigned to it during configuration. When you deactivate a protection group, its application resource groups are also unmanaged.

You can activate or deactivate a protection group in the following ways:

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

When a protection group is activated on the primary cluster, the application resource groups that are configured for the protection group are also started. The Geographic Edition

software uses the following Oracle Solaris Cluster commands on the primary cluster to bring the resource groups online:

```
# clresourcegroup online -eM resource-group-list
```

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

When you activate a protection group, the data replication product that you are using determines the clusters on which data replication can start. The following sections describe additional behaviors when you activate or deactivate a particular data replication product:

- [“Effects of Activating an Availability Suite Protection Group” on page 70](#)
- [“Effects of Activating and Deactivating an EMC SRDF Protection Group” on page 70](#)
- [“Effects of Activating an Oracle Data Guard Protection Group” on page 71](#)

Effects of Activating an Availability Suite Protection Group

The Availability Suite feature allows data replication to start only from the primary cluster. So, if you activate a protection group from the secondary cluster, data replication does not start.

Effects of Activating and Deactivating an EMC SRDF Protection Group

Activating and deactivating protection group on a cluster has the following effect on the data replication layer:

- When activated, the data replication configuration of the protection group is validated. During validation, the current local role of a protection group is compared with the configuration of the EMC SRDF device groups.
 - If the EMC SRDF device group is not in a `Failedover` state, the local role of the protection group should match the role of the EMC SRDF device group.
 - If the EMC SRDF device group is in a `Failedover` state, then the local role of the protection group becomes secondary while the role of the EMC SRDF device group remains primary.

Data replication is started on the data replication device groups that are configured for the protection group, no matter whether the activation occurs on a primary or secondary cluster. Data is always replicated from the cluster on which the local role of the protection group is primary to the cluster on which the local role of the protection group is secondary.

- Deactivating an EMC SRDF protection group on a cluster has the following effect on the data replication layer:
 - The data replication configuration of the protection group is validated. During validation, the current local role of the protection group is compared with the aggregate device group state. If validation is successful, data replication is stopped.

- Data replication is stopped on the data replication device groups that are configured for the protection group, whether the deactivation occurs on a primary or secondary cluster.

Deactivating an EMC SRDF protection group has the following effect on the application layer:

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

The EMC SRDF command that is used to stop data replication depends on the RDF state of the EMC SRDF device group.

The following table describes the EMC SRDF command that is used to stop data replication for each of the possible combinations of factors.

TABLE 7-1 Commands Used to Stop EMC SRDF Data Replication

Aggregate Device Group State	Valid Local Protection Group Role	EMC SRDF Command
Split, Suspended, Partitioned, or Failover	primary or secondary	No command is run because no data is being replicated.
Synchronized or R1Updated	primary or secondary	The <code>symrdf split</code> command is run.

Effects of Activating an Oracle Data Guard Protection Group

The Oracle Data Guard software allows data replication to start only if you activate a protection group in one of the following ways:

- Locally from the primary cluster.
- Globally from either the primary or the standby cluster.

So, if you attempt to activate an Oracle Data Guard protection group locally from the standby cluster, data replication does not start. However, if you activate the protection group globally from the standby cluster, data replication does start.

▼ How to Activate a Protection Group

This procedure activates, or starts, the protection group on the primary and secondary clusters, depending on the scope of the command. When you activate a protection group on the primary cluster, its application resource groups are also brought online.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section highlight the protection group name and click Start Protection Group. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in “Oracle Solaris Cluster System Administration Guide”.

1. Assume the root role or assume a role that is assigned the Geo Management RBAC rights profile.

For more information about RBAC, see “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the root role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Activate the protection group.

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

```
phys-node-n# geopg start -e scope [-n] protection-group-name
```

`-e scope`

Specifies the scope of the command.

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

Note - The property values `global` and `local` are *not* case sensitive.

`-n`

Prevents the start of replication at protection group startup.

If you omit this option, the replication subsystem starts at the same time as the protection group. The following data replication products have additional behaviors:

- For Oracle Data Guard, if you omit this option, the `geopg start` command performs the following operations on each Oracle Data Guard Broker configuration in the protection group:

- Verifies that the resource group that is named in the `local_oracle_svr_rg_name` property contains a resource of type `SUNW.scalable_rac_server_proxy` for a scalable resource group or a resource of type `SUNW.oracle_server` for a failover resource group.
- Verifies that the Oracle Data Guard `dgmgrl` command can connect using the values that are given for `sysdba_username`, `sysdba_password`, and `local_db_service_name`. Or if the `sysdba_username` and `sysdba_password` properties are null, verifies that the `dgmgrl` command can connect using the Oracle wallet connection format, `dgmgrl /@local_db_service_name`.
- Verifies that the role configured for the replication resource is the same as the role of the protection group on the local cluster.
- Verifies that the Oracle Data Guard Broker configuration details match those that are held by Geographic Edition. The details to check include which cluster is primary, the configuration name, the database mode (for both the primary and standby clusters), the replication mode, the standby type, that `FAST_START FAILOVER` is disabled, and that `BystandersFollowRoleChange` is equal to `NONE`.
- For MySQL, if you omit this option, the `geopg start` command performs the following actions, if the role of the protection group is secondary on the local cluster:
 - Starts the MySQL slave threads
 - Prevents modification by an unprivileged database user if this option is configured
 - Prepares the `my.cnf` file to start the database with modifications prevented for an unprivileged database user if this option is configured
- For Availability Suite, if you omit this option, the data replication subsystem starts at the same time as the protection group and the `geopg start` command performs the following operations on each device group in the protection group:
 - Verifies that the role configured for the replication resource is the same as the role of the protection group on the local cluster.
 - Verifies that the role of the volume sets associated with the device group is the same as the role of the protection group on the local cluster.
 - If the role of the protection group on the local cluster is secondary, unmounts the local volumes defined in all volume sets associated with the device group.
 - If the role of the protection group on the local cluster is primary, enables the autosynchronization feature of the Availability Suite remote mirror feature. Also, resynchronizes the volume sets associated with the device group.

protection-group-name

Specifies the name of the protection group.

The `geopg start` command uses the `clresourcegroup online -eM resourcegroup` command to bring resource groups and resources online. For more information about using this command, see the [clresourcegroup\(1CL\)](#) man page.

- If the role of the protection group is primary on the local cluster, the `geopg start` command performs the following operations:
 - Runs a script that is defined by the `RoleChange_ActionCmd` property.
 - Brings the application resource groups in the protection group online on the local cluster. For Oracle Data Guard, this includes the shadow Oracle database server resource groups.

The following data replication products have additional behaviors:

- For MySQL, the `geopg start` command also performs the following operations:
 - Prepares the `my.cnf` file to start the database without the slave threads
 - Brings online the application resource groups in the protection group on the local cluster
- For Availability Suite, the `geopg start` command also performs the following operations:
 - If the application resource group is a failover type resource group that shares affinities with a device group in the same protection group, the command adds strong, positive affinities and failover delegation between the application resource group and the lightweight resource group.

The application resource group must not have strong, positive affinities with failover delegation. Otherwise, the attempt to add strong, positive affinities with failover delegation with the lightweight resource group will fail.

- The command creates strong dependencies between the `HASStoragePlus` resource in the application resource group and the `HASStoragePlus` resource in the lightweight resource group for this device group.

If the command fails, the `Configuration` status might be set to `Error`, depending on the cause of the failure. The protection group remains deactivated, but data replication might be started and some resource groups might be brought online. Run the `geoadm status` command or Oracle Solaris Cluster Manager to obtain the status of your system.

If the `Configuration` status is set to `Error`, revalidate the protection group by using the procedures that are described in [“How to Validate a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Example 7-1 Globally Activating a Protection Group

This example globally activates a protection group.

```
phys-paris-1# geopg start -e global sales-pg
```

Example 7-2 Locally Activating a Protection Group

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

```
phys-paris-1 geopg start -e local sales-pg
```

▼ How to Deactivate a Protection Group

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section highlight the protection group name and click Stop Protection Group. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

You can deactivate, or stop, a protection group in the following ways:

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

This procedure deactivates the protection group on all nodes of the primary and secondary clusters, depending on the scope of the command. When you deactivate a protection group, its application resource groups are also unmanaged.

The result of deactivating a protection group on the primary or standby cluster depends on the type of data replication that you are using.

- If you are using Oracle Data Guard software, you can stop the Oracle Data Guard configuration from the primary or the standby cluster when the configuration is enabled because the Oracle Data Guard command-line interface (dgmgrl) on both clusters still accepts commands.
- Deactivating an EMC SRDF protection group on a cluster has the following effect on the data replication layer:
 - The data replication configuration of the protection group is validated. During validation, the current local role of the protection group is compared with the aggregate device group state. If validation is successful, data replication is stopped.
 - Data replication is stopped on the data replication device groups that are configured for the protection group, whether the deactivation occurs on a primary or secondary cluster.
- If you are using the Availability Suite feature, data replication can be stopped only from the primary cluster. So, when you deactivate a protection group on the secondary cluster, this deactivate command does not stop data replication.

1. Assume the root role or assume a role that is assigned the Geo Management RBAC rights profile.

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the root role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Deactivate the protection group.

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

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

`-e scope`

Specifies the scope of the command.

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

Note - The property values, such as `global` and `local`, are *not* case sensitive.

`-D`

Specifies that only replication should be stopped and the protection group should be online.

If you omit this option, the replication subsystem and the protection group are both stopped. If the role of the protection group on the local cluster is primary, omitting the `-D` option also results in taking the application resource groups offline and putting them in an unmanaged state.

`protection-group-name`

Specifies the name of the protection group.

The following data replication products have additional behaviors:

- For Oracle Data Guard, if the role of the protection group is primary on the local cluster, the `geopg stop` command disables the Oracle Data Guard Broker configuration.
 - If the `geopg stop` command fails, run the `geoadm status` command or Oracle Solaris Cluster Manager to see the status of each data replication component. For example, the

Configuration status might be set to Error depending on the cause of the failure. The protection group might remain activated even though some resource groups might be unmanaged. The protection group might be deactivated with data replication running.

- If the Configuration status is set to Error, revalidate the protection group by using the procedures described in [“How to Validate a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).
- For Availability Suite, if the role of the protection group is primary on the local cluster, the `geopg stop` command disables the autosynchronization of each device group and places the volume sets into logging mode.

Example 7-3 Deactivating a Protection Group on All Clusters

This example deactivates a protection group on all clusters.

```
# geopg stop -e global sales-pg
```

Example 7-4 Deactivating a Protection Group on the Local Cluster

This example deactivates a protection group on the local cluster.

```
# geopg stop -e local sales-pg
```

Example 7-5 Stopping Remote Replication While Leaving the Protection Group Online

This example stops replication on the local cluster only.

```
# geopg stop -e local -D sales-pg
```

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

```
# geopg stop -e local sales-pg
```

Example 7-6 Deactivating a Protection Group While Keeping Application Resource Groups Online

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

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

```
# geopg remove-resource-group apprg1,apprg2 sales-pg
```

2. Deactivate the protection group.

```
# geopg stop -e global sales-pg
```

Troubleshooting If the `geopg stop` command fails, run the `geoadm status` command or Oracle Solaris Cluster Manager to obtain the status of each data replication component. For example, the configuration status might be set to Error depending on the cause of the failure. The protection group might remain activated even though some resource groups might be unmanaged. The protection group might be deactivated with replication running.

If the configuration status is set to Error, revalidate the protection group. See [“Validating a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Resynchronizing a Protection Group

You can resynchronize the configuration information of the local protection group with the configuration information retrieved from the partner cluster. You need to resynchronize a protection group when its Synchronization status in the output of the `geoadm status` command or in Oracle Solaris Cluster Manager is Error. For example, you might need to resynchronize protection groups after booting the cluster. For more information, see [Booting a Cluster](#).

Resynchronizing a protection group updates only entities that are related to Geographic Edition. To update Oracle Solaris Cluster resource groups, resource types, and resources, use the `cluster export -t rg,rt,rs` command to generate an XML cluster configuration file, modify the XML file for the expected configuration on the secondary cluster, and run the `clresource create` command with the `-a` option to apply the configuration updates. For more information, see [“How to Configure Oracle Solaris Cluster Software on All Nodes \(XML\)”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#) and the `cluster(1CL)` and `clresource(1CL)` man pages.

▼ How to Resynchronize a Protection Group

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section highlight the protection group name and click Update Protection Group. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Before You Begin The protection group must be deactivated on the cluster where you run the `geopg update` 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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Resynchronize the protection group.

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

protection-group-name

Specifies the name of the protection group

Modifying Protection Groups and Data Replication Components

This section contains the following procedures:

- [“How to Modify a Protection Group”](#) on page 79
- [“How to Modify a Data Replication Component”](#) on page 80

▼ How to Modify a Protection Group

Perform this procedure to modify the properties of a protection group.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section click the protection group name. In the Protection Group page, click the Properties tab, then click Edit. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Before You Begin Ensure that the protection group you want to modify exists locally.

1. Log in to one of the cluster nodes.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Modify the configuration of the protection group.

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

```
# geopg set-prop -p property-setting [-p...] protection-group
```

-p property-setting

Specifies the properties of the protection group.

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

protection-group

Specifies the name of the protection group.

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

For more information about the `geopg` command, see the [geopg\(1M\)](#) man page.

▼ How to Modify a Data Replication Component

Use this procedure to modify a property of a data replication component.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section click the protection group name. In the Protection Groups panel, click the data replication component name. Then in the Data Replication component panel, click Edit. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in “Oracle Solaris Cluster System Administration Guide”.

1. Assume the root role or assume a role that is assigned the Geo Management RBAC rights profile.

For more information about RBAC, see “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the root role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and Oracle ZFS Storage Appliance software.

2. Modify the data replication component.

The following command modifies the properties of a data replication component in a protection group on the local cluster. The command then propagates the new configuration to the partner cluster if the partner cluster contains a protection group with the same name.

```
# geogg modify-replication-component -p property [-p...] replication-component protection-group
```

-p property

Specifies the properties of the data replication component.

replication-component

Specifies the name of the data replication component.

protection-group

Specifies the name of the protection group that contains the data replication component.

Example 7-7 Modifying the Properties of an Oracle ZFS Storage Appliance Remote Replication Component

The following example modifies the Timeout property of the remote replication component `trancos` which is part of the appliance protection group `zfssa-p`.

```
# geogg modify-replication-component -p Timeout=215 trancos zfssa-pg
```

Deleting Protection Groups and Data Replication Components

This section provides the following procedures:

- [“How to Delete a Protection Group” on page 81](#)
- [“How to Remove a Data Replication Component From a Protection Group” on page 83](#)

▼ How to Delete a Protection Group

Use the following procedure to delete a protection group from the local cluster.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section highlight the protection group name and click Delete. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,” in “Oracle Solaris Cluster System Administration Guide”](#).

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

To keep the application resource groups online while deleting the protection group, you must remove the application resource groups from the protection group. See [Example 7-9](#) for examples of this procedure.

Before You Begin Ensure that the following conditions are met:

- The protection group you want to delete exists locally.
- The protection group is offline on all clusters from which you want to delete it.

1. Log in to one of the nodes on the primary cluster, `cluster-paris`.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

2. Delete the protection group.

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

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

```
protection-group-name
```

Specifies the name of the protection group

3. To also delete the protection group on the secondary cluster, repeat Step 1 and Step 2 from a node of the secondary cluster.

Example 7-8 Deleting a Protection Group

This example deletes a protection group `srdfpg` from both partner clusters. The protection group is offline on both partner clusters.

In this example, `phys-paris-1` is a node of the primary cluster and `phys-newyork-1` is a node of the secondary cluster.

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

Example 7-9 Deleting an EMC SRDF Protection Group While Keeping Application Resource Groups Online

This example keeps online two application resource groups, `apprg1` and `apprg2`, while deleting their protection group, `srdpfg`, from both partner clusters. The following commands remove the application resource groups from the protection group, then delete the protection group.

```
phys-paris-1# geopg remove-resource-group apprg1,apprg2 srdpfg
phys-paris-1# geopg stop -e global srdpfg
phys-paris-1# geopg delete srdpfg
phys-newyork-1# geopg delete srdpfg
```

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

For Oracle Data Guard, unlike other data replication modules, the Oracle database-server resource group is not added to the protection group. Instead, a shadow Oracle database-server resource group is added to represent this resource group. You can add and remove the shadow Oracle database-server resource group to and from the protection group at any time without affecting the Oracle Data Guard data replication. Consequently, the application resource groups that are shown in this example would have no data to replicate in an Oracle Data Guard protection group. Application resource groups that might meet this criteria can be scalable web servers, where their data is static or held on some remote storage that is not controlled by this cluster.

▼ How to Remove a Data Replication Component From a Protection Group

Use the following procedure to remove a data replication component from a protection group.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section click the protection group name. In the Data Replication components section, highlight the data replication component name and click Remove. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Before You Begin Ensure that the following conditions are met:

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

1. Assume the root role or assume a role that is assigned the Geo Management RBAC rights profile.

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the root role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwX:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and Oracle ZFS Storage Appliance software.

2. Remove the data replication component.

```
# geopg remove-replication-component replication-component protection-group
```

replication-component

Specifies the name of the data replication component.

protection-group

Specifies the name of the protection group.

Example 7-10 Removing a Remote Replication Component From an Oracle ZFS Storage Appliance Protection Group

In the following example, the remote replication component `trancos` is removed from the appliance protection group, `zfssa-pg`.

```
# geopg remove-replication-component trancos zfssa-pg
```

Administering Sites

You administer a site from a cluster that is a controller of the site. Changes to the site are propagated by the issuing controller cluster to all other clusters in the site.

This chapter contains the following information about administering a site:

- [“Adding a Cluster to a Site” on page 85](#)
- [“Changing the Site Role of a Site Cluster” on page 89](#)
- [“Synchronizing Site Configuration Information” on page 92](#)
- [“Removing a Cluster From a Site” on page 93](#)
- [“Deleting a Site” on page 95](#)

Adding a Cluster to a Site

This section contains the following procedures to add a cluster to a site.

- [“How to Add a Cluster to a Site” on page 85](#)
- [“How to Accept an Invitation to Join a Site” on page 87](#)

▼ How to Add a Cluster to a Site

Perform this command to add one or more clusters to an existing site. You can add a global cluster or a zone cluster. For procedures to create a site, see [“How to Create a Site” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

1. **From a node of a site controller cluster, assume the root role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Ensure that all nodes of the site controller cluster are online.

```
phys-schost-1# cluster status -t node
=== Cluster Nodes ===

--- Node Status ---

Node Name                               Status
-----
phys-schost-2                           Online
phys-schost-1                           Online
```

If any node is offline, wait until the node is brought back up before you issue the site invitation. The addition of a new cluster to a site will fail if any node in the issuing cluster is not online.

3. Invite the cluster to join the site.

- To add the cluster as a site controller, use the following command.

```
site-controller# geosite add-member -c cluster site
```

- To add the cluster as a member, use the following command.

```
site-controller# geosite add-member -m cluster site
```

`-c`
Specifies the site controller role.

`-m`
Specifies the site member role.

cluster

The name of the cluster to add to the site. To add multiple clusters, separate the cluster names with a comma (.). You can use both the `-c` and `-m` options in the same `geosite add-member` command.

site

The name of the site that you are adding the cluster to.

4. Verify the invitation.

Command output is similar to the following, where the new cluster is added as a site controller.

```
site-controller## geosite status

Site : site

Controller "issuing-cluster"
Configuration                : OK

Controller "added-cluster"
Configuration                : OK
Synchronization              : Unknown
...
```

Troubleshooting If the `geosite add-member` command fails with a timeout, update the site's Timeout value to a larger number and retry the command.

Next Steps To complete the addition of the cluster to the site, the invited cluster must join the site. Go to [“How to Accept an Invitation to Join a Site” on page 87](#).

▼ How to Accept an Invitation to Join a Site

Perform this procedure to complete the addition of a cluster to a site.

- Before You Begin**
- Ensure that an invitation to join a site has been issued for the cluster. See [“How to Add a Cluster to a Site” on page 85](#).
 - Ensure that all nodes of the cluster that issued the invitation are online.
 - Ensure that the common agent container is started on all nodes of both the cluster that issued the invitation and the invited cluster.

1. From a node of the site cluster that issued the invitation, ensure that all nodes are online.

```
phys-schost-1# clnode status
=== Cluster Nodes ===

--- Node Status ---

Node Name                Status
-----                -
phys-schost-2            Online
phys-schost-1            Online
```

If any node of the issuing cluster is offline, wait until the node is back online before you accept the site invitation. The acceptance of a site invitation will fail if any node in the issuing cluster is not online.

2. **From a node of the invited cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

3. **Ensure that all nodes of the invited cluster are online.**

```
invited-cluster-node# cluster status -t node
```

4. **Accept the invitation to join the site.**

```
invited-cluster-node# geosite join issuing-cluster site
```

issuing-cluster

The name of the cluster that issued the `geosite` command to add the invited cluster.

site

The name of the site that the invited cluster is joining.

5. **Verify that the cluster is a member of the site.**

The following command lists all sites that the issuing cluster is a member of.

```
invited-cluster-node# geosite list
site
```

Troubleshooting If a `geosite join` operation fails due to a failure to verify trust, ensure that the common agent container is running on all nodes of both the invited cluster and the issuing cluster. If a node is not running the common agent container, issue the `/usr/sbin/cacaoadm start` command on that node.

If a `geosite join` operation fails due to a timeout, re-issue the command with a `joinTimeout` property value larger than the default of 30. For example, `geosite join cluster -p joinTimeout=120 site`.

Changing the Site Role of a Site Cluster

A cluster that is part of a site can have one of two roles, controller or member:

- A cluster that has the controller role for a site is authorized to perform operations on the other site members and the multigroups they manage. The controller authorization is based on mutual agreement, any cluster can propose itself as a controller, but it can only control those clusters which have then agreed to accept its authority.
- A cluster that is a simple site member has agreed to accept commands from the site controllers, but cannot issue commands with site-wide effect.

This section provides the following procedures to change the role of a site cluster:

- [“How to Change a Site Membership Role” on page 89](#)
- [“How to Forcibly Change a Site Member to a Site Controller” on page 90](#)

▼ How to Change a Site Membership Role

Perform this procedure to change the role of a site cluster between member and controller.

Note - If the site has no reachable site controller, instead perform [“How to Forcibly Change a Site Member to a Site Controller” on page 90](#).

1. **From a node of a site controller cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Change the site role of the cluster.**

- **To change a site member to a site controller, use the following command.**

```
site-controller# geosite set-role -c cluster site
```

- **To change a site controller to a site member, use the following command.**

```
site-controller# geosite set-role -m cluster site
```

-c

Specifies the site controller role.

-m

Specifies the site member role.

cluster

The name of the cluster to change roles. You can specify more than one cluster from the same site in a single `geosite set-role` command by using the `-c` or `-m` option for each cluster that you want to change roles.

site

The name of the site with which the cluster has membership.

▼ How to Forcibly Change a Site Member to a Site Controller

Perform this procedure to change the role of a site member to a site controller when no existing site controller is accessible. There must be at least one other accessible member in the site.

The site member submits itself or another site member as a candidate for site controller. Each site member cluster must individually accept the candidate cluster as a new controller. No controller operations that are issued by the candidate cluster are obeyed by a site member cluster that has not yet accepted the candidate cluster as a controller.

If a site controller cluster was unavailable at the time a site member cluster submits itself to be a controller, the controller automatically accepts and synchronizes the site configuration change when it again becomes available.

If the site member is the only cluster in the site, the controller role is immediately assigned to the cluster. If later another cluster is added to the site, the new cluster must then accept the existing cluster's state as a site controller.

1. **From a node of the site member cluster that you want to make a site controller, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwX:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Submit the request to become a site controller.

```
candidate-cluster-node# geosite set-role -c candidate-cluster site
```

```
-c candidate-cluster
```

The name of the cluster that is requesting to be made a site controller. The candidate cluster can be the issuing cluster or another site member cluster.

```
site
```

The name of the site for which the candidate cluster is requesting the site controller role.

3. From another member cluster in the site, accept the request from the candidate cluster.

A site member cluster in the site must accept the role change request before the change will take effect for that cluster.

Note - If the site has no other clusters, omit the remaining steps in this procedure.

```
site-member-cluster-node# geosite accept candidate-cluster site
```

```
candidate-cluster
```

The name of the cluster that requested to be made a controller of the specified site.

4. Verify on the accepting cluster that the candidate cluster is now a site controller.

```
site member-cluster-node# geosite show -v site
```

```
...
```

```
candidate-cluster controller
```

5. Repeat Step 3 and Step 4 for each remaining cluster in the site.

Troubleshooting If more than one site member cluster proposes itself as a site controller and different site clusters accept a different new site controller, this creates a conflict for control of the site. To avoid such conflicts, promote only one cluster as a site controller and accept it on all accessible site clusters before you promote another cluster as a new site controller.

If a conflict for control of the site does occur, identify which cluster should be the site controller, then issue the `geosite update desired-site-controller site` command from each accessible site cluster that is in a synchronization error with that cluster.

If any cluster is down at the time that a site cluster is promoted to site controller, after the site cluster returns to service, ensure that it has no synchronization conflicts and, if necessary, use the `geosite update` command to resolve any conflicts.

Synchronizing Site Configuration Information

Under normal circumstances, a membership change or a change to a site property is automatically propagated to all clusters in the site.

This section provides procedures to manually update site configuration information on the local cluster with configuration information from the remote site cluster. The local site configuration information is overwritten by information from the remote cluster.

▼ How to Synchronize Site Configuration Information

Perform this procedure to update site configuration information on the local cluster.

1. **From a node of a site controller cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Synchronize the site configuration information with the remote site cluster.**

Site configuration information from the remote site cluster overwrites the site configuration information on the local cluster.

```
local-site-cluster-node# geosite update remote-site-cluster site
```

```
remote-site-cluster
```

The name of the remote site cluster whose site configuration information to import to the local site cluster. Both clusters must be configured in the same site.

```
site
```

The name of the site.

See Also For details about the synchronization statuses, see the Synchronization Status section of the [geoadm\(1M\)](#) man page.

Removing a Cluster From a Site

This section contains the following procedures to remove a cluster from a site:

- [“How to Remove a Cluster From a Site” on page 93](#)
- [“How to Remove an Unreachable Cluster From a Site a Cluster” on page 94](#)

▼ How to Remove a Cluster From a Site

Perform this procedure to make a cluster remove itself from a site. If the cluster you want to remove from a site is unreachable, instead follow "How to Remove From a Site a Cluster That Is Unreachable".

1. **From a node of the cluster that you want to remove from a site, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.** For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Detach the issuing cluster from the site.**

- **If the cluster is not the last controller in the site, use the following command.**

```
node-of-cluster-to-remove# geosite leave site
```

site

The name of the site to remove the cluster from.

- **If the cluster is the last controller in the site, include the -f option in the command.**

```
node-of-cluster-to-remove# geosite leave -f site
```

- If the site contains other member clusters, the site continues to exist without a site controller.
- If the site has no other member clusters, removal of the last controller also deletes the site.

Note - If the last site member being removed is not a site controller, the remove command fails with an error.

3. **From a remaining site cluster, verify that the removed cluster is no longer listed as a site member.**

```
# geosite show -v site
```

```
*** Site "site" is not configured ***
```

▼ How to Remove an Unreachable Cluster From a Site a Cluster

Perform this procedure to remove from a site a cluster that is unreachable. If a cluster you want to remove is reachable, instead follow "How to Remove a Cluster From a Site".

Note - To remove the last cluster from a site, instead follow [“How to Delete a Site” on page 95](#).

1. **From a node of a site controller cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Remove the cluster from the site.

```
site-controller-cluster-node# geosite remove-member cluster site
```

```
cluster
```

The name of the cluster to remove from the site. You can remove multiple clusters by specifying each cluster name separated by a comma (,).

3. From a remaining site cluster, verify that the removed cluster is no longer listed as a site member.

```
# geosite show -v site
```

Deleting a Site

This section contains procedures to delete a site. A site is automatically deleted when its last cluster is detached from the site. Each cluster must remove itself from the site. Because the last cluster to remove must be a site controller, first detach all clusters except one site controller. If the last cluster you remove is not a site controller, the command fails with an error.

When a site is deleted, all multigroups that the site references are automatically deleted on all clusters where the site was defined. The deletion of a multigroup has no effect on the protection groups that were configured in the deleted multigroup.

▼ How to Delete a Site

This procedure deletes a site by removing all of its member clusters. Perform this procedure from one site controller cluster.

Before You Begin Ensure that the site has at least one site controller cluster. If the site has no site controller cluster, change the role of a site member cluster to site controller before you perform this procedure. See [“How to Forcibly Change a Site Member to a Site Controller” on page 90](#).

1. **From a node of a site controller cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwX:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Remove all other clusters from the site.**

```
# geosite remove-member list-of-members site
```

list-of-members

A comma-separated list of each cluster in the site, excluding the local cluster.

site

The name of the site to detach the specified clusters from.

3. **Verify that the local cluster is the only remaining member of the site.**

```
# geosite show -v site  
last-site-controller-cluster controller
```

4. **Forcibly detach the local cluster from the site.**

```
# geosite leave -f site
```

The site is automatically deleted when the last site cluster is removed.

5. **Verify that the site no longer exists.**

```
# geosite show -v site  
  
*** Site "site" is not configured ***
```


Administering Multigroups

This chapter contains the following information about administering multigroups:

- [“Modifying Protection Groups in a Multigroup” on page 97](#)
- [“Starting and Stopping Multigroups” on page 101](#)
- [“Synchronizing Multigroup Configuration Information” on page 103](#)
- [“Deleting a Multigroup” on page 104](#)

For procedures to create a new multigroup, see [“Configuring Sites and Multigroups”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

For procedures to switch over or take over multigroups, see [Chapter 11, “Migrating Services”](#).

Modifying Protection Groups in a Multigroup

When you make a change to a multigroup, the site controller cluster that issues the change propagates the change to all other clusters in the site where the changed multigroup is active.

This section provides the following procedures to add, change, remove, and synchronize protection groups in a multigroup:

- [“How to Add a Protection Group to a Multigroup” on page 97](#)
- [“How to Change the Protection Groups in a Multigroup Dependency Chain” on page 99](#)
- [“How to Remove a Protection Group From a Multigroup” on page 100](#)
- [“How to Synchronize Multigroup Configuration Information” on page 103](#)

▼ How to Add a Protection Group to a Multigroup

Perform this procedure to add one or more protection groups to a multigroup.

Before You Begin Ensure that a partner cluster of the protection group you are adding is configured in the site.

1. **From a node of a controller cluster for the site to configure with the new multigroup, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [Chapter 4, “Administering RBAC”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwX:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Add the protection group to the multigroup.**

```
site-controller-cluster-node# geomg add-protection-group protection-group-list multigroup
```

The following describes the syntax choices for *protection-group-list*:

cluster:protection-group

Specifies a single protection group: The colon (:) separates the cluster name *cluster* from the name of the protection group that is configured in that cluster.

cluster:protection-group/cluster:protection-group

Specifies a protection group that has a dependency on another protection group, called a *dependency chain*: The protection group that is specified before the slash (/) in the dependency chain depends on the protection group that is specified after the slash.

cluster1:protection-group1,cluster1:protection-group2,cluster2:protection-group1/cluster3:protection-group1

The comma (,) separates multiple protection group names in the protection-group list.

(cluster1:protection-group2,cluster2:protection-group1)/cluster3:protection-group1

Specifies that multiple protection groups, *cluster1:protection-group2* and *cluster2:protection-group1*, all have a dependency on the *cluster3:protection-group1* protection group. In this form of dependency chain, parentheses are only used to enclose the multiple protection groups that have a dependency on another, single protection group. Only one protection group can be specified as the depended-on protection group in the dependency chain.

For more information, see the [geomg\(1M\)](#) man page.

3. **Verify the addition of the protection groups.**

```
site-controller-cluster-node# geomg show multigroup
```

▼ How to Change the Protection Groups in a Multigroup Dependency Chain

Perform this procedure to change a set of protection groups in a multigroup that are configured in a dependency to or from each other. Such interdependent protection groups are referred to as a *dependency chain*. For more information about configuring dependency chains, see [“How to Add a Protection Group to a Multigroup” on page 97](#).

You must first delete the entire dependency chain of protection groups from the multigroup. You then add back any protection groups in the deleted dependency chain that you want to continue as part of the multigroup.

This procedure does not affect the protection groups themselves.

1. **From a node of a site controller cluster, assume the root role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the root role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Remove the protection-group dependency chain from the multigroup.**

Specify each protection-group dependency chain that you want to change.

```
# geomg remove-protection-group protection-group-list multigroup
```

protection-group-list

Name of one or more interdependency protection-groups to remove from the multigroup. Specify the full protection-group dependency that contains the protection group that you want to remove.

```
# geomg remove-protection-group dependee-protection-group/depended-on-protection-group multigroup
```

To remove multiple protection groups, with and without a dependency relationship, separate each protection group name with a comma (,).

```
# geomg remove-protection-group protection-group,dependee-protection-group/depended-on-protection-group multigroup
```

multigroup

The name of the multigroup from which you are removing the protection groups.

3. **Add back those protection groups that you want to still be in the multigroup.**

```
# geomg add-protection-group protection-group-list multigroup
```

▼ How to Remove a Protection Group From a Multigroup

Perform this procedure to remove a protection group from a multigroup. This procedure does not affect the protection group.

Note - If the protection group to remove is part of set of interdependent protection groups, called a dependency chain, do not perform this procedure. Go instead to [“How to Change the Protection Groups in a Multigroup Dependency Chain” on page 99](#).

1. **From a node of a site controller cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Remove the protection group from the multigroup.**

- **If the protection group to remove is not part of a dependency chain with one or more other protection groups, use the following command:**

```
# geomg remove-protection-group cluster:protection-group multigroup
```

- **If the protection group to remove is part of a dependency chain with one or more other protection groups, use the following command to specify the complete dependency chain of protection groups:**

```
# geomg remove-protection-group protection-group-dependency-chain multigroup
```

cluster

Name of the cluster where the protection group to remove is configured.

protection-group-dependency-chain

Name of the complete dependency chain of protection groups that contains the protection group that you want to remove. For example, `newyork:pg1/newyork:pg2` is a dependency chain in which the `newyork:pg1` protection group depends on the `newyork:pg2` protection group.

To remove multiple protection groups, separate each protection group or dependency chain name with a comma (,).

3. **If a protection group that you removed was part of a dependency chain, add back any protection groups in the dependency chain that you want to continue to be in the multigroup.**

```
# geomg add-protection-group protection-group-list multigroup
```

Starting and Stopping Multigroups

This section contains the following procedures:

- [“How to Start All Protection Groups in a Multigroup” on page 101](#)
- [“How to Stop All Protection Groups in a Multigroup” on page 102](#)

▼ How to Start All Protection Groups in a Multigroup

Perform this procedure to start all protection groups that are configured a multigroup. The protection groups to start can be either those on only the clusters that are specified in the protection-group list of the multigroup or those on both partner clusters where the protection groups are configured.

1. **From a node of a site controller cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Start all protection groups in the multigroup.

- To start all protection groups on only the clusters specified in the protection-group list of the multigroup, include the `-e local` option in the command.

```
# geomg start -e local multigroup
```

- To start all protection groups on both partner clusters where the protection groups are configured, include the `-e global` option in the command.

```
# geomg start -e global multigroup
```

See the [geomg\(1M\)](#) man page for information about additional options for the `start` subcommand.

▼ How to Stop All Protection Groups in a Multigroup

Perform this procedure to stop all protection groups that are configured a multigroup. The protection groups to stop can be either those on only the clusters that are specified in the protection-group list of the multigroup or those on both partner clusters where the protection groups are configured.

1. From a node of a site controller cluster, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.

For more information about RBAC, see “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Stop all protection groups in the multigroup.

- To stop all protection groups on only the clusters specified in the protection-group list of the multigroup, include the `-e local` option in the command.

```
# geomg stop -e local multigroup
```

- To stop all protection groups on both partner clusters where the protection groups are configured, include the `-e global` option in the command.

```
# geomg stop -e global multigroup
```

See the [geomg\(1M\)](#) man page for information about additional options for the `stop` subcommand.

Synchronizing Multigroup Configuration Information

Under normal circumstances, information about a multigroup configuration change is automatically synchronized among all clusters in the multigroup.

This section provides procedures to manually update the locally known configuration information of a multigroup with the configuration information known to another cluster in the site. Synchronization of a multigroup is normally performed automatically after you add or remove a protection group. You would perform this procedure if synchronization failed for one or more clusters of the multigroup site. The configuration information known to the local cluster is overwritten by information from the specified remote cluster in the site.

▼ How to Synchronize Multigroup Configuration Information

Perform this procedure to manually synchronize multigroup configuration information within a site.

1. **From a node of a controller cluster of the local site, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Synchronize the multigroup configuration information with the remote cluster in the site.**

```
# geomg update remote-site-cluster multigroup-name
```

For more information about multigroup synchronization states, see the [`geoadm\(1M\)`](#) man page.

Deleting a Multigroup

This section describes procedures to delete a multigroup. This operation has no effect on the protection groups that are configured in the deleted multigroup.

Alternatively, if you intend to delete a site that references the multigroup you want to delete, instead follow [“How to Delete a Site”](#) on page 95. When you delete a site, all multigroups that the site references are automatically deleted as well.

▼ How to Delete a Multigroup

Perform this procedure to delete a multigroup from all clusters of a site where the multigroup exists.

1. **From a node of a controller cluster of the site where the multigroup exists, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Delete the multigroup.

```
controller-cluster-node# geomg delete multigroup
```

3. Verify that the multigroup is deleted.

```
controller-cluster-node# geomg show multigroup
```


Monitoring and Validating the Geographic Edition Software

This chapter describes the files and tools that you can use to monitor and validate the Geographic Edition software.

This chapter contains the following sections:

- [“Monitoring the Runtime Status of the Geographic Edition Software” on page 107](#)
- [“Viewing the Geographic Edition Log Messages” on page 114](#)
- [“Displaying Configuration Information for Partnerships and Protection Groups” on page 114](#)

Also see the applicable Geographic Edition data replication guide for additional procedures to check the runtime status of that particular type of protection group.

Monitoring the Runtime Status of the Geographic Edition Software

You can display the runtime status of the local 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 about all partnerships, 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
- **Site** – Provides information about all sites, including the names of the sites, the name and role of each site cluster, site heartbeat, and site heartbeat plug-in,
- **Multigroup** – Provides information about all multigroups, including the names of the multigroups, the name of associated sites, and synchronization status
- **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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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
  Plug-in "ping_plugin" : Inactive
  Plug-in "tcp_udp_plugin" : OK

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

Cluster cluster-paris      : OK
  Role                       : Primary
  Activation state           : Deactivated
  Configuration              : OK
  Data replication           : OK
  Resource groups            : None

Cluster cluster-newyork    : OK
  Role                       : Secondary
  Activation state           : Deactivated
  Configuration              : OK
  Data replication           : OK
  Resource groups            : None

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

Cluster cluster-paris      : OK
  Role                       : Primary
  Activation state           : Deactivated
  Configuration              : OK
  Data replication           : OK
  Resource groups            : None

Cluster cluster-newyork    : OK
  Role                       : Secondary
  Activation state           : Deactivated
  Configuration              : OK
  Data replication           : OK
```

```

Resource groups      : None

Site : site1

Controller "cluster-paris"
  Configuration      : OK

Controller "cluster-madrid"
  Configuration      : OK
  Synchronization    : OK

Heartbeat "paris-to-madrid-site1-cluster-madrid" monitoring "cluster-madrid": OK
  Plug-in "tcp_udp_plugin" : OK

Member "cluster-london"
  Configuration      : OK
  Synchronization    : OK

Heartbeat "paris-to-london-site1-cluster-london" monitoring "cluster-london": OK
  Plug-in "tcp_udp_plugin" : OK

Multigroup "mg2"    :
  Site            : site1
  Configuration    : OK
  Synchronization with cluster cluster-madrid : OK
  Synchronization with cluster cluster-london : OK

Multigroup "mg1"    :
  Site            : site1
  Configuration    : OK
  Synchronization with cluster cluster-madrid : OK
  Synchronization with cluster cluster-london : OK

Pending operations:

Protection group "tcpg" operation: Start

Pending multigroup operations:

Multigroup "mg1" operation: Start

```

The information displayed shows that the protection group, `tcpg`, is started 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 10-1 Status Value Descriptions

Field	Value Descriptions
Partnership	OK – The partners are connected. Error – The connection between the partner clusters is lost.

Field	Value Descriptions
	<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> <ul style="list-style-type: none"> ■ For information about resynchronizing a partnership, see “Resynchronizing a Partnership” on page 49. ■ For information about resynchronizing a protection group, see “Resynchronizing a Protection Group” on page 78. <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>geogg 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>
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 heartbeat plug-ins is degraded or is 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 – No component of the protection group on either partner is in the Degraded, Error, or Unknown status, and the protection group configuration is the same on both partner clusters.</p> <p>Degraded – Data replication is either not running or is in a partial error state.</p> <p>Error – At least one component of the partnership is in an error state on at least one partner, or the protection group configuration is different between the partner clusters.</p> <p>Unknown – The status for at least one component of the protection group is unknown, or the status of the protection group is not accessible.</p>

Field	Value Descriptions
Protection group > Synchronization (state of protection group configuration information between partner clusters)	<p>OK – The configuration is synchronized between partner clusters.</p> <p>Error – The configuration on the partner clusters is different. You must synchronize the protection group again.</p> <p>Mismatch – The protection group has been configured on each partner cluster individually. You must remove the configuration from one cluster and copy the configuration of the partner cluster.</p> <p>Unknown – Information is not accessible because the partners are disconnected.</p>
Protection group > Cluster (state of protection group on each cluster)	<p>None – The data replication or resource group component is not configured in the protection group.</p> <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>Error – The state of some components of the protection group, such as configuration data, data replication, or resource groups, is in Error.</p> <p>Unknown – The state of some components of the protection group, such as configuration data, data replication, or resource groups, is unavailable.</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>
Protection group > Cluster > 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 “Validating a Protection Group” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”.</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 > 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>

Field	Value Descriptions
	<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 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>
Site > Controller or Member > Configuration	<p>OK – The site configuration is correct.</p> <p>Error – The site configuration has an error.</p> <p>Unknown – The site configuration is not checked.</p>
Site > Controller or Member > Synchronization	<p>DIFFERENT – Site configuration information on two compared clusters is different but compatible. This is a transitory status. Eventually, the older of the two configurations is automatically synchronized with the newer configuration, at which time the synchronization status becomes OK. If for some reason the status remains at DIFFERENT, use the <code>geosite validate</code> subcommand to synchronize the site configurations.</p> <p>ERROR – The site configuration on two compared clusters is different and cannot be automatically resolved. Configuration changes from an issuing controller cluster are not accepted by a cluster while it is in the ERROR synchronization status with the issuing controller. If configuration information between two clusters do not resolve automatically and the clusters are in the ERROR status, use the <code>geosite update</code> subcommand to resolve site configuration conflicts between the two clusters.</p> <p>OK – Site configuration information matches on the compared clusters.</p>

Field	Value Descriptions
	<p>UNKNOWN – Site configuration information cannot be compared because the Geographic Edition software cannot reach a cluster. Configuration changes from an issuing controller cluster are not propagated to a cluster while it is in the UNKNOWN synchronization status. If the cluster is not automatically synchronized, use the <code>geosite validate</code> subcommand to update the cluster with the latest site configuration information.</p>
Site > Controller or Member > Heartbeat	<p>OK – Heartbeat monitoring is enabled, and the partner cluster is responding within timeout and retry periods.</p> <p>Degraded – Heartbeat checks are running but one of the primary heartbeat plug-ins is degraded or is not running.</p> <p>Error – Heartbeat monitoring is running but the partner cluster is not responding and retries have timed out.</p> <p>Offline – Heartbeat monitoring is offline.</p>
Site > Controller or Member > Heartbeat > Plug-in	<p>OK – The partner cluster is responding.</p> <p>Inactive – The plug-in is not in use. It is a standby plug-in that is used for retrying if other plug-ins do not respond.</p> <p>No-response – The partner cluster is not responding.</p>
Multigroup > Configuration	<p>OK – The configuration of the multigroup is correct and validated without errors.</p> <p>Error – The configuration of the multigroup has errors. Run the <code>geomg validate</code> command on the multigroup to identify the error.</p> <p>Unknown – The configuration is not accessible.</p>
Multigroup > Synchronization	<p>DIFFERENT – Multigroup configuration information on two compared clusters is different but compatible. This is a transitory status. Eventually, the older of the two configurations is automatically synchronized with the newer configuration, at which time the synchronization status becomes OK. If for some reason the status remains at DIFFERENT, use the <code>geomg validate</code> command to synchronize the multigroup configurations.</p> <p>ERROR – The multigroup configuration on two compared clusters is different and cannot be automatically resolved. Configuration changes from an issuing controller cluster are not accepted by a cluster while it is in the ERROR synchronization status with the issuing controller. If configuration information between two clusters do not resolve automatically and the clusters are in the ERROR status, use the <code>geomg update</code> command to resolve multigroup configuration conflicts between the two clusters.</p> <p>OK – Multigroup configuration information matches on the compared clusters.</p> <p>UNKNOWN – Multigroup configuration information cannot be compared because the Geographic Edition software cannot reach a cluster. Configuration changes from an issuing controller cluster are not propagated to a cluster while it is in the UNKNOWN synchronization status. If the cluster is not automatically synchronized after the Geographic Edition framework is started on that cluster, use the <code>geomg validate</code> command to update the cluster with the latest multigroup configuration information.</p>

Viewing the Geographic Edition Log Messages

All the Geographic Edition components produce messages that are stored in log files.

Information about the loading, running, and stopping 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 Geographic Edition modules are updated on the node where the 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`.

For data replication modules that are based on script-based plug-ins, you can set the `DEBUG` property to `TRUE` for more verbose messages to aid in troubleshooting.

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.

This section provides the following procedures:

- [“How to Display Configuration Information About Partnerships” on page 115](#)
- [“How to Display Configuration Information About Protection Groups” on page 115](#)

▼ How to Display Configuration Information About Partnerships

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, then click the partnership name. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Display information about the partnership.

```
# geops list partnership-name
```

partnership-name

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

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

Example 10-1 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

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name to go to its page, then in the Protection Groups section click the protection group name. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

2. Display information about a protection group.

```
# geopg list [protection-group]
```

protection-group

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 10-2 Displaying Configuration Information About a Protection Group

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

```
# geopg list avspg
```

Migrating Services

This chapter provides information about detecting cluster failure and moving services to an accessible cluster.

- [“Detecting Cluster Failure” on page 117](#)
- [“Migrating Replication Services by Switching Over Protection Groups” on page 118](#)
- [“Forcing a Takeover of a Protection Group” on page 123](#)
- [“Recovering a Protection Group to a Cluster” on page 127](#)

Detecting Cluster Failure

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

- [“Detecting Primary Cluster Failure” on page 117](#)
- [“Detecting Secondary Cluster Failure” on page 118](#)

Detecting Primary Cluster Failure

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

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

- Heartbeat failure is detected by a partner cluster.
- The heartbeat is activated in emergency mode to verify that the heartbeat loss is not transient and that the primary cluster has failed. The heartbeat remains in the `Online` state during this default time-out interval, while the heartbeat mechanism continues to retry the primary cluster.

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

For more information about logging, see [“Viewing the Geographic Edition Log Messages” on page 114](#).

- If the partnership is configured for heartbeat-loss notification, then one or both of the following actions occurs:
 - An email is sent to the address specified in the `Notification_emailaddr`s property.
 - The script defined in `Notification_actioncmd` is executed.

For more information about configuring heartbeat-loss notification, see [“Configuring Heartbeat-Loss Notification” on page 64](#).

Detecting Secondary Cluster Failure

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

During failure detection, the following actions take place:

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

Migrating Replication Services by Switching Over Protection Groups

Perform a switchover of a protection group when you want to migrate services to the partner cluster in an orderly fashion. You can switch over an individual protection group or switch over a multigroup of multiple protection groups in a single operation.

This section contains the following information:

- [“Actions Performed by the Geographic Edition Software During a Switchover” on page 119](#)

- [“How to Switch Over Replication From the Primary Cluster to the Secondary Cluster” on page 120](#)
- [“How to Switch Over a Multigroup” on page 122](#)

Actions Performed by the Geographic Edition Software During a Switchover

A switchover consists of the following actions:

- Application services are unmanaged on the former primary cluster.
- The replication role is reversed and now continues to run from the new primary.
- Application services are brought online on the new primary cluster.

When you run the `geopg switchover` command, the software confirms that the secondary cluster does indeed allow the replication role reversal. The command creates copies of the replicated source projects and exports the shares for use on the target appliance to ensure that no projects and mount point conflicts can occur. The command then destroys such created clones on the target appliance and confirms that the actual reverse replication operation can be performed on the target appliance.

Note - When migrating EMC SRDF replication services, basic Geographic Edition operations such as `geopg switchover` perform a `symrdf swap` operation. The `symrdf swap` operation requires significantly more time for static RDF than dynamic RDF. Therefore, you might need to increase the value of the timeout property of the protection group when using static RDF.

For Oracle Data Guard, before the switchover, the command checks that the remote database is in an `enabled` state in the Oracle Data Guard Broker configuration. The command also confirms that the configuration is healthy by issuing the Oracle Data Guard command-line interface (`dgmgrl`) `show configuration` command to ensure that the command returns a `SUCCESS` state. If the output from this command indicates that Oracle Data Guard Broker is busy performing its own health check, the Oracle Data Guard command-line interface retries the command until it receives a `SUCCESS` response or until two minutes have passed. If the command-line interface is unable to get a `SUCCESS` response, the command fails.

For Oracle ZFS Storage Appliance, the command creates copies of the replicated source projects and exports the shares for use on the target appliance to ensure that no projects and mount point conflicts can occur. The command then destroys such created clones on the target appliance and confirms that the actual reverse replication operation can be performed on the target appliance.

The software performs the following actions on the original primary cluster:

- For Availability Suite, removes affinities and resource dependencies between all the application resource groups in the protection group and the internal resource group, such as the lightweight resource groups
- Takes offline the application resource groups in the protection group and places them in the Unmanaged state.
- Performs a switchover to the secondary cluster for each replication configuration in the protection group.

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

- Runs the script that is defined in the `RoleChange_ActionCmd` property
- Brings online all application resource groups in the protection group

If the command completes successfully, the secondary cluster becomes the new primary cluster for the protection group. The original primary cluster becomes the new secondary cluster. The application resource groups in the protection group are brought online on the new primary cluster and replication from the appliance that is connected from the new primary cluster to the new secondary cluster begins.

For Oracle Data Guard, databases that are associated with the Oracle Data Guard Broker configurations of the protection group have their role reversed according to the role of the protection group on the local cluster. For HA for Oracle configurations, the `dataguard_role` resource property is also updated with the status of the new primary and standby clusters. The shadow Oracle database server resource group and any other application resource groups in the protection group are online on the new primary cluster. Data replication from the new primary cluster to the new standby cluster begins.

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

If the Configuration status of the protection group is set to `Error`, revalidate the protection group by using the procedures that are described in [“Validating a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures that are described in [Resynchronizing a Protection Group](#).

▼ How to Switch Over Replication From the Primary Cluster to the Secondary Cluster

Perform this procedure to switch over a protection group to the secondary cluster.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI, if you do not need to use the `-f` force option. Click Partnerships, click the partnership name to go to its page, highlight the protection group name, and click `Switchover`. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Note - To switch over a set of protection groups that are configured as a multigroup, instead follow procedures in [“How to Switch Over a Multigroup”](#) on page 122.

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

- For Oracle ZFS Storage Appliance, data replication is active between the primary cluster and the secondary cluster. That is, the replication is enabled from the source to the target appliances.
- For Oracle ZFS Storage Appliance, the Geographic Edition replication resource for this appliance replication shows the `Online` state.
- For Oracle Data Guard, the Oracle Data Guard Broker `show configuration` command must show a `SUCCESS` state. This state is reflected in the state of the Geographic Edition replication resource for this Oracle Data Guard Broker configuration, which should show the `online` state.
- The Geographic Edition software is running on both clusters.
- The secondary cluster is a member of a partnership.
- Both cluster partners can be reached.
- The overall state of the protection group is set to `OK`.



Caution - For EMC SRDF, if you have configured the `cluster_dgs` property, only applications that belong to the protection group can write to the device groups specified in the `Cluster_dgs` property.

1. **Assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and Oracle ZFS Storage Appliance software.

2. Initiate the switchover.

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

```
phys-paris-1# geopg switchover [-f] -m new-primary-cluster protection-group-name
```

`-f`

Forces the command to perform the operation without asking you for confirmation.

`-m new-primary-cluster`

Specifies the name of the cluster that is to be the new primary cluster for the protection group.

protection-group-name

Specifies the name of the protection group.

Example 11-1 Forcing a Switchover From the Primary Cluster to the Secondary Cluster

The following example performs a switchover to the secondary cluster `cluster-newyork`.

```
phys-paris-1# geopg switchover -f -m cluster-newyork example-protection-group
```

Next Steps To fail back a protection group to the original cluster, see [“Recovering a Protection Group to a Cluster” on page 127](#).

▼ How to Switch Over a Multigroup

Perform this procedure to switch operation of all protection groups in a multigroup to partner clusters in the specified site.

Note - To switch over an individual protection group, instead follow procedures in [“How to Switch Over Replication From the Primary Cluster to the Secondary Cluster” on page 120](#).

Before You Begin Ensure that the multigroup is started.

1. **From a node of a controller cluster in the target site, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.**

For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwx:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. **Switch the multigroup to the partner clusters in the specified site.**

```
# geomg switchover -s site multigroup
```

`-s site`

The name of the site that the multigroup is to switch to.

`multigroup`

The name of the multigroup to switch over.

See the [geomg\(1M\)](#) man page for information about additional options for the `switchover` subcommand.

Next Steps To fail back a protection group to the original cluster, see [“Recovering a Protection Group to a Cluster”](#) on page 127.

Forcing a Takeover of a Protection Group

Perform a takeover when applications need to be brought online on the secondary cluster, regardless of whether the data is completely consistent between the primary and the secondary. You can take over an individual protection group or take over multiple protection groups that are configured as a multigroup in a single operation.

The information in this section assumes that the protection group has been started.

This section contains the following information:

- [“Actions Performed by the Geographic Edition Software After a Takeover”](#) on page 124

- [“How to Force Immediate Takeover of a Protection Group by a Secondary Cluster” on page 124](#)
- [“How to Take Over a Multigroup” on page 126](#)

Actions Performed by the Geographic Edition Software After a Takeover

The following steps occur after you initiate a takeover:

- If the former primary cluster can be reached and the protection group is not locked for notification handling or some other reason, the protection group is deactivated.
- The data of the former primary cluster `cluster-paris` is taken over by the new primary cluster.

Note - The data on the former primary cluster `cluster-parisC` might not be consistent with the data on the original primary cluster. Data replication from the new primary cluster to the former primary cluster is stopped.

- Application services are brought online on the new primary cluster.
- The protection group is activated without replication enabled.

For more details about takeover and the effects of the `geopg takeover` command, see [Appendix C, Disaster Recovery Administration Example](#).

For details about the possible conditions of the primary and secondary cluster before and after a takeover, see [Appendix D, Takeover Postconditions](#).

▼ How to Force Immediate Takeover of a Protection Group by a Secondary Cluster

Perform this procedure to take over an individual protection group by the secondary cluster.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI, if you do not need to use the `-f` force option. Click Partnerships, click the partnership name to go to its page, highlight the protection group name, and click Takeover. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,” in “Oracle Solaris Cluster System Administration Guide”](#).

Note - To take over a set of protection groups that are configured as a multigroup, instead follow procedures in [“How to Take Over a Multigroup” on page 126](#).

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

- The Geographic Edition software is up and running on the secondary cluster.
- The secondary cluster is a member of a partnership.
- The Configuration status of the protection group is OK on the secondary cluster.
- If you are using EMC Symmetrix Remote Data Facility, ensure that no LUNs in the data replication device group are in the SyncInProg state. Otherwise, the application might fail to start on the new primary cluster due to data inconsistencies.

Perform this procedure from a node in the secondary cluster.

1. Assume the root role or assume a role that is assigned the Geo Management RBAC rights profile.

For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the root role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rxw:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and the data replication software.

2. Initiate the takeover.

```
phys-newyork-1# geopg takeover [-f] protection-group-name
```

`-f`

Forces the command to perform the operation without your confirmation, even if the primary cluster for a protection group in the multigroup is reachable and the protection group is active on that cluster.

protection-group-name

Specifies the name of the protection group.

Example 11-2 Forcing a Takeover by a Secondary Cluster

The following example forces the takeover of the protection group `example-pg` by the secondary cluster `cluster-newyork`.

The node `phys-newyork-1` is the first node of the secondary cluster. For a reminder of which node is `phys-newyork-1`, see [Example Geographic Edition Cluster Configuration](#).

```
phys-newyork-1# geopg takeover -f example-pg
```

Next Steps For information about the state of the primary and secondary clusters after a takeover, see [Appendix D, Takeover Postconditions](#).

To fail back a protection group to the original cluster, see [“Recovering a Protection Group to a Cluster” on page 127](#).

▼ How to Take Over a Multigroup

Perform this procedure to transfer all operation of protection groups in a multigroup to partner clusters in the specified site.

Note - To take over a protection groups that is not configured as a multigroup, instead follow procedures in [“How to Force Immediate Takeover of a Protection Group by a Secondary Cluster” on page 124](#).



Caution - To avoid possible loss of replicated data, use instead the `geomg switchover` command. The `geomg takeover` command is intended only for situations where the potential risk of data loss is justified, such as when the primary site is not available or during unplanned downtime.

1. **From a node of a controller cluster of the target site, assume the `root` role or assume a role that is assigned the Geo Management RBAC rights profile.** For more information about RBAC, see [“Securing Geographic Edition Software” in “Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

Note - If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.

```
# chmod A+user:username:rwX:allow /var/cluster/geo
```

The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and data replication software.

2. Initiate takeover of the multigroup by the standby site.

```
# geomg takeover [-f] -s site multigroup
```

-f

Forces takeover, even if the primary cluster for a protection group in the multigroup is reachable and the protection group is active on that cluster.

-s *site*

The name of the standby site which is to take over the specified multigroup.

multigroup

The name of the multigroup to take over.

See the [geomg\(1M\)](#) man page for information about additional options for the takeover subcommand.

Next Steps For information about failing back a protection group to the original cluster, see [“Recovering a Protection Group to a Cluster” on page 127](#).

Recovering a Protection Group to a Cluster

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

The Geographic Edition software supports the following kinds of failback:

- **Failback-switchover.** During a failback-switchover, applications are brought online again on the original primary cluster after the data of the original primary cluster was resynchronized with the data on the secondary cluster.
- **Failback-takeover.** During a failback-takeover, applications are brought online again on the original primary cluster and use the current data on the original primary cluster. Any updates that occurred on the secondary cluster while it was acting as primary are discarded.

If you want to leave the new primary as the primary cluster and the original primary cluster as the secondary after the original primary restarts, you can resynchronize and revalidate the protection group configuration without performing a switchover or takeover.

For procedures to perform a failback-switchover or a failback-takeover, see the Geographic Edition documentation for the data replication product you are using.

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 129](#)
- [“Configuring a Protection Group to Run a Script at Switchover or Takeover” on page 132](#)

Creating a Role-Change Action Script

You can configure the 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.

Observe the following requirements for the role-change action script:

- The path to this script must be valid on all nodes of all partner clusters that can host the protection group.
- The script must have execute permissions for the user that launches the script. The intended user can be the root role, or an administrator with the necessary `solaris.cluster.*` authorization to execute a switchover or takeover operation plus any other actions that the script performs. For more information about user rights for Geographic Edition software, see [Chapter 4, “Administering RBAC”](#).
- The script must begin with the path to the shell, such as `#!/bin/ksh`.

The following command-line runs the script:

```
# custom-action-command-path -o primary -c cluster-name \
```

-s partnership-name protection-group user-arguments

custom-action-command-path

Specifies a path to the action command you have created.

-o primary

Specifies that the role being assumed by the cluster is primary.

-c cluster-name

Specifies the name of the secondary cluster that is assuming the new role of primary cluster.

-s partnership-name

Specifies the name of the partnership that hosts the protection group.

protection-group

Specifies the name of the protection group that is undergoing the role change.

user-arguments

Specifies static arguments that are passed after all the Geographic Edition supplied options.

This free-form string can be parsed by the script as required. For example, you could specify a list of key=value pairs, such as name=oracle.com, ip=10.1.2.3. You could also specify a sequence of options, such as -n oracle.com -a 10.1.2.3.4. The format of these arguments is not restricted by the 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 exit status of the action script can impact the `geomg switchover` or `geomg takeover` commands. During switchover or takeover of a multigroup, if the action script returns a nonzero exit status for a protection group that has dependents, the switchover or takeover of those dependent protection groups is not performed.

The 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 12-1 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"
# 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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

```
# geopg set-prop -p RoleChange_ActionCmd=fully-qualified-script \  
-p RoleChange_ActionArgs=script-arguments protection-group-name
```

-p property-setting

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 Geographic Edition Properties”](#).

protection-group

Specifies the name of the protection group.

Example 12-2 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
```

Script-Based Plug-Ins

This chapter provides information about Geographic Edition script-based plug-ins. It covers the following topics:

- [“Introduction to Geographic Edition Script-Based Plug-Ins” on page 133](#)
- [“Property Descriptions for Script-Based Plug-Ins” on page 137](#)
- [“Internals for Script-Based Plug-Ins” on page 148](#)

Introduction to Geographic Edition Script-Based Plug-Ins

Geographic Edition software provides modules to support data replication using the Availability Suite feature of Oracle Solaris, Oracle Data Guard, Oracle ZFS Storage Appliance, Oracle MySQL, and EMC SRDF. However, the creation of these modules requires detailed knowledge of both the replication software and the internals of the Geographic Edition product. Geographic Edition uses the common agent container with a number of Java management beans (MBeans) that form the interface for the Geographic Edition monitoring and management infrastructure and the replication control software.

By providing a more generic interface module analogous to the Oracle Solaris Cluster Generic Data Service (GDS), the Geographic Edition script-based plug-in enables you to rapidly integrate additional replication technologies by supplying a few interface scripts to fulfill the necessary control functions. This capability frees you from needing to learn the internals of Geographic Edition or needing any knowledge of Java technology or MBeans. Instead, you can focus on the replication technology you need to protect your enterprise data. For more information on the Generic Data Service, see [“Oracle Solaris Cluster Generic Data Service \(GDS\) Guide”](#).

For simplicity, the term script is used throughout this document to represent any compiled binary or script-based executable.

This section provides the following information:

- [“Advantages and Disadvantages of Using Script-Based Plug-Ins” on page 134](#)

- [“Script-Based Plug-In Architecture” on page 134](#)
- [“Restrictions of Script-Based Plug-Ins” on page 136](#)
- [“Creating Script-Based Plug-Ins” on page 136](#)

Advantages and Disadvantages of Using Script-Based Plug-Ins

The main advantage of using the script-based plug-in comes from reducing the barriers to implementing new replication mechanisms. Rather than spending time learning about the Java, JMX, MBeans, or common agent container technologies, you can focus on the critical logic needed, for example, to set up a replicated configuration or change the direction of the replication flow.

The disadvantage of this approach stems from the very generic nature of the plug-in that makes it so easy to use. Generic plug-ins lack some of the tight integration that a custom module can offer. For example, the arguments that you supply on the command line to a script-based plug-in configuration are at the script argument level rather than the highly specific replication variable level. So, whereas the Geographic Edition Oracle Data Guard module has separate, specific arguments for properties like `standby_type` and `replication_mode`, an equivalent script-based plug-in version would pass these properties and their value as part of a single, anonymous bundle to a script. The script would need to determine the arguments and whether each argument is valid.

Script-Based Plug-In Architecture

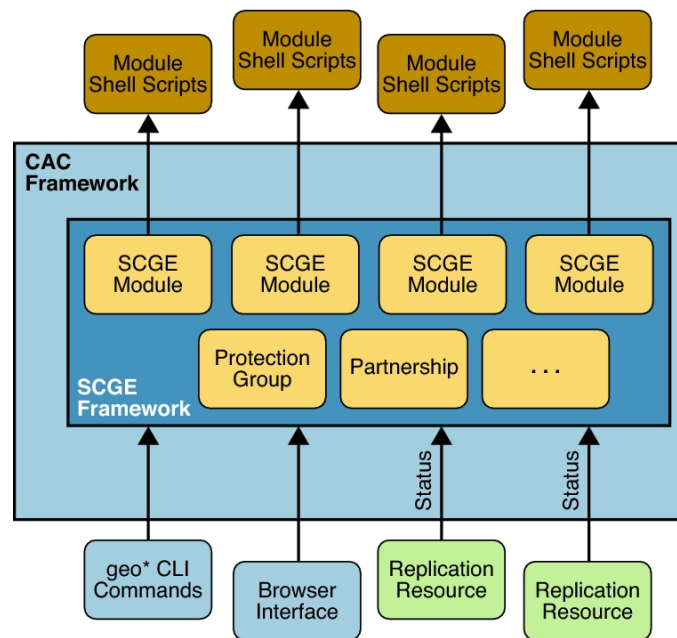
Unlike other replication modules, the script-based plug-in is generic and capable of supporting a wide range of replication technologies. Consequently, the script-based plug-in does not contain a specific set of scripts to control a particular piece of replication software. Instead, it provides a framework for integrating a set of scripts or programs that you, the developer, write and that a system administrator will later use.

This flexibility means that the script-based plug-in cannot directly enforce the inclusion or exclusion of application resource groups in a protection group. Furthermore, the script-based plug-in cannot even restrict the node lists of these entities, nor the relationship with the replication resource group that contains the replication resource needed to supply the replication status, or indeed any other resource group that is required.

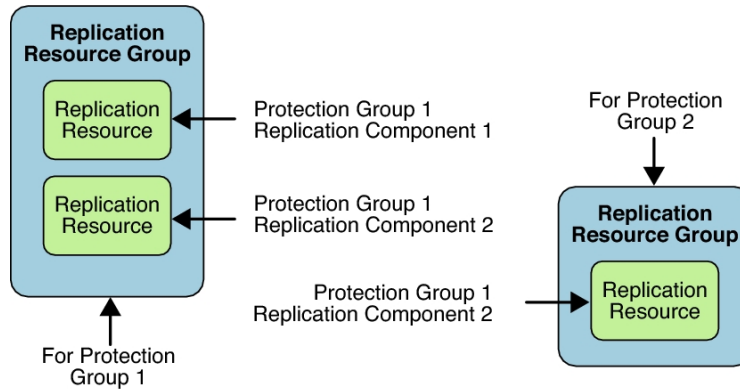
The following figure outlines the relationships between the various components within the Geographic Edition system. Commands issued through the command-line interface (CLI) call the Geographic Edition modules through their relevant common agent container modules. These modules then call out to shell scripts to perform specific tasks. Once a protection group has been instantiated, the replication resource, representing a particular replicated object entity,

reports its status back to the module through the event framework. This process enables the overall replication status to be reflected in the Geographic Edition output on the command line or in the browser interface.

FIGURE 13-1 Script-Based Plug-Ins Framework



The script-based plug-in developer therefore is free to govern the relationships between any or all of these entities: application resource group, data replication resource group, and replication status resource group. As the following figure shows, the only constraints are the requirements to have a named replication resource group per protection group and a named replication resource per device group or replicated component.

FIGURE 13-2 Script-Based Plug-In Replication Resource Group

The consequence of these requirements is that the administrator must provide script-based plug-in configuration file for each protection group that is accessible from all cluster nodes and that details which nodes pertain to each script-based plug-in configuration. The purpose of this configuration file is to ensure that any subsequent developer-written scripts are called on one or more nodes on which the service is present.

In addition to the standard protection properties, the script-based plug-in enables the developer to name one or more scripts to perform the actions required by the Geographic Edition framework. These actions fall into two separate groups: those actions that operate at a per protection level and those actions that operate at a per replicated component level.

Restrictions of Script-Based Plug-Ins

There are no inherent restrictions regarding what you can do when creating script-based plug-in modules. However, using the script-based plug-in does not enable you to circumvent or overcome any inherent limitations present in the replication technology you intend to use.

Creating Script-Based Plug-Ins

The scripts can be written using an integrated development environment (IDE), such as the NetBeans IDE. For more information on NetBeans IDE, see [NetBeans IDE \(https://netbeans.org/\)](https://netbeans.org/).

Property Descriptions for Script-Based Plug-Ins

This section contains the following information:

- [“Protection Group Properties - Overview” on page 137](#)
- [“Replicated Component Properties - Overview” on page 138](#)
- [“Protection Group Property Descriptions” on page 139](#)

Protection Group Properties - Overview

The table in this section lists the protection group properties, along with a brief description, type of property, and default value for each property.

The scripts named by the developer in these properties can reference independent executables, a single common executable, or a combination of the two. No restrictions are placed on the language used to implement these scripts with the exception that the scripts must be able to run by root, from the command line, without a graphical display, and they must return either a zero (success) or nonzero (failure) exit code. The script-based plug-in Mbean returns any error code resulting from a failure. For more information, see [Appendix F, “Error Return Codes for Script-Based Plug-Ins”](#).

Protection groups that use script-based plug-in replication have the global properties provided in the following table. Note that all of these properties are tunable when you are offline.

TABLE 13-1 Protection Group Global Policies

Property Name	Description	Type	Default Value
add_app_rg_args	The arguments that are provided to the script, add_app_rg_script.	Optional	Not applicable
add_app_rg_script	The script used to validate and perform tasks relevant for adding an application resource group to a protection group.	Required	/bin/true
configuration_file	The per protection group script-based plug-in configuration file containing details of the nodes pertinent to script-based plug-in replicated components held in the protection group.	Required	/etc/opt/SUNWscgrepsbp/configuration
create_config_script	The script used to create, modify, and validate	Required	/bin/false

	a script-based plug-in replicated component instance.		
remove_app_rg_args	The arguments that are provided to the script, remove_app_rg_script.	Optional	Not applicable
remove_app_rg_script	The script used to validate and perform tasks relevant for removing an application resource group from a protection group.	Required	/bin/true
remove_config_script	The script used to remove a script-based plug-in replicated component instance.	Required	/bin/true
start_replication_script	The script used to start the data replication for a script-based plug-in replicated component instance.	Required	/bin/true
stop_replication_script	The script used to stop the data replication for a script-based plug-in replicated component instance.	Required	/bin/true
switchover_script	The script used to switch over the data replication direction for a script-based plug-in replicated component instance.	Required	/bin/true
takeover_script	The script used to take over the data replication for a script-based plug-in replicated component instance.	Required	/bin/true

The [“Protection Group Property Descriptions” on page 139](#) section describes in detail the actions that each script and its associated arguments should perform when called by the script-based plug-in Mbean. [“Standardized Script Command-Line Arguments” on page 149](#) explains how scripts can discriminate between the steps being performed.

Replicated Component Properties - Overview

Each replication component added to a particular protection group uses the scripts named in [“Protection Group Properties - Overview” on page 137](#). Individual replications distinguish themselves by varying the properties passed to these scripts.

The script-based plug-in module provides for two site-specific password properties:

- A local service password property (local_service_password)

- A remote service password property (`remote_service_password`)

These properties enable administrators of a script-based plug-in deployment to supply passwords to log in to services or remote systems without having to provide these passwords at switchover or takeover time. For more information, see [“How Geographic Edition Handles Password Properties” on page 151](#).

The script-based plug-in module requires the developer to provide a property naming the replication resource contained in the replication resource group that holds the status of the replication.

Replicated components in script-based plug-in protection groups have the optional properties provided in the following table. Note that all of these properties are tunable when you are offline.

TABLE 13-2 Optional Replicated Component Properties

Property Name	Description	Type
<code>create_config_args</code>	The arguments passed to the script named by the <code>create_config_script</code> protection group property.	Global
<code>local_service_password</code>	A password that might be needed by the scripts to perform some function on the local system that requires the entry of a password.	Local
<code>remote_service_password</code>	A password that might be needed by the scripts to perform some function on the remote system that requires the entry of a password.	Local
<code>remove_config_args</code>	The arguments passed to the script named by the <code>remove_config_script</code> protection group property.	Global
<code>start_replication_args</code>	The arguments passed to the script named by the <code>start_replication_script</code> protection group property.	Global
<code>stop_replication_args</code>	The arguments passed to the script named by the <code>stop_replication_script</code> protection group property.	Global
<code>switchover_args</code>	The arguments passed to the script named by the <code>switchover_script</code> protection group property.	Global
<code>takeover_args</code>	The arguments passed to the script named by the <code>takeover_script</code> protection group property.	Global

Protection Group Property Descriptions

This section describes the following protection group properties:

- [“add_app_rg_script Property” on page 140](#)
- [“configuration_file Property” on page 141](#)
- [“create_config_script Property” on page 142](#)
- [“remove_app_rg_script Property” on page 143](#)

- [“remove_config_script Property” on page 144](#)
- [“start_replication_script Property” on page 145](#)
- [“stop_replication_script Property ” on page 145](#)
- [“switchover_script Property” on page 146](#)
- [“takeover_script Property” on page 147](#)

add_app_rg_script Property

The script referenced by the `add_app_rg_script` property is responsible for checking that one or more application resource groups selected by the administrator are suitable for addition to the protection group. These checks might require that certain resource types be present or absent. Furthermore, the script must also set up any resource group affinities or dependencies within the confines of what is allowed by Geographic Edition. These affinities or dependencies are needed for the application resource group to produce the correct behavior.

Application resource groups must be in the unmanaged state when they are added to the configuration.

The `add_app_rg_script` is called at other points within the protection group life cycle, not just on the addition of application resource groups, to ensure that application resource groups continue to conform to the required rules. The script should be written to ensure that these rules are met at all times.

Resource groups are offline and unmanaged on the standby site so certain application resource groups that represent services with embedded data replication might be unsuitable for addition to the protection group directly. An example is database data replication such as MySQL and Oracle RAC. The `add_app_rg_script` script must accommodate such validation.

The script must also be able to validate the `add_app_rg_args` property supplied to it with the `validate_parameters=true` option without actually performing any of the steps associated with this task. This operation is called only at the time of protection group update and creation, as opposed to at the time of device group update, modification, or validation.

When executed with `validate_parameters=false`, the script must perform any task required to add the resource groups listed in the final comma-separated `rgList` parameter. These actions might include altering one or more of these resource group properties. The script is called on the local cluster to where the `geopg add-resource-group` command is run and called asynchronously on the remote cluster in response to the internal application resource group table being updated.

For example, if `add_app_rg_script = /var/tmp/addRGs` and `add_app_rg_args = -u root -d /mydir`, the resulting command looks like the following example:

```
# /var/tmp/addRGs -u root -d /mydir function=add_application_rgs \
```

```

validate_parameters={true|false} \
currentRole={PRIMARY|SECONDARY} pg=protection-group \
rgList=resource-group1,resource-group2,resource-group3,...

```

where the `rgList` parameter is the comma-separated list of application resource groups that the administrator has opted to add. The script is not responsible for creating these resource groups. Instead, the resource groups must already exist on both clusters. Furthermore, these resource groups must have the `auto_start_on_new_cluster` property set to `false`.

The function name for this step is `add_application_rgs`.

configuration_file Property

The `configuration_file` property specifies the file name of the configuration file used to drive the execution of replicated component-level scripts described in [“Plug-In Script Functional Requirements” on page 148](#). Because individual script-based plug-ins inside a protection group might be on disjoint node sets or individual nodes, you should call the user scripts only on the appropriate cluster node or nodes. For more information, see [“Plug-In Script Functional Requirements” on page 148](#).

The configuration file must exist on all cluster nodes on both the primary and standby clusters. The script-based plug-in module tries to read the file from each node in turn until it finds a readable copy, but makes no effort to determine whether all copies are identical.

The format of the configuration file is as follows:

```

script-based-plug-in-configuration-name | nodes-that-must-succeed-running-script | comma-separated-node-list

```

For example:

```

foo.com | any | phys-node1, phys-node2
bar.com | all | phys-node1, phys-node3
baz.com | any | phys-node4
boo | any | phys-node4
biff | all | phys-node2

```

The script-based plug-in configuration name field must match the name of the replicated component being added to the protection group through the `geopg add-device-group` command.

For `foo.com`, a particular function step is tried on `phys-node1` and then, if it fails on `phys-node2`. The function step can succeed on either node. This configuration assumes that the service is a multinode service like Oracle RAC.

For `bar.com`, a particular function step must succeed on both `phys-node1` and `phys-node3` for the step to complete. Again, this configuration is only relevant to multinode services like Oracle

RAC. This function step enables a script to perform a task on multiple nodes without needing to connect to a remote node using `rsh` or `ssh` between the nodes.

create_config_script Property

The script referenced by the `create_config_script` property is responsible for creating, modifying, and validating a script-based plug-in configuration. The script must be able to validate the `create_config_args` property supplied to it with the `validate_parameters=true` option without actually performing the configuration creation.

When executed with `validate_parameters=false`, the script must create a replication group and an associated replication resource for the particular script-based plug-in. There must be only one replication resource group per script-based plug-in protection group and only one replication resource per replicated component. For example, a configuration with two script-based plug-in protection groups (`hr-pg` and `sales-pg`), each with two replicated components (`hr-west` and `hr-east` for `hr-pg`, and `sales-north` and `sales-south` for `sales-pg`), would have two resource groups (`hr_pg_rep-rg` and `sales_pg_rep-rg`). These resource groups would then have the following two resources:

- `hr-west-rep-rs` and `hr-east-rep-rs` in `hr_pg_rep-rg`
- `sales-north-rep-rs` and `sales-south-rep-rs` in `sales_pg_rep-rg`

When creating the second replicated component or validating either configuration, the script must handle the case where the resource group already exists.

On completion, the script must write the resource group name and resource to standard output. This task is checked by the script-based plug-in framework to both validate that the objects exist and to set up the appropriate notification handling for state change events. The format for the output is as follows:

```
reprg=replication-resource-group-name  
reprs=replication-resource-name
```

For example, for the case where the replication resource group is called `hr_pg_rep-rg` and the replication resource is called `hr-west-rep-rs`, the output would be as follows:

```
reprg=hr_pg_rep-rg  
reprs=hr-west-rep-rs
```

The script must also write a list of resource groups to standard output that it has either created, or that exist already, or that it considers internal to the protection group. The format of the output must be as follows, with a carriage return at the end of the line:

```
rglist=comma-separated-list-of-resource-groups
```

For example, for the case where `foo-rg` and `bar-rg` are internal, the output would be as follows:

rgList=foo-rg,bar-rg

If no resource groups exist, the output would be as follows:

rgList=

Examples of such internal resource groups are the lightweight resource groups in the Availability Suite module or the shadow RAC proxy server resource groups in the Oracle Data Guard module.

This script is called for each script-based plug-in created in any specific protection group because `create_config_script` is a global protection group property. For example, if a protection group has script-based plug-in configurations `foobar.com` and `baz.com`, the `create_config_script` script is called once when `foobar.com` is added with the `create_config_args` property given for the `foobar.com` property. The script is later called for `baz.com` when it is added to the protection group with the `baz.com create_config_args` property value. This process results in a replication resource group with two resources: one resource monitoring `foobar.com` replication and the other resource monitoring `baz.com`.

If the protection group is known to both the primary and standby sites, then adding the script-based plug-in configuration to the protection group will cause the `create_config_script` script to be executed on the site that the `geopg` command is run from and then on the remote site as a result of the internal Oracle Solaris Cluster Geographic protection group table transfer. The latter step happens asynchronously.

The `create_config_script` script is called with the `create_config_args` property followed by the standard command-line arguments and an additional `isModify` parameter. This parameter is set to `false` when the command has been called as a result of a `geopg create-device-group` or `geopg validate pgcommand`. This parameter is set to `true` when the command has been called as a result of a `geopg modify-device-group` command.

For example, if `create_config_script = /var/tmp/add` and `create_config_args = "-u root -d /mydir"`, the resulting command looks like the following example:

```
/var/tmp/add -u root -d /mydir function=create_configuration \  
validate_parameters={true|false} currentRole={PRIMARY|SECONDARY} \  
pg=protection-group isModify={true|false}
```

The function name for this step is `create_configuration`.

remove_app_rg_script Property

The script referenced by the `remove_app_rg_script` property is responsible for removing one or more application resource groups, selected by the administrator, from the protection group. A comma-separated list of resource groups to remove is passed to the script through the `rgList`

parameter. The script is called on the local cluster to where the `geopg remove-resource-group` command is run and called asynchronously on the remote cluster in response to the internal application resource group table being updated.

The script must also be able to validate the `remove_app_rg_args` property supplied to it with the `validate_parameters=true` option without actually performing any of the steps associated with this task. This operation is called only at the time of protection group update and creation, as opposed to at the time of device group update, modification, or validation.

For example, if `remove_app_rg_script = /var/tmp/removeRGs` and `remove_app_rg_args = "-u root -d /mydir"`, the resulting command looks like the following example:

```
#!/var/tmp/removeRGs -u root -d /mydir\  
function=remove_application_rgs \  
validate_parameters={true|false} \  
currentRole={PRIMARY|SECONDARY} pg=protection-group\  
rgList=resource-group1,resource-group2,resource-group3,...
```

where the `rgList` parameter is the comma-separated list of application resource groups that the administrator has opted to remove. The script is not responsible for removing these resource groups, only for making the necessary changes to their properties that might be required as a result of removing them from Geographic Edition protection group control.

The function name for this step is `remove_application_rgs`.

remove_config_script Property

The script referenced by the `remove_config_script` property is responsible for reversing the work of the `create_config_script` script. The script must be able to validate the `remove_config_args` property supplied to it with the `validate_parameters=true` option without actually performing the configuration removal.

When executed with `validate_parameters=false`, the script must remove the replication resource (originally named by the `create_config_script` script `reprs=` output for the specific script-based plug-in) from the replication resource group given by the `create_config_script` script `reprg=` output. If the resource is the last in the resource group, the script must also remove the resource group.

For example, if `remove_config_script = /var/tmp/remove` and `remove_config_args = "-u root -d /mydir"`, the resulting command looks like the following example:

```
# /var/tmp/remove -u root -d /mydir function=remove_configuration \  
validate_parameters={true|false} \  
currentRole={PRIMARY|SECONDARY} pg=protection-group
```

The function name for this step is `remove_configuration`.

start_replication_script Property

The script referenced by the `start_replication_script` property is responsible for starting the data replication process and enabling the replication resource that is used to monitor the replication. The script must also be able to validate the `start_replication_args` property supplied to it with the `validate_parameters=true` option without actually starting the data replication.

When executed with `validate_parameters=false`, the script must start the actual data replication and enable the replication resource that is used to monitor the replication.

For example, if `start_replication_script = /var/tmp/start` and `start_replication_args = "-u root -d /mydir"`, the resulting command looks like the following example:

```
# /var/tmp/start -u root -d /mydir function=start_replication \
validate_parameters={true|false} \
currentRole={PRIMARY|SECONDARY} pg=protection-group
```

The `start_replication_script` script is called on one or both clusters depending on which of the following commands the administrator specifies:

For local clusters only:

```
# geopg start -e local protection-group # local cluster only
```

For both clusters:

```
# geopg start -e global protection-group # both clusters
```

The function name for this step is `start_replication`.

stop_replication_script Property

The script referenced by the `stop_replication_script` property is responsible for stopping the data replication process and disabling the replication resource that is used to monitor the replication. The script must also be able to validate the `stop_replication_args` property supplied to it with the `validate_parameters=true` option without actually starting the data replication.

When executed with `validate_parameters=false`, the script must stop the actual data replication and disable the replication resource that is used to monitor the replication.

For example, if `stop_replication_script = /var/tmp/stop` and `stop_replication_args = "-u root -d /mydir"`, the resulting command looks like the following example:

```
# /var/tmp/stop -u root -d /mydir function=start_replication \  
validate_parameters={true|false} \  
currentRole={PRIMARY|SECONDARY} pg=protection-group
```

The `stop_replication_script` script is called on one or both clusters depending on which of the following commands the administrator specifies:

For local cluster only:

```
# geopg stop -e local protection-group # local cluster only
```

For both clusters:

```
# geopg stop -e global protection-group # both clusters
```

The function name for this step is `stop_replication`.

switchover_script Property

The script referenced by the `switchover_script` property is responsible for two functions:

- Checking that the service is in a position to switch over
- Performing the actual data replication switchover

The second step is only performed if the first step is completed successfully, meaning that the step exits with a zero exit code. In each case, the script is called on both clusters.

The `switchover_script` script is first called on the cluster on which the `geopg switchovercommand` is executed. Subsequent changes in Geographic Edition status trigger an event on the remote cluster, causing the script to be executed asynchronously on that cluster, too. The arguments for the two calls are different.

For example:

If `switchover_script = /var/tmp/switchover` and `switchover_args = "-u root -d /mydir"`, the resulting command looks like the following example:

```
# /var/tmp/switchover -u root -d /mydir function=check_switchover \  
validate_parameters=false currentRole={PRIMARY|SECONDARY} \  
pg=protection-group newRole={PRIMARY|SECONDARY}
```

If that step succeeds:

```
# /var/tmp/switchover -u root -d /mydir \  
function=perform_switchover \  
validate_parameters=false \  
currentRole={PRIMARY|SECONDARY} pg=protection-group \  
\
```

```
newRole={PRIMARY|SECONDARY}
```

The argument `newRole` is the target role of the cluster after a successful switchover.

The function names for these steps are `check_switchover` and `perform_switchover` and just `switchover` for the `validate_parameter` step, which is called as follows:

```
# developer-switchover-program developer-switchover-program-arguments \
function=switchover validate_parameters=true \
currentRole={PRIMARY|SECONDARY} pg=protection-group
```

takeover_script Property

The script referenced by the `takeover_script` property is responsible for two functions:

- Checking that the service is in a position to be taken over
- Performing the actual data replication takeover

The second step is only performed if the first step is completed successfully, meaning that the step exits with a zero exit code. In each case, the script is called on both clusters. If the original primary cluster is available, the protection group is deactivated on that cluster. Deactivation involves stopping the application resource groups.

The `takeover_script` script must be called on the standby cluster by executing the `geopg takeover` command on that cluster. The arguments for the two calls are different.

For example, if `takeover_script = /var/tmp/switchover` and `takeover_args = "-u root -d /mydir"`, the resulting command looks like the following example:

```
# /var/tmp/switchover -u root -d /mydir function=check_takeover \
validate_parameters=false currentRole={PRIMARY|SECONDARY} \
pg=protection-group newRole={PRIMARY|SECONDARY}
```

Then, if that step succeeds:

```
# /var/tmp/switchover -u root -d /mydir function=perform_takeover \
validate_parameters=false currentRole={PRIMARY|SECONDARY} \
pg=protection-group newRole={PRIMARY|SECONDARY}
```

The argument `newRole` is the target role of the cluster after a successful takeover.

The function names for these steps are `check_takeover` and `perform_takeover` and just `takeover` for the `validate_parameter` step, which is called as follows:

```
# developer-takeover-program developer-takeover-program-arguments \
function=takeover validate_parameters=true \
currentRole={PRIMARY|SECONDARY} pg=protection-group
```

Internals for Script-Based Plug-Ins

This section describes the internals for the script-based plug-ins. It covers the following topics:

- [“Plug-In Script Functional Requirements” on page 148](#)
- [“Script-Based Plug-In Replication Resource Groups and Resources” on page 150](#)
- [“Protection Group Status Mapped from Replication Resource Status” on page 150](#)
- [“How Geographic Edition Handles Password Properties” on page 151](#)

Plug-In Script Functional Requirements

A protection group has several global properties that are valid and relevant to both the primary and secondary clusters, and by extension all cluster nodes. Additionally, each replicated component has a set of local and global properties. Together, these properties describe and control the replication pertaining to one or more replicated services.

This section describes the following topics:

- [“Plug-In Script Argument Validation” on page 148](#)
- [“Standardized Script Command-Line Arguments” on page 149](#)

Plug-In Script Argument Validation

Each script provided in one of the protection group properties must be capable of validating the arguments with which it has been called in order to determine whether the arguments are complete and acceptable. Validation ensures that scripts such as `switchover_script` and `takeover_script`, that are not called regularly, do not fail because their arguments have become incompatible. Failing to validate the arguments could lead to the inability to switch over or take over in an emergency.

Scripts must therefore be able to validate the arguments defined by the administrator through the Geographic Edition command-line interface (CLI), and issue a return code of zero, if they are correct. The script must not perform its real function at this stage, for example, to switch over, take over, or create a script-based plug-in configuration. If you do not want to perform these checks, the script must still return without performing any additional work in response to the `validate arguments` call.

The `validate arguments` step is denoted by the Geographic Edition script-based plug-in Mbean passing `validate_parameters=true` as one of the command-line arguments. When a script-based plug-in replication component is added to a protection group, all the replicated component-specific scripts listed in [“Protection Group Properties - Overview” on page 137](#)

are called on to validate their arguments. This call is made on one or more nodes per cluster depending on the particular script-based plug-in replicated component configuration as defined in the configuration file. For more information, see [“configuration_file Property” on page 141](#) and [“Protection Group Properties - Overview” on page 137](#).

The same validation calls are made under the following circumstances:

- When the replication component is modified because the modification might result in program argument changes
- When there are protection group validation calls in response to the `geopg validate protection-group` command
- When the Geographic Edition software is starting and recreating the initial script-based plug-in replicated component objects that are stored in the Cluster Configuration Repository (CCR)

There are also two protection group level program properties, `add_app_rg_script` and `remove_app_rg_script`, that have associated protection group argument properties.

Standardized Script Command-Line Arguments

All scripts are called using a standardized command-line structure. The format of the command line is as follows:

```
# developer-program-name administrator-supplied-program-arguments \  
function=step-name \  
validate_parameters={true|false} \  
currentRole={PRIMARY|SECONDARY} \  
pg=protection-group \  
additional-function-dependent-arguments
```

where *developer-program-name* is the name of one of the externally developed scripts and *administrator-supplied-program-arguments* provides the arguments given for this script by the administrator when setting up a script-based plug-in configuration.

The use of the `function=step-name` argument enables scripts to determine what action they are being called on to perform. This function is especially important if a single script has been written to perform one or more tasks. Two scripts in particular need to be concerned with this argument: `switchover_script` and `takeover_script`.

The `currentRole` argument indicates the current role of the local cluster, while the `pg` argument denotes the name of the protection group containing the script-based plug-in configuration. Scripts should be prepared to deal with values in either uppercase or lowercase. The same is true of the `newRole` argument for `switchover_script` and `takeover_script`.

All scripts, if successful, must return a zero exit code. On failure, all scripts must return a nonzero exit code and generate a localized error message on standard error (`stderr`).

Any output sent to standard output (stdout) is generally ignored (with the exception of `create_config_script`), unless common agent container logging is turned on. In that case, the output is saved in the `/var/cacao/instances/default/logs/cacao.0` log file, along with other common agent container debugging information. Do not save debugging information as a matter of course because the volume of output can be substantial.

Script-Based Plug-In Replication Resource Groups and Resources

The name of the replication resource group for a particular protection group is defined by the value returned by `create_config_script` in the `reprg=` string sent to standard output. This string contains one or more replication resources referenced by individual replication resources named by `create_config_script` in the `reprs=` string sent to standard output. For any one protection group, the value returned by `create_config_script` must be identical.

The function of the replication resources is to monitor the state of the replication associated with the resource and thus the replicated component. The replication resource status, which is set by a probe method, is used to determine the overall status of the protection group. The start and stop methods of the replication resource do not start and stop the actual data replication.

The replication resource must be enabled and disabled by `start_replication_script` and `stop_replication_script`.

Protection Group Status Mapped from Replication Resource Status

The protection group status reflects the aggregated status of all replication resources in the replication resource group created by the developer-written `create_config_script` program.

The following table illustrates the mapping from the status of each replication resource to the protection group status. An X represents any possible status for the resource and demonstrates that the most restrictive status governs the overall status of the protection group.

Unknown	Faulted	Degraded	Online	Protection Group Status
True	X	X	X	UNKNOWN
False	True	X	X	FAULTED
False	False	True	X	DEGRADED
False	False	False	True	ONLINE

How Geographic Edition Handles Password Properties

This section describes the mechanism by which Geographic Edition handles password properties, when the entity added to a protection group (for example, an Oracle Data Guard or script-based plug-in configuration) requires a password property.

The password properties are read during the execution of the `geopg` command. These password properties are recognized by their conformance to the pattern `*_password`. When `geopgi` (a back-end program called by the `geopg` command) parses the protection group properties list, it looks for such arguments. If the password has been supplied in cleartext, as shown in the following example, then `geopg` warns the user that the password is insecure, but continues processing the password.

```
... -p sysdba_password=password ...
```

For any password properties that have been specified, the `geopgi` program enters non-echo mode and prompts for these passwords, as shown in the following example:

```
... -p local_service_password= -p remote_service_password= ...
```

Once all the arguments have been processed, these pairs are written into an internal password file on the local node, which is root readable only. A separate `internalPasswordField` argument is inserted into the properties list with the value `hostname:filename`.

Once in the core Geographic Edition Java code, the `internalPasswordField` argument is unpacked, and the file is read remotely through an internal common agent container to common agent container call. For security, the passwords are then converted into the hexadecimal representation of their character codes before they are written to the Oracle Solaris Cluster CCR, if the rest of the properties are correct and complete, and the validation succeeds.

When required, the passwords can be queried and converted back from the CCR and supplied to the appropriate programs to achieve the relevant switchovers, takeovers, or status queries.

◆◆◆ **A P P E N D I X A**

Standard Geographic Edition Properties

This appendix provides the standard properties of Geographic Edition heartbeats, heartbeat plug-in, partnerships, protection groups and data replication device groups, sites, and multigroups.

This appendix contains the following sections:

- [“Heartbeat Properties” on page 153](#)
- [“Heartbeat Plug-in Properties” on page 154](#)
- [“Partnership Properties” on page 155](#)
- [“Protection Group Properties” on page 156](#)
- [“Site Properties” on page 157](#)
- [“Multigroup Properties” on page 158](#)

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

Heartbeat Properties

The following table describes the heartbeat properties that the Geographic Edition software defines.

TABLE A-1 Geographic Edition Heartbeat Properties

Property Name	Description
Query_interval (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>

Heartbeat Plug-in Properties

The following table describes the heartbeat plug-in properties that the Geographic Edition software defines.

TABLE A-2 Geographic Edition 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>remote-IP-address/UDP/2084/ipsec</code>, <code>remote-IP-address/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>
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 <code>Responder_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>

Property	Description
Type (enum)	<p>Designates the type of plug-in. Set to either <code>primary</code> or <code>backup</code>.</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 <code>ping_plugin</code>. If using this plug-in, the default value is <code>backup</code>.</p>

Partnership Properties

The following table describes the partnership properties that the Geographic Edition software defines.

TABLE A-3 Geographic Edition 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>
Notification_EmailAdrs (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>

Protection Group Properties

The following table describes the protection group properties that the Geographic Edition software defines.

TABLE A-4 Geographic Edition Protection Group Properties

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>
External_Dependency_Allowed (Boolean)	<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 <code>true</code>.</p> <p>Tuning recommendations: This property can be tuned at any time.</p> <p>Category: Optional</p> <p>Default: <code>false</code></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, <code>RoleChange_ActionCmd</code>, is run.</p> <p>Tuning recommendations: This property can be tuned at any time.</p> <p>Category: Optional</p> <p>Default: Empty string</p>
RoleChange_ActionCmd (string)	<p>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>This file should be valid on all nodes of all partner clusters that can host the protection group, and have execute permissions for the user that launches the script.</p> <p>Tuning recommendations: This property can be tuned at any time.</p> <p>Category: Optional</p> <p>Default: Empty string</p>
Timeout (integer)	<p>Specifies the timeout period for the protection group in seconds. The timeout period is the longest time Geographic Edition waits for a response after you run a <code>geopg</code> command, such as <code>geopg start</code>, <code>geopg stop</code>, <code>geopg switchover</code>, and <code>geopg takeover</code>. If the command does not respond within the timeout period, the Geographic Edition software reports the operation as timed out, even if the underlying command eventually completes successfully.</p>

Property	Description
	<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 170.</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>

Site Properties

The following table describes the site properties that the Geographic Edition software defines.

TABLE A-5 Geographic Edition Site Properties

Property	Description
Description (string)	<p>Describes the site. The system sets this property on the local cluster, then propagates the value to the other clusters in the site.</p> <p>Tuning recommendations: This property is assigned at creation and tunable at runtime.</p> <p>Category: Optional</p> <p>Default: None</p>
joinTimeout (integer)	<p>Specifies, in seconds, the longest time that the JMX client, the site, waits for the invocation of an MBean-server method to return. If a method does not return by the end of the timeout period, the client moves to its next set of instructions and reports the operation as timed out. By default, a client waits indefinitely for a method to return. If the MBean server is unable to complete an invocation, the JMX client will hang indefinitely. The timeout period is site-wide and applies to operations on a per-cluster basis.</p>

Property	Description
	<p>The <code>joinTimeout</code> property is used when the join operation requires more time to complete than the default setting allows before the <code>geosite joincommand</code> times out.</p> <p>Tuning recommendations: This property is assigned at creation, but it is not tunable at runtime.</p> <p>Category: Optional</p> <p>Default: 30 seconds</p> <p>Minimum: 20 seconds</p> <p>Maximum: 3600 seconds</p>
<p><code>Query_interval</code> (integer)</p>	<p>Specifies, in seconds, the frequency between heartbeat status requests used by the site clusters. The plug-in enters emergency mode if three <code>Query_interval</code> periods pass without response. The plug-in times out and enters error mode if a further <code>Query_interval</code> period passes with no response.</p> <p>Tuning recommendations: This property is assigned at creation and tunable at runtime.</p> <p>Category: Optional</p> <p>Default: 120 seconds</p> <p>Minimum: 20 seconds</p> <p>Maximum: 300 seconds</p>
<p><code>Timeout</code> (integer)</p>	<p>Specifies, in seconds, the longest time that the JMX client, the site, waits for the invocation of an MBean-server method to return. If a method does not return by the end of the timeout period, the client moves to its next set of instructions and reports the operation as timed out. By default, a client waits indefinitely for a method to return. If the MBean server is unable to complete an invocation, the JMX client will hang indefinitely. The timeout period is site-wide and applies to operations on a per-cluster basis.</p> <p>Tuning recommendations: This property is assigned at creation and tunable at runtime.</p> <p>Category: Optional</p> <p>Default: 30 seconds</p> <p>Minimum: 30 seconds</p> <p>Maximum: 3600 seconds</p>

Multigroup Properties

The following table describes the multigroup properties that the Geographic Edition software defines.

TABLE A-6 Geographic Edition Multigroup Properties

Property	Description
Description (string)	<p>Describes the multigroup. The system sets this property on the local cluster, then propagates the value to the other clusters in the site where the multigroup is defined.</p> <p>Tuning recommendations: This property is assigned at creation and tunable at runtime.</p> <p>Category: Optional</p> <p>Default: None</p>
Timeout (integer)	<p>Specifies, in seconds, the longest time that the JMX client, the site, waits for the invocation of an MBean-server method to return. If a method does not return by the end of the timeout period, the client moves to its next set of instructions and reports the operation as timed out. By default, a client waits indefinitely for a method to return. If the MBean server is unable to complete an invocation, the JMX client will hang indefinitely. The timeout period is site-wide and applies to operations on a per-cluster basis.</p> <p>The timeout period applies to operations on a per-cluster basis. An operation with a local scope times out if the operation is not completed 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 actions are timed separately. So, an operation with a global scope times out if the local operation is not completed after the specified timeout period or if the remote operation is not completed after the specified timeout period.</p> <p>For example, the following command is started with a local scope:</p> <pre># geomg start -e local multigroup</pre> <p>If you set the Timeout property to 3000 seconds, the <code>geomg start</code> command times out if the operation does not complete after 3000 seconds.</p> <p>You can start the same command with a global scope as follows:</p> <pre># geomg start -e global multigroup</pre> <p>If the Timeout property is set to 3000 seconds, the <code>geomg start</code> command times out if the operation is not completed on the local cluster after 3000 seconds or if the operation is not completed on the remote cluster after 3000 seconds. If the local action takes 1500 seconds and the remote action takes 1500 seconds, the operation is not timed out.</p> <p>The multigroup timeout value is an estimated value that is applied to some sub-operations. The timeout value does not apply to the entire operation, so not every operation on a protection group is timed against the timeout period. For example, the time taken to initialize the data structure and to check for the precondition of the operation are not timed in the timeout period.</p> <p>Tuning recommendations: This property is assigned at creation and tunable at runtime.</p> <p>Category: Optional</p>

Multigroup Properties

Property	Description
	Default: 30 seconds Minimum: 20 seconds Maximum: 3600 seconds

◆◆◆ APPENDIX B

Legal Names and Values of Geographic Edition Entities

This appendix lists the requirements for legal characters for the names and values of Geographic Edition entities.

This appendix contains the following sections:

- “Legal Names for Geographic Edition Entities” on page 161
- “Legal Values for Geographic Edition Entities” on page 162

Legal Names for Geographic Edition Entities

Geographic Edition entity names consist of the following:

- Host names
- Cluster names, which must follow the naming requirements for host names
- Partnership names
- Protection group names
- Custom heartbeat names

All names must comply with the following rules:

- Must start with a letter
- Must not exceed 255 characters
- Can contain the following:
 - Upper and lowercase letters
 - Digits
 - Dashes (-), except as the last character of a host name or cluster name
 - Underscores (_), except in a host name or cluster name

For more information about host name requirements, see RFC 1123 at <http://www.rfcs.org/>.

Legal Values for Geographic Edition Entities

The 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

Disaster Recovery Administration Example

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

Example Company 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 [“Modifying Partnership Properties” on page 39](#).

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, Example Company does not allow takeovers unless the primary cluster cannot be repaired within two hours.

For information about verifying a disconnect on a system, see [“Detecting Cluster Failure” on page 117](#).

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 [“Forcing a Takeover of a Protection Group” on page 123](#).

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 the Geographic Edition manual for your data replication product.

Takeover Postconditions

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

This appendix contains the following sections:

- [“Results of a Takeover When the Partner Cluster Can Be Reached” on page 165](#)
- [“Results of a Takeover When the Partner Cluster Cannot Be Reached” on page 166](#)

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 D-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
cluster-paris: primary, deactivated cluster-newyork: secondary, deactivated	cluster-paris: secondary, deactivated cluster-newyork: primary, deactivated
cluster-paris: primary, activated cluster-newyork: secondary, deactivated	cluster-paris: secondary, deactivated cluster-newyork: primary, deactivated
cluster-paris: primary, deactivated cluster-newyork: secondary, activated	cluster-paris: secondary, deactivated cluster-newyork: primary, activated, with data replication stopped
cluster-paris: primary, activated cluster-newyork: secondary, activated	cluster-paris: secondary, deactivated cluster-newyork: primary, activated, with data replication stopped

The following table describes the states when you run the `geopg takeover` command on the primary cluster, `cluster-paris`.

TABLE D-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
<code>cluster-paris: primary, deactivated</code> <code>cluster-newyork: secondary, deactivated</code>	<code>cluster-paris: primary, deactivated</code> <code>cluster-newyork: secondary, deactivated</code>
<code>cluster-paris: primary, activated</code> <code>cluster-newyork: secondary, deactivated</code>	<code>cluster-paris: primary, activated, with data replication stopped</code> <code>cluster-newyork: secondary, deactivated</code>
<code>cluster-paris: primary, deactivated</code> <code>cluster-newyork: secondary, activated</code>	<code>cluster-paris: primary, deactivated</code> <code>cluster-newyork: secondary, deactivated</code>
<code>cluster-paris: primary, activated</code> <code>cluster-newyork: secondary, activated</code>	<code>cluster-paris: primary, activated, with data replication stopped</code> <code>cluster-newyork: secondary, deactivated</code>

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 D-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
<code>cluster-paris: primary, deactivated, synchronization status Unknown</code> <code>cluster-newyork: secondary, deactivated, synchronization status Unknown</code>	<code>cluster-paris: primary, deactivated, synchronization status Error</code> <code>cluster-newyork: primary, deactivated, synchronization status Error</code>

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

◆◆◆ APPENDIX E

Troubleshooting Geographic Edition Software

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

This appendix contains the following sections:

- [“Troubleshooting Monitoring and Logging” on page 169](#)
- [“Troubleshooting Migration Problems” on page 170](#)
- [“Troubleshooting Cluster Start and Restart” on page 171](#)

Troubleshooting Monitoring and Logging

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

- [“Configuring the Logger File to Avoid Too Many Traces” on page 169](#)
- [“Configuring the Logger File to Avoid Detailed Messages From the gc r Agent” on page 170](#)
- [“Configuring the Logger File to Avoid jmx Remote Traces” on page 170](#)

For information about logging, see [“Viewing the Geographic Edition Log Messages” on page 114](#).

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 Logger 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 logger 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 `SUNWscmasa` package.

Configuring the Logger 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 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 Oracle Solaris 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 Oracle Solaris Cluster procedures. Refer to [“Clearing the STOP_FAILED Error Flag on Resources”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#) about recovering from the `ERROR_STOP_FAILED` state and restarting the application.

The 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 Geographic Edition Software”](#) on page 107.

Troubleshooting Cluster Start and Restart

This section provides the following information about troubleshooting problems that you might encounter with starting and restarting the Geographic Edition software:

- [“Validating Protection Groups in an Error State”](#) on page 171
- [“Restarting the Common Agent Container”](#) on page 172
- [“Matching the NodeList Property of an Availability Suite Protection Group to Those of Its Device Group and Resource Group”](#) on page 172

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 Oracle Solaris Cluster software enables the common agent container only during the Oracle Solaris 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.

Matching the `nodeList` Property of an Availability Suite Protection Group to Those of Its Device Group and Resource Group

When you add resource groups, or Availability Suite device groups to a protection group, or when you run the command `geopg get` on a protection group, the order of the hosts in the `nodeList` property of each device group and resource group in the protection group must match the order of the hosts in the `nodeList` property of the protection group, or the operation will fail with a message similar to:

```
Application resource group app-rg must have a nodeList whose physical host components match those of protection group app-pg and the resources it contains.
```

The Geographic Edition software requires that the entries in the `nodeList` property of an Availability Suite protection group match those of any device group or resource group added to the protection group. The order of the entries in their `nodeList` properties must also be identical.

Solution or Workaround

Ensure that the entries, and the order of the entries in the `nodeList` properties of a protection group, of its device groups, and of its resource groups are identical.

Error Return Codes for Script-Based Plug-Ins

Error Codes That Are Returned in Script-Based Plug-Ins

The script-based plug-in MBean can return any of the error codes shown in the following table.

Return Code	Error Message	Description
101	E_SBP_PROGRAM_FAILED_TO_READ_CCR	Program {0} failed to read the cluster configuration repository (CCR)
110	E_SBP_PROGRAM_EXITED_NON_ZERO	Program {0} returned a nonzero exit code.
112	E_SBP_UNEXPECTED_ERROR	Unexpected error - {0}.
125	E_SBP_ONE_OR_MORE_RGS_NON_EXISTENT	One or more of the resource groups ({0}) returned by program {1} do not exist.
126	E_SBP_RG_LIST_WRONG_FORMAT,	The output {0} returned by program {1} is invalid. The output must conform to the format <i>rglist=comma-separated-resource-groups</i>
127	E_SBP_NO_SUCH_FILE	An attempt was made to execute a null or nonexistent command. Check the logs for more details.
128	E_SBP_CANNOT_READ_CONFIG_FILE	Unable to read configuration file {0} from any cluster node. This file must be available on all cluster nodes.
129	E_SBP_ENTRY_NOT_FOUND_IN_CONFIG_FILE	No entry for script-based plug-in configuration {0} exists in configuration file {1}.
130	E_SBP_CONFIG_FILE_FORMAT_ERROR	Field {0} in configuration file {1} must be {2}.
131	E_SBP_CONFIG_FILE_FIELD_FORMAT_ERROR	Configuration file {0} must have three fields per script-based plug-in entry. The fields must be separated by " ".
132	E_SBP_CONFIG_FILE_INVALID_NODE	The entry for script-based plug-in {0} in configuration file {1} contains an invalid cluster node {2} in the node list field.
133	E_SBP_FAILED_TO_CHECK_X_BIT	Failed to check whether {0} is executable on cluster node {1}.

Error Codes That Are Returned in Script-Based Plug-Ins

134	E_SBP_SCRIPT_DOES_NOT_EXIST	Script or program {0} does not exist on cluster node {1}.
135	E_SBP_SCRIPT_FILE_IS_NOT_EXECUTABLE	Script or program {0} is not executable on cluster node {1}.
136	E_SBP_INTERRUPTED_OR_TIMED_OUT	The command or probe was interrupted or timed out.
138	E_SBP_COULD_NOT_GET_MBEAN_PROXY	Unable to get MBean proxy for {0} on node {1}.
139	E_SBP_FILE_SECURITY_ACCESS_REFUSED	The Java Security Manager refused access to file {0} on node {1}.
140	E_SBP_NULL_FILE_NAME	File name for property {0} was null.
141	E_SBP_UNABLE_TO_CREATE_SBP_CONFIG	Unable to create script-based plug-in configuration {0}.
142	E_SBP_UNABLE_TO_MODIFY_SBP_CONFIG	Unable to modify script-based plug-in configuration {0}.
143	E_SBP_UNABLE_TO_DELETE_DG	Unable to delete script-based plug-in configuration {0}.
144	E_SBP_UNABLE_TO_CREATE_PROPERTY	Unable to create property {0}.
146	E_SBP_UNABLE_TO_UPDATE_PG_PROPERTY	Unable to update protection group property {0}.
148	E_SBP_UNABLE_TO_UPDATE_PROPERTY	Unable to update property {0}.
150	E_SBP_UNABLE_TO_GET_PROPERTY	Unable to retrieve property.
151	E_SBP_UNABLE_TO_GET_CLUSTER_NODELIST	Unable to get cluster nodelist.
200	E_SBP_CONFIG_ERROR	Configuration error detected for protection group {0}.
201	E_SBP_SCRIPT_FAILED	The user supplied script-based plug-in command {0} failed with error code {1} on node {2}. The script error message is {3}.
210	E_SBP_INVALID_PROPERTY_FILE	Invalid property file {0}.
221	E_SBP_MISSING_PROPERTY	Property {0} is not set.
222	E_SBP_DUPLICATE_PROPERTY	Duplicate property {0}.
223	E_SBP_INVALID_PROPERTY	Invalid property {0}.
224	E_SBP_INVALID_PROPERTY_VALUE	Invalid value for property {0}.
225	E_SBP_SBP_CONFIG_ALREADY_IN_PG	Script-based plug-in configuration {0} already in protection group {1}.
226	E_SBP_SBP_CONFIG_NOT_FOUND_IN_PG	Script-based plug-in configuration {0} is not found in protection group {1}.
231	E_SBP_UNABLE_TO_NOTIFY_STATUS_CHANGE	Unable to send change notification for data replication status.
233	E_SBP_RG_OFFLINE_EXCEPTION	Failed to take resource group {0} offline.
234	E_SBP_SAME_PROPERTY_VALUE	Property value already set. No modification is needed.
235	E_SBP_UNEXPECTED_EXCEPTION	Unexpected exception - {0}.

Error Codes That Are Returned in Script-Based Plug-Ins

236	E_SBP_SERVER_REQUEST_FAILED_DUE_TO_TIMEOUT	Error in running control script on host {0}. Operation timed out after {1} seconds.
237	E_SBP_SERVER_REQUEST_FAILED_WITH_REASON	Error in running control script on host {0} due to system error - {1}\

Index

A

accepting an invitation to join a site, 87

activating, 30

See also enabling

See also starting

Geographic Edition software, 30

protection groups, 69, 71

guidelines, 69

adding

clusters to a site, 85

custom heartbeat plug-in

to a custom heartbeat, 62

to the default heartbeat, 61

new cluster nodes to a partnership, 40

protection groups to a multigroup, 97

administering

access, 37

Geographic Edition tasks, 25

heartbeats, 51

Oracle Solaris Cluster tasks, 23

overview, 19

protection groups, 69

security, 37

B

booting a cluster, 36

C

clusters

adding a new node, 40

adding to a site, 85

administration concepts, 23

booting, 36

changing a site role, 89

forcibly changing a site role, 90

joining a site, 87

leaving a partnership, 48

removing from a site, 93

removing from a site when unreachable, 94

sample configuration, 26

status of, 107

command-line interface

overview, 21

common agent container

restarting, 172

configuring

logger.properties file, 169

logical hostname, 24

role-change action script, 132

creating

custom heartbeat plug-in, 62

heartbeat plug-ins, 54

heartbeats, 52

role-change action script, 129

custom heartbeats

action script, 65

D

data replication components

modifying, 80

removing from a protection group, 83

deactivating

protection groups, 69, 75

guidelines, 69

DEBUG property

script-based plug-ins, 114

deleting, 48

See also removing

heartbeat plug-ins, 56

- heartbeats, 56
- multigroups, 104
- partnerships, 48
- protection groups, 81
- sites, 95
- dependency chains
 - modifying protection groups in, 99
 - syntax, 98
- detecting failure, 117
 - primary cluster, 117
 - secondary cluster, 118
- device groups
 - overview, 25
- disabling Geographic Edition software, 33
- disaster recovery overview, 163
- displaying
 - heartbeat configuration, 57
 - partnership configuration, 115
 - protection group configuration, 115

E

- enabling, 30
 - See also* activating
 - See also* starting
- enabling Geographic Edition software, 30
- examples
 - activating a protection group
 - globally, 74
 - locally, 75
 - adding a custom heartbeat plug-in to a new custom heartbeat, 64
 - adding a custom heartbeat plug-in to the default heartbeat, 62
 - configuring a protection group custom command, 132
 - configuring heartbeat-loss notification, 65
 - creating a heartbeat, 53
 - creating a heartbeat plug-in, 54
 - deactivating a protection group
 - globally, 77
 - keeping application resource groups online, 77
 - locally, 77
 - deleting a heartbeat, 56
 - deleting a partnership, 49
 - deleting a plug-in from a heartbeat, 57

- deleting a protection group, 82
- deleting a protection group with application resource groups online, 83
- disabling a cluster, 34
- displaying heartbeat configuration information, 58
- displaying partnership configuration information, 115
- displaying protection-group configuration information, 116
- displaying the infrastructure status, 36
- enabling the infrastructure, 32
- leaving a partnership, 49
- modifying a data replication component, 81
- modifying heartbeat plug-in properties, 55
- modifying partnership properties, 40
- modifying properties of the default heartbeat, 60
- notification action script, 66
- removing a data replication component from a protection group, 84
- renaming a cluster in a partnership, 45
- resynchronizing a partnership, 50
- sample clusters configuration, 26
- starting a protection group
 - globally, 74
 - locally, 75
- stopping a protection group
 - globally, 77
 - keeping application resource groups online, 77
 - locally, 77
- stopping data replication on an online protection group, 77
- switchover action script, 131
- takeover of a protection group, 126

F

- failback-switchover, 127
- failback-takeover, 127
- failure
 - detecting, 117
 - detecting for secondary cluster, 118
 - primary cluster, 117

G

- geo-cluster-name, 29

- geo-clusterstate, 29
- geo-failovercontrol, 29
- geo-hbmonitor, 29
- geo-infrastructure, 29
- geoadm show, 35
- geoadm status, 107
- Geographic Edition software
 - disabling, 33
 - enabling, 30
- graphical user interface (GUI)
 - overview of, 20
- guidelines
 - activating a protection group, 69
 - deactivating a protection group, 69
 - starting a protection group, 69
 - stopping a protection group, 69

H

- heartbeat plug-ins
 - adding a custom plug-in
 - to a custom heartbeat, 62
 - to the default heartbeat, 61
 - creating, 54
 - creating a custom plug-in, 60, 62
 - deleting, 56
 - modifying properties, 55
- heartbeat-loss notification, 64
- heartbeats, 54
 - See also* heartbeat plug-ins
 - adding a custom heartbeat plug-in
 - to a custom heartbeat, 62
 - to the default heartbeat, 61
 - administering, 51
 - creating, 52
 - creating a custom heartbeat plug-in, 60
 - custom action script, 65
 - deleting, 56
 - description, 51
 - displaying configuration, 57
 - heartbeat plug-in properties, 154
 - loss notification, 64
 - modifying properties, 59
 - properties, 153
 - tuning, 58

J

- joining a site, 87

L

- leaving
 - partnerships, 48
- logger.properties file
 - configuring, 169
- logging, 114
 - troubleshooting, 169
- logical hostname
 - configuring, 24
- loss of heartbeat notification, 64
 - creating action shell script, 65
 - properties, 65

M

- migration
 - troubleshooting, 170
- modifying
 - data replication components, 80
 - heartbeat plug-in properties, 55
 - heartbeat properties, 59
 - partnerships, 39
 - protection groups, 79
 - protection groups in a dependency chain, 99
 - RBAC rights, 37
- monitoring
 - Geographic Edition, 107
 - infrastructure resource groups, 29
 - troubleshooting, 169
- multigroups
 - actions during switchover, 119
 - actions during takeover, 124
 - adding protection groups, 97
 - deleting, 104
 - properties, 158
 - removing a protection group, 100
 - starting, 101
 - status of, 107
 - stopping, 102
 - switchover, 122
 - synchronizing configuration information, 103

- takeover, 126
- N**
- nodes
 - renaming, 42
 - notification_actioncmd, 64
 - notification_emailaddrs, 64
- O**
- operations
 - status of, 107
 - Oracle Solaris Cluster Manager, 20
 - Oracle Solaris Cluster software
 - administration concepts, 23
 - resources, 23
 - Oracle wallet, 73
- P**
- partnerships
 - adding a new cluster node, 40
 - deleting, 48
 - displaying configuration information, 115
 - leaving, 48
 - modifying, 39
 - properties, 155
 - removing trust, 46
 - renaming a cluster, 42
 - renaming a node in a cluster, 42
 - resynchronizing, 49
 - status of, 107
 - primary cluster
 - failure detection, 117
 - recovering from a switchover, 127
 - recovering from a takeover, 127
 - switchover, 118
 - properties
 - heartbeat, 153
 - heartbeat plug-in, 154
 - multigroup, 158
 - partnership, 155
 - protection group, 156
 - sites, 157
 - tuning heartbeat, 58
 - protection groups
 - actions during switchover, 119
 - actions during takeover, 124
 - activating, 69, 71
 - guidelines, 69
 - adding to a multigroup, 97
 - administering, 69
 - configuring
 - role-change action script, 132
 - deactivating, 69, 75
 - guidelines, 69
 - deleting, 81
 - displaying configuration information, 115
 - modifying, 79
 - modifying in a dependency chain, 99
 - properties, 156
 - removing data replication components, 83
 - removing from a multigroup, 100
 - resynchronizing, 78
 - starting, 71
 - guidelines, 69
 - in a multigroup, 101
 - status of, 107
 - stopping, 75
 - guidelines, 69
 - in a multigroup, 102
 - switchover, 120
 - takeover, 124
 - troubleshooting error state, 171
- R**
- RBAC
 - modifying rights, 37
 - recovering
 - after a switchover, 127
 - after a takeover, 127
 - removing, 48
 - See also* deleting
 - all clusters from a site, 95
 - clusters from a site, 93
 - clusters from a site when unreachable, 94
 - data replication components from a protection group, 83
 - protection groups from a multigroup, 100

- trust, 46
 - renaming
 - a cluster, 42
 - cluster nodes, 42
 - requirements
 - naming conventions, 161
 - role-change action script, 129
 - script-based plug-in scripts, 148
 - resource groups
 - configuring, 23
 - Geographic Edition infrastructure, 29
 - resources
 - configuring, 23
 - resynchronizing
 - partnerships, 49
 - protection groups, 78
 - role-change action script, 129
 - configuring protection group for, 132
 - creating, 129
 - requirements, 129
 - runtime status
 - Geographic Edition, 107
- S**
- script-based plug-ins, 133
 - DEBUG property, 114
 - error codes, 173
 - scripts
 - custom loss of heartbeat action, 65
 - switchover and takeover action, 129
 - secondary cluster
 - failure detection, 118
 - switchover, 118
 - security
 - administering, 37
 - sites
 - accepting an invitation to join, 87
 - adding a cluster, 85
 - changing a cluster role, 89
 - deleting, 95
 - forcibly changing a cluster role, 90
 - joining, 87
 - properties, 157
 - removing a cluster, 93
 - removing all clusters, 95
 - removing an unreachable cluster, 94
 - status of, 107
 - synchronizing configuration information, 92
 - troubleshooting
 - joining a site, 88
 - starting, 30
 - See also* activating
 - See also* enabling
 - cluster troubleshooting, 171
 - protection groups, 71
 - guidelines, 69
 - protection groups in a multigroup, 101
 - status
 - descriptions, 107
 - Geographic Edition infrastructure, 35
 - stopping
 - all protection groups in a multigroup, 102
 - protection groups, 75
 - guidelines, 69
 - switchover, 118
 - actions performed, 119
 - configuring role-change action script, 132
 - custom action script, 129
 - multigroups, 122
 - protection groups, 120
 - synchronizing
 - multigroup configuration information, 103
 - site configuration information, 92
- T**
- takeover
 - actions performed, 124
 - configuring role-change action script, 132
 - custom action script, 129
 - multigroups, 126
 - protection groups, 124
 - results after, 165
 - timeout
 - description of, 156, 159
 - troubleshooting
 - cluster restart, 171
 - cluster start, 171
 - forcibly changing a site cluster role, 91
 - joining a site, 88
 - log messages, 114

- logging, 169
- migration problems, 170
- monitoring, 169
- protection group error state, 171
- restarting common agent container, 172

trust

- removing, 46

tuning

- heartbeat properties, 58