

# Oracle® Solaris Cluster Geographic Edition Reference Manual

Copyright © 2004, 2012, 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.

# Contents

---

<b>Preface</b> .....	5
<b>OSC4GEO 1M</b> .....	9
geoadm(1M) .....	10
geohb(1M) .....	17
geopg(1M) .....	23
geops(1M) .....	44
<b>Index</b> .....	51



# Preface

---

The *Oracle Solaris Cluster Geographic Edition Reference Manual* provides reference information about commands that are supplied with Oracle Solaris Cluster Geographic Edition software. This book is intended for experienced system administrators with extensive knowledge of Oracle software and hardware. This book is not to be used as a planning or presales guide. The information in this book assumes knowledge of the Oracle Solaris Operating System and expertise with the data replication software that is used with Oracle Solaris Cluster Geographic Edition software.

Both novice users and those familiar with the Oracle Solaris Operating System can use online man pages to obtain information about their SPARC based system or x86 based system and its features.

A man page is intended to answer concisely the question “What does this command do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

---

**Note** – Oracle Solaris Cluster software runs on two platforms, SPARC and x86. The information in this book pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

---

## Overview

The following contains a brief description of each man page section and the information it references.

Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.

The following is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if no bugs can be reported, no BUGS section is included. See the `int ro` pages for more information and detail about each section, and `man(1)` for general information about man pages.

NAME	This section gives the names of the commands or functions that are documented, followed by a brief description of what they do.
------	---

## SYNOPSIS

This section shows the syntax of commands or functions. If a command or file does not exist in the standard path, its full path name is shown. Options and arguments are alphabetized, with single-letter arguments first, and options with arguments next, unless a different argument order is required.

The following special characters are used in this section:

- [ ] Brackets. The option or argument that is enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.
- . . . Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, “filename . . .”.
- | Separator. Only one of the arguments separated by this character can be specified at a time.
- { } Braces. The options and/or arguments enclosed within braces are interdependent. All characters within braces must be treated as a unit.

## PROTOCOL

This section occurs only in subsection 3R and indicates the protocol description file.

## DESCRIPTION

This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. DESCRIPTION does not discuss OPTIONS or cite EXAMPLES. Interactive commands, subcommands, requests, macros, and functions are described under USAGE.

## IOCTL

This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the `ioctl(2)` system call is called `ioctl` and generates its own heading. `ioctl` calls for a specific device are listed alphabetically (on the man page for that specific device). `ioctl` calls are used for a particular class of devices. All these calls have an `io` ending, such as `mtio(7I)`.

## OPTIONS

This section lists the command options with a concise summary of what each option does. The options are listed

---

	literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.
OPERANDS	This section lists the command operands and describes how they affect the actions of the command.
OUTPUT	This section describes the output – standard output, standard error, or output files – generated by the command.
RETURN VALUES	If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or –1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions that are declared void do not return values, so they are not discussed in RETURN VALUES.
ERRORS	On failure, most functions place an error code in the global variable <code>errno</code> that indicates why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.
USAGE	This section lists special rules, features, and commands that require in-depth explanations. The subsections that are listed here are used to explain built-in functionality:  Commands Modifiers Variables Expressions Input Grammar
EXAMPLES	This section provides examples of usage or of how to use a command or function. Wherever possible, a complete example, which includes command-line entry and machine response, is shown. Whenever an example is given, the prompt is shown as <code>example%</code> , or if the user must be superuser, <code>example#</code> . Examples are followed by explanations, variable substitution rules, or returned

	values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.
ENVIRONMENT VARIABLES	This section lists any environment variables that the command or function affects, followed by a brief description of the effect.
EXIT STATUS	This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion, and values other than zero are returned for various error conditions.
FILES	This section lists all file names that are referred to by the man page, files of interest, and files created or required by commands. Each file name is followed by a descriptive summary or explanation.
ATTRIBUTES	This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See <a href="#">attributes(5)</a> for more information.
SEE ALSO	This section lists references to other man pages, in-house documentation, and outside publications.
DIAGNOSTICS	This section lists diagnostic messages with a brief explanation of the condition that caused the error.
WARNINGS	This section lists warnings about special conditions that could seriously affect your working conditions. WARNINGS is not a list of diagnostics.
NOTES	This section lists additional information that does not belong anywhere else on the page. NOTES covers points of special interest to the user. Critical information is never covered here.
BUGS	This section describes known bugs and, wherever possible, suggests workarounds.



**R E F E R E N C E**

**OSC4GEO 1M**

**Name** geoadm – enable or disable the Oracle Solaris Cluster Geographic Edition infrastructure on the local cluster

**Synopsis** /usr/cluster/bin/geoadm [*subcommand*] -?  
 /usr/cluster/bin/geoadm -V  
 /usr/cluster/bin/geoadm show  
 /usr/cluster/bin/geoadm start  
 /usr/cluster/bin/geoadm status  
 /usr/cluster/bin/geoadm stop [-f | -i]

**Description** The geoadm command enables or disables the Oracle Solaris Cluster Geographic Edition infrastructure on the local cluster. You can also use this command to check whether the Oracle Solaris Cluster Geographic Edition software is enabled on the cluster.

After installation, the Oracle Solaris Cluster Geographic Edition product must be enabled by running the command `geoadm start`. This command enables the cluster to participate in partnerships and to host protection groups. The `geoadm start` command configures the highly available infrastructure that unlocks the Oracle Solaris Cluster Geographic Edition module on the cluster.

The `geoadm stop` command stops the highly available infrastructure that leaves the Oracle Solaris Cluster Geographic Edition module locked on the cluster. Run the `geoadm stop` command before you remove product packages.

**Note** – The `geoadm stop` command can successfully disable the Oracle Solaris Cluster Geographic Edition module on a cluster only under one condition. All local states of the protection groups in which the cluster is involved must be `Offline` or `Unknown`.

Running the `geoadm start` or the `geoadm stop` command on one node of the cluster affects the entire cluster.

To use the `geoadm` command to enable or disable the Oracle Solaris Cluster Geographic Edition infrastructure, you must be assigned the proper role-based access control (RBAC) rights profile.

If you have root access, you have permissions to perform any operation. If you do not have root access, the following RBAC rights apply:

- **Basic Solaris User.** You can read information about Oracle Solaris Cluster Geographic Edition entities by running commands such as `geopg list`, `geohb list`, and `geops list`.
- **Geo Management.** You can perform all the read operations that someone with Basic Solaris User access can perform. You can also perform administrative and configuration tasks such as `geohb create`, `geopg switchover`, `geoadm start`, and `geoadm stop`.

For more information, see the `rbac(5)` man page and “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

The general form of this command is as follows:

```
geoadm [subcommand] [options] [operands]
```

You can omit *subcommand* only if *options* specifies the option `-?`, `-o`, `-V`, or `-v`.

Each option of this command has a long form and a short form. Both forms of each option are given with the description of the option in the OPTIONS section of this man page.

**Subcommands** The following subcommands are supported:

`show`

Display whether the Oracle Solaris Cluster Geographic Edition software is enabled on the cluster.

`start`

Configure and enable the Oracle Solaris Cluster Geographic Edition infrastructure on the cluster.

`status`

Display the runtime status of the Oracle Solaris Cluster Geographic Edition entities on the local cluster.

The Oracle Solaris Cluster Geographic Edition software must be installed on the local cluster before you can run the `status` subcommand. You can run the `status` subcommand whether or not the cluster has been enabled for partnership.

The `status` subcommand displays the following information:

- Whether the local cluster is enabled for partnership.
- Whether the local cluster is involved in a partnership. If the cluster is involved in a partnership, the `status` subcommand lists all partnership members.
- Heartbeat status.
- Protection group status.
- Status of ongoing transactions.

See the EXTENDED DESCRIPTION section for details about the possible values for each status.

`stop`

Disable the Oracle Solaris Cluster Geographic Edition infrastructure and configuration on the cluster.

**Caution** – All protection groups on the cluster must be in the `Offline` state to run the `stop` subcommand successfully.

The `stop` subcommand removes the cluster state and Oracle Solaris Cluster Geographic Edition infrastructure resource groups on the cluster, but the `stop` subcommand does not remove data replication resource groups.

The Oracle Solaris Cluster Geographic Edition infrastructure is enabled and configured again the first time you run the `start` subcommand after running the `stop` subcommand.

**Options** The following options are supported:

-?

--help

Displays help information. When this option is used, no other processing is performed.

You can specify this option without a subcommand or with a subcommand.

If you specify this option without a subcommand, the list of subcommands for this command is displayed.

If you specify this option with a subcommand, the usage options for the subcommand are displayed.

The question mark might be interpreted as a special character by some shells. Use quotes (`-"?"`) or an escape character to avoid pattern matching.

-f

--force

Indicates that you want to bypass the command confirmation questions while the Oracle Solaris Cluster Geographic Edition infrastructure is being disabled.

-i

--interactive

Indicates that you want to display the interactive command confirmation questions while the Oracle Solaris Cluster Geographic Edition infrastructure is being disabled.

-V

--version

Displays the version of the command.

Do not specify this option with subcommands, operands, or other options. The subcommands, operands, or other options are ignored. The `-V` option only displays the version of the command. No other operations are performed.

**Extended Description** The following sections list the status descriptions.

**Partnership Status** Partnership status displays status for the local cluster only. The partnership status can be one of the following:

OK                      The partner clusters are connected.

Error                    The partner clusters are disconnected.

Synchronization Status	<p>Synchronization status displays status for all the clusters in a partnership. The partnership synchronization status can be one of the following:</p> <p>OK                      The configuration is synchronized between partner clusters.</p> <p>Error                     The configuration on the partner clusters is different. You must synchronize the partnership again.</p> <p>Mismatch                The partner clusters have been configured individually. Therefore, you must remove the configuration on one cluster and copy the configuration of the partner cluster.</p> <p>Unknown                 Information is not accessible because the partners are disconnected.</p>
Intercluster Resource Management Connection Status	<p>Intercluster Resource Management (ICRM) Connection status displays status for the cluster management agent that is running on the remote cluster. The ICRM connection status can be one of the following:</p> <p>OK                        The ICRM module on each of the partner clusters can communicate with each other.</p> <p>Error                     The ICRM module on the local cluster is unable to communicate with the ICRM module on the remote cluster.</p>
Heartbeat Status	<p>Heartbeat status displays status for the local cluster only. The heartbeat status can be one of the following:</p> <p>OK                        Heartbeat monitoring is enabled, and the partner cluster is responding within timeout and retry periods.</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>
Heartbeat Plug-in Status	<p>Heartbeat plug-in status displays status for the local cluster only. The heartbeat plug-in status can be one of the following:</p> <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>
Protection Group Status	<p>Protection group status displays the overall status for all the clusters in the protection group. The overall protection group status can be one of the following:</p> <p>OK                        The protection group is online, application resource groups are online, and data replication is running.</p>

	Offline	The protection group is inactive.
	Degraded	The protection group and the application resource groups are online, but data replication is either not running or is in a partial error state.
	Error	The protection group is online, but at least one component of the partnership, such as configuration, data replication, or resource groups, is in an error state.
	Unknown	The protection group is online, but the status for at least one component of the partnership, such as configuration, data replication, or resource groups, is unknown.
Configuration Status	The configuration status of a protection group displays the status for a single cluster only. The configuration status of a protection group can be one of the following:	
	OK	The protection group configuration has been validated without errors on the cluster.
	Error	The validation of the protection group configuration resulted in an error. You must validate the protection group again.
	Unknown	Information is not accessible because the partner clusters are disconnected.
Data Replication Status	Data replication status displays the status of data replication that has been configured for the protection group on a single cluster. The data replication status of a protection group can be one of the following:	
	OK	Data replication is configured, online, and running without errors.
	Degraded	Data replication is either off or in a partial error state.
	Error	Data replication is not operating because of an error.
	None	Data replication has not been configured.
	Unknown	Information is not accessible because the partners are disconnected.
Resource Groups Status	Resource group status displays the status of resource groups that have been configured for the protection group on a single cluster. The resource group status of a protection group can be one of the following:	
	OK	All resource groups are online on the primary cluster, and all resource groups are offline or unmanaged on the secondary cluster.
	Error	Not all resource groups are online on the primary cluster, or not all resource groups are offline or unmanaged on the secondary cluster.
	Unknown	Information is not accessible because the partners are disconnected.

**Exit Status** The following exit values are returned:

0	The command completed successfully.
nonzero	An error has occurred.

**Examples** **EXAMPLE 1** Enabling the Oracle Solaris Cluster Geographic Edition Infrastructure

The following `geoadm` command enables the Oracle Solaris Cluster Geographic Edition control module on the local cluster.

```
# geoadm start
```

**EXAMPLE 2** Disabling the Oracle Solaris Cluster Geographic Edition Infrastructure

The following `geoadm` command removes the infrastructure resource groups that were created when you enabled the Oracle Solaris Cluster Geographic Edition infrastructure.

```
# geoadm stop
```

**EXAMPLE 3** Checking the Status of the Oracle Solaris Cluster Geographic Edition Infrastructure

The following `geoadm` command displays whether the Oracle Solaris Cluster Geographic Edition infrastructure is enabled on the local cluster and, if enable, the nodes on which the infrastructure is active.

```
# geoadm show
```

```
--- CLUSTER LEVEL INFORMATION ---
```

```
Oracle Solaris Cluster Geographic Edition is active on:  
node phys-paris-2, cluster cluster-paris
```

```
Command execution successful
```

```
#
```

**EXAMPLE 4** Monitoring the Runtime Status of the Oracle Solaris Cluster Geographic Edition Software

The following `geoadm` command displays the runtime status of the local Oracle Solaris Cluster Geographic Edition enabled cluster.

```
# 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
```

**EXAMPLE 4** Monitoring the Runtime Status of the Oracle Solaris Cluster Geographic Edition Software *(Continued)*

```

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

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

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

Pending operations
Protection group "tcpg" : Start

```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes.

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	ha-cluster/geo/framework
Interface Stability	Evolving

**See Also** [rbac\(5\)](#), [geohb\(1M\)](#), [geopg\(1M\)](#), [geops\(1M\)](#)



**Name** geohb – configure and manage the heartbeat mechanism

**Synopsis** `/usr/cluster/bin/geohb -?`  
`/usr/cluster/bin/geohb -V`  
`/usr/cluster/bin/geohb add-plugin pluginname heartbeatname -p property [-p...]`  
`/usr/cluster/bin/geohb create -r clusterlist [-p property [-p...]] heartbeatname`  
`/usr/cluster/bin/geohb delete heartbeatname`  
`/usr/cluster/bin/geohb list heartbeatname ...`  
`/usr/cluster/bin/geohb modify-plugin -p property [-p...] pluginname heartbeatname`  
`/usr/cluster/bin/geohb remove-plugin pluginname heartbeatname`  
`/usr/cluster/bin/geohb set-prop -p property [-p...] heartbeatname ...`

**Description** The geohb command enables you to configure and manage the heartbeat mechanism.

A heartbeat is a monitor between two clusters: a requester cluster and a responder cluster. Creating a partnership establishes two heartbeats, one in each direction. For example, a partnership between a primary cluster, `cluster-paris`, and a secondary cluster, `cluster-newyork`, contains two heartbeats. One heartbeat has `cluster-paris` as the requester and `cluster-newyork` as the responder. The other heartbeat has `cluster-newyork` as the requester and `cluster-paris` as the responder.

The Oracle Solaris Cluster Geographic Edition software provides a default heartbeat mechanism that is based on the TCP/UDP plug-in as the primary plug-in and the ping plug-in as backup. The geohb command enables you to configure and maintain heartbeats. You can perform the following tasks:

- Configuring a heartbeat between clusters that participate in a partnership. Configuring a heartbeat includes the configuration of associated plug-ins.
- Creating or deleting a heartbeat property.
- Adding, modifying, and removing plug-ins that are associated with a heartbeat.
- Retrieving the current configuration of a heartbeat and its associated plug-ins.

Run the geohb command on a cluster that has been enabled for partnership.

To run the geohb command to configure and manage the heartbeat mechanism, you must be assigned the proper role-based access control (RBAC) rights profile.

If you have root access, you have permissions to perform any operation. If you do not have root access, the following RBAC rights apply:

- **Basic Solaris User.** You can read information about Oracle Solaris Cluster Geographic Edition entities by running commands such as `geopg list`, `geohb list`, and `geops list`.

- **Geo Management.** You can perform all the read operations that someone with Basic Solaris User access can perform. You can also perform administrative and configuration tasks such as `geohb create`, `geopg switchover`, `geoadm start`, and `geoadm stop`.

For more information, see the [rbac\(5\)](#) man page and “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

The general form of this command is as follows:

```
geohb [subcommand] [options] [operands]
```

You can omit *subcommand* only if *options* specifies the option `-?`, `-o`, `-V`, or `-v`.

Each option of this command has a long form and a short form. Both forms of each option are given with the description of the option in the OPTIONS section of this man page.

**Subcommands** The following subcommands are supported:

`add-plugin`

Add a plug-in to an existing heartbeat. If you specify a custom plug-in, you must also specify the path to your custom plug-in command by using the `Query_cmd` property.

`create`

Create a heartbeat. You can monitor heartbeat status by running the [geoadm\(1M\)](#) command. You must configure the remote cluster to make the heartbeat operational.

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

`delete`

Delete a heartbeat.

`list`

Display existing configuration information.

`modify-plugin`

Modify heartbeat plug-in properties.

`remove-plugin`

Remove a plug-in from a heartbeat.

`set-prop`

Modify heartbeat properties.

**Options** The following options are supported:

`-?`

`--help`

Displays help information. When this option is used, no other processing is performed.

You can specify this option without a subcommand or with a subcommand.

If you specify this option without a subcommand, the list of subcommands for this command is displayed.

If you specify this option with a subcommand, the usage options for the subcommand are displayed.

The question mark might be interpreted as a special character by some shells. Use quotes ("-?") or an escape character to avoid pattern matching.

-p *property*

--property *property*

Specifies the properties of a heartbeat or heartbeat plug-in.

A heartbeat property is assigned a value by using a *name=statement* pair. Multiple properties might be set at one time by using multiple statements.

The values for these properties are assigned at creation and tunable at runtime.

See the EXTENDED DESCRIPTION section for currently defined properties.

-r *clusterlist*

--remote-cluster *clusterlist*

Specifies the name of a remote cluster with which the local cluster should establish heartbeat monitoring.

-V

--version

Displays the version of the command.

Do not specify this option with subcommands, operands, or other options. The subcommands, operands, or other options are ignored. The -V option only displays the version of the command. No other operations are performed.

**Operands** The following operands are supported:

*heartbeatname*

Specifies an identifier for the heartbeat on the local cluster. If you are trying to create a new heartbeat, and the specified identifier already exists, the geohb create command fails.

*pluginname*

Specifies the name of the heartbeat plug-in.

### Extended Description

#### Heartbeat Properties

The following sections list the heartbeat and heartbeat plug-in properties.

You can specify the following heartbeat property:

`Query_interval`

Specifies the frequency between heartbeat status requests in seconds. The plug-in enters emergency mode if three `Query_interval` periods pass without response. The plug-in times out and enters error mode if a further `Query_interval` period passes with no response.

Optional property.

Type: Integer.

Tuning recommendations: The value of this property is assigned at creation and tunable at runtime.

Minimum value: 20 seconds

Maximum value: 300 seconds

Default value: 120 seconds.

Heartbeat Plug-in Properties Heartbeat plug-in properties determine how a heartbeat functions.

`Plugin_properties`

Specifies a property string that is specific to the plug-in.

Optional property.

Type: String.

Tuning recommendations: The value of this property is assigned at creation and tunable at runtime.

Default value: None, except for heartbeats that use the default heartbeat plug-ins, `tcp_udp_plugin` and `ping_plugin`.

For the `tcp_udp_plugin` plug-in, the format of this string is predefined as *remote\_IP\_address/UDP/2084/ipsec*, *remote\_IP\_address/TCP/2084/ipsec*. The *remote\_IP\_address* argument specifies the IP address of the partner cluster. The optional */ipsec* string indicates that the plug-in uses IPsec.

For the `ping_plugin`, the format of this string is predefined as *remote\_IP\_address*, where *remote\_IP\_address* specifies the IP address of the partner cluster.

`Query_cmd`

Specifies the path to the command for a heartbeat status request.

Required property if the plug-in does not specify a predefined plug-in.

Type: String.

Tuning recommendations: The value of this property is assigned at creation and tunable at runtime.

Default value: None.

`Requester_agent`

Specifies the absolute path to requester agent.

Optional property.

Type: String.

Tuning recommendations: The value of this property for the default plug-in should not be tuned except for testing purposes.

Default value: None.

#### Responder\_agent

Specifies the absolute path to the responder agent.

Optional property.

Type: String.

Tuning recommendations: The value of this property for the default plug-in should not be tuned except for testing purposes.

Default value: None.

#### Type

Specifies the type of plug-in. Set to either `Primary` or `Backup`.

Required property.

Type: Enum.

Tuning recommendations: The value of this property is assigned at creation and tunable at runtime.

Default value: None, except for heartbeats with default heartbeat name `ping_plugin`. In this case, the default value is `Backup`.

**Exit Status** The following exit values are returned:

0	The command completed successfully, indicating that the remote cluster is alive.
nonzero	An error has occurred, meaning that the remote cluster did not respond to the heartbeat check.

#### Examples

##### EXAMPLE 1 Creating a Heartbeat

The following `geohb` command creates a heartbeat that is named `paris-to-newyork`, which communicates between the local cluster and the cluster `cluster-newyork`.

```
# geohb create -r cluster-newyork paris-to-newyork
```

##### EXAMPLE 2 Creating a Heartbeat Plug-in

The following `geohb` command creates a heartbeat plug-in that is named `command1` for the heartbeat `paris-to-newyork`.

```
# geohb add-plugin command1 -p Query_cmd=/usr/bin/hb paris-to-newyork
```

**EXAMPLE 3** Modifying a Heartbeat

The following `geohb` command modifies the properties for the default heartbeat between `cluster-paris` and `cluster-newyork`.

```
# geohb set-prop -p Query_interval=60 hb_cluster-paris~cluster-newyork
```

**EXAMPLE 4** Deleting a Plug-in From a Heartbeat

The following `geohb` command deletes the plug-in that is named `command1`, from the heartbeat that is named `paris-to-newyork`.

```
# geohb remove-plugin command1 paris-to-newyork
```

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes.

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	ha-cluster/geo/framework
Interface Stability	Evolving

**See Also** [rbac\(5\)](#), [geops\(1M\)](#)

**Name** geopg – create or manage protection groups

**Synopsis** /usr/cluster/bin/geopg -?

/usr/cluster/bin/geopg -V

/usr/cluster/bin/geopg add-device-group -p *property* [-p...] *devicegroupname*  
*protectiongroupname*

/usr/cluster/bin/geopg add-replication-component -p *property* [-p...] *configurationname*  
*protectiongroupname*

/usr/cluster/bin/geopg add-resource-group *resourcegroup* [,*resourcegroup*...]  
*protectiongroupname*

/usr/cluster/bin/geopg create -s *partnershipname* -o *localrole* [-d *datareplicationtype*]  
[-p *property* [-p...]] *protectiongroupname*

/usr/cluster/bin/geopg delete *protectiongroupname*

/usr/cluster/bin/geopg get -s *partnershipname* [*protectiongroupname*]

/usr/cluster/bin/geopg list [*protectiongroupname*] [...]

/usr/cluster/bin/geopg modify-device-group *devicegroupname* -p *property* [-p...]  
*protectiongroupname*

/usr/cluster/bin/geopg modify-replication-component *configurationname*  
-p *property* [-p...] *protectiongroupname*

/usr/cluster/bin/geopg remove-device-group *devicegroupname* *protectiongroupname*

/usr/cluster/bin/geopg remove-replication-component *configurationname* *protectiongroupname*

/usr/cluster/bin/geopg remove-resource-group *resourcegroup* [,*resourcegroup*...]  
*protectiongroupname*

/usr/cluster/bin/geopg set-prop -p *property* [-p...] *protectiongroupname*

/usr/cluster/bin/geopg start -e {local | global} [-n] *protectiongroupname*

/usr/cluster/bin/geopg stop -e {local | global} [-D] *protectiongroupname*

/usr/cluster/bin/geopg switchover -m *newprimaryclustername* [-f] *protectiongroupname*

/usr/cluster/bin/geopg takeover [-f] *protectiongroupname*

/usr/cluster/bin/geopg update *protectiongroupname*

/usr/cluster/bin/geopg validate *protectiongroupname*

**Description** The geopg command enables you to configure and maintain protection groups. You can perform the following tasks:

- Configure a protection group between clusters that participate in a partnership. This task includes the configuration of associated data replication parameters.
- Add or remove resource groups and data-replication disk device groups.

- Add, modify, and remove data replication parameters that are associated with a protection group property.
- Retrieve the current configuration of a specific protection group or all defined protection groups.
- Activate a protection group.
- Deactivate a protection group.
- Switch over the role of a protection group.
- Take over the primary role of a protection group.

Before you create a protection group, ensure that the clusters that will host the protection group are already placed in a partnership.

To run the `geopg` command to create, modify, or delete a protection group, or a data-replication disk device group, you must be assigned the proper role-based access control (RBAC) rights profile.

If you have root access, you have permission to perform any operation. If you do not have root access, the following RBAC rights apply:

- **Basic Solaris User.** You can read information about Oracle Solaris Cluster Geographic Edition (Geographic Edition) entities by running commands such as `geopg list`, `geohb list`, and `geops list`.
- **Geo Management.** You can perform all the read operations that someone with Basic Solaris User access can perform. You can use commands `geohb create`, `geopg switchover`, `geoadm start`, and `geoadm stop` to perform administrative and configuration operations.

For more information, see the `rbac(5)` man page and “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

The general form of this command is as follows:

```
geopg [subcommand] [options] [operands]
```

You can omit *subcommand* only if *options* specifies the option `-?`, `-o`, `-V`, or `-v`.

Each option of this command has a long form and a short form. Both forms of each option are given with the description of the option in the OPTIONS section of this man page.

#### How Protection Groups Start

The `geopg start` command activates the protection group on both the primary and standby clusters. This activation starts the Geographic Edition management of the resource groups in the protection group. Depending on the form of the command that you use and on the cluster where you issue the command, activating the protection group might not start the resource groups on that cluster. The resource groups start, or are brought online, only on the primary cluster.



The `-e` option defines the scope of the `geopg start` command. If you specify `-e local`, the `geopg start` command runs on the cluster where you issue the command. If you specify `-e global`, the `geopg start` command runs on both clusters in the partnership.

The different forms of the command have the following effects:

- If you run the `geopg start -e local` command on the primary cluster, and if you run the `geopg start -e global` command on either cluster, the Geographic Edition software activates the protection group on both clusters in the partnership. The software brings online resource groups on the primary cluster only. The resource groups are put in a managed state.
- If you run the `geopg start -e local` command on the standby cluster, the Geographic Edition software activates the protection group on the standby cluster only. Resource groups are not started on the standby cluster. The resource groups are put in the Unmanaged state.

**Subcommands** The following subcommands are supported:

`add-device-group`

Enables a data-replication disk device group to be part of a protection group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

`add-replication-component`

Enables a data-replication component to be part of a protection group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

`add-resource-group`

Enables an application resource group to be part of a protection group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

If a protection group is active when you add a resource group, the resource group must be in either the `Unmanaged` or `Online` state. To start an unmanaged resource group, run the `geopg start` command.

If a protection group is inactive, the resource group must be in the `Unmanaged` state when you add the resource group. If the resource group is offline, bring the resource group to the `Unmanaged` state before adding the resource group to an inactive protection group.

`create`

Creates a protection group within an existing partnership. The system performs this action on only the local cluster.

`delete`

Deletes a protection group. The system performs this action on only the local cluster.

`get`

Creates the local configuration for a protection group that already exists on the remote partner cluster by retrieving the configuration from the partner cluster. If you do not

specify a protection group name, all the protection groups that have been configured for the partnership on the remote cluster will be created on the local cluster. The system performs this action on only the local cluster.

#### list

Displays the following information about the protection group:

- Defined protection groups
- Resource groups that are wrapped into protection groups
- Operation status information

If you do not specify the name of a protection group, the `list` subcommand displays information for all the protection groups configured on the local cluster. The system performs this action on only the local cluster.

#### modify-device-group

Modifies the properties of a data-replication disk device group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

#### modify-replication-component

Modifies the properties of a data-replication component. The system performs this action on the local cluster, then propagates the action to the partner cluster.

#### remove-device-group

Removes data-replication disk device groups from a protection group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

Removing a data-replication disk device group does not stop data replication or change the data-replication status for that data-replication disk device group.

#### remove-replication-component

Removes data-replication components from a protection group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

Removing a data-replication component does not stop data replication or change the data-replication status for that data-replication component.

#### remove-resource-group

Removes resource groups from a protection group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

Removing an application resource group does not change the application resource group to the Unmanaged state.

#### set-prop

Modifies the properties of a protection group. The system performs this action on the local cluster, then propagates the action to the partner cluster.

**start**

Starts a protection group. When you start a protection group for the first time, the role of the protection group on a cluster is the role that was assigned to the protection group when you created the protection group. After that, the protection group starts with the latest role it was given.

The `start` subcommand uses Oracle Solaris Cluster commands to bring resources and resource groups under the management of the protection group.

You can start a protection group on the following levels, depending on whether you specify `local` or `global` to the `-e` option:

- With the `-e global` option, starts the protection group on all clusters where the protection group has been configured
- With the `-e local` option, starts the protection group on either the primary cluster or standby cluster, whichever cluster the command is issued from

For more details about activating a protection group and the effects on the resource groups in that protection group, see “How Protection Groups Start” in the DESCRIPTION section of this man page.

Starting a protection group enables the following events to occur:

- The protection group configuration is validated.
- Clusters can determine whether an operation has been completed.
- Data replication can be started, depending on the data replication that you use.
  - When using Oracle Solaris Availability Suite (Availability Suite) software, you can start data replication from the primary cluster only.
  - When using EMC Symmetrix Remote Data Facility (SRDF) software, you can start data replication from either the primary or standby cluster.
  - When using the Oracle Data Guard software, you can start data replication from the primary cluster by using the local or the global scope.

For all supported data replication software, you must not specify the `-n` option when using this form of the command.

- Protected applications can be started.
- If the partner can be reached, the partner cluster can be notified of the protection group activation.

**Note** – Protection groups that are referred to as “online” are active protection groups.

**stop**

Stops a protection group. You can stop a protection group on the following levels, depending on whether you specify `local` or `global` to the `-e` option:

- With the `-e global` option, stops the protection group on all clusters where the protection group has been configured
- With the `-e local` option, stops the protection group on either the primary or standby cluster, whichever cluster the command is issued from

Stopping a protection group enables the following events to occur:

- The protection group becomes inactive.
- Data replication and application resource groups stop. You can prevent application resource groups from stopping by taking the resource groups out of the protection group *before* you stop the protection group.
  - When using Availability Suite software, you can stop data replication from the primary cluster only.
  - When using EMC Symmetrix Remote Data Facility (SRDF) software, you can stop data replication from either the primary or standby cluster.
  - When using Oracle Data Guard, you can stop data replication from either the primary or standby cluster, by using either the local or the global scope.
- If the partner cluster is reachable, the partner cluster is notified that the protection group has been stopped.

#### `switchover`

Switches the assigned role of a cluster in the protection group. You must start the protection group before you can perform a switchover. The system performs this action on the local cluster, then propagates the action to the partner cluster.

#### `takeover`

Forces a cluster to become the primary cluster without considering the partner cluster state. The system performs this action on the local cluster, then propagates the action to the partner cluster.

If you issue a takeover from the standby cluster and the cluster is able to communicate with the partner cluster, the Geographic Edition software switches the role of the partner cluster so that the partner cluster becomes the new standby. This operation is not the same as a switchover. The protection group is deactivated on the new standby.

After the `geopg takeover` command successfully completes, reactivating the protection group on the standby cluster might require that you recover and synchronize data.

**Note** – When possible, run the `geopg switchover` command instead of the `geopg takeover` command to coordinate between partner clusters to avoid loss of replicated data.

Run the `geopg takeover` command only in situations when the primary cluster is not available or during unplanned downtime. The situation must justify a potential loss of data and the time required to repair and to reactivate the protection group on the standby cluster.

**update**

Resynchronizes the configuration information of the local protection group with the partner's configuration information. The `geopg update` command overwrites the configuration information of the protection group on the local cluster with the configuration information from the partner cluster. This command does not overwrite replicated data. The system performs this action on only the local cluster.

**validate**

Validates a protection group on the local cluster by performing a sanity check of the dynamic values. The system performs this action on only the local cluster.

**Options** The following options are supported:

-?

--help

Displays help information. When this option is used, no other processing is performed.

You can specify this option with or without a subcommand.

If you specify this option without a subcommand, the list of subcommands for this command is displayed.

If you specify this option with a subcommand, the usage options for the subcommand are displayed.

The question mark might be interpreted as a special character by some shells. Use quotes (`-"?"`) or an escape character to avoid pattern matching.

-D

--only-datarep

Specifies that only the data replication should be deactivated, leaving the protection group active. If you do not use this option, the entire protection group is deactivated. To stop a protection group that has already had its data replication subsystem stopped, you must run the `geopg stop` command again and omit this option.

-d *datareplicationtype*

--datarep-type *datareplicationtype*

Specifies the data replication mechanism for data replication between the clusters of the protection group.

The *datareplicationtype* must be one of the following strings:

**avs** Specifies that the data replication mechanism is Availability Suite Remote Mirror, with patches.

**odg** Specifies that the data replication mechanism is Oracle Data Guard.

**srdf** Specifies that the data replication mechanism is the dynamic EMC Symmetrix Remote Data Facility (SRDF).

- e {local|global}
- scope {local|global}  
Specifies whether the command operates only on the local cluster, `local`, or on both clusters, `global`, where the protection group has been configured.
- f
- force  
Forces the command to perform the operation without asking you for confirmation.
- m *newprimarycluster*
- primary *newprimarycluster*  
Specifies the name of the cluster that is to be the new primary cluster for the protection group.
- n
- nodatarep  
This option is only valid when starting a protection group that is not active. If the protection group is not active, it is started without data replication. If data replication for this protection group is already running independently, it is stopped before the protection group is started.  
  
If the protection group is already active with data replication enabled, data replication is not stopped. An error message is displayed that states that the `geopg stop -D` command should be used to stop data replication for an active protection group.  
  
If this option is omitted, data replication starts at the same time as the protection group.
- o *localrole*
- role *localrole*  
Specifies the role of the local cluster as either `primary` or `secondary` when used with the `create` subcommand. You can change the role of a cluster by running the `geopg switchover` command.
- p *property*
- property *property*  
Sets the properties of a protection group.  
  
A protection group property is assigned a value by using a *name=statement* pair statement. You can set multiple properties at one time by using multiple statements.  
  
See the EXTENDED DESCRIPTION section for defined properties.
- s *partnershipname*
- partnership *partnershipname*  
Specifies the name of the existing partnership where the protection group is configured.
- V
- version  
Displays the version of the command.

Do not specify this option with subcommands, operands, or other options, as they are ignored. The `-V` option only displays the version of the command. No other operations are performed.

**Operands** The following operands are supported:

*configurationname*

Specifies the name of any device group or replicated component that you want to add.

*devicegroupname*

Specifies the name of the disk device group to be created or added to the protection group.

*protectiongroupname*

Specifies the name of the protection group. The `create` subcommand fails if the protection group that you specify with this option already exists. The `modify` subcommand fails if the specified protection group is not valid.

*resourcegroupname*

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

If a protection group is active when you add a resource group, the resource group must be in either the `Unmanaged` or `Online` state. To start an unmanaged resource group, run the `geopg start` command.

If the protection group is inactive when you add a resource group, the resource group must be in the `Unmanaged` state before you add the resource group.

**Extended Description** This section contains descriptions of properties that you can specify.

**Note** – Property values, such as `true` and `false`, are *not* case sensitive.

General Protection  
Group Properties

Description (string)

Describes the protection group. The system sets this property on the local cluster, then propagates the value to the partner cluster.

**Category:**

Optional

**Default:**

None

**Tunable:**

Assigned at creation and tunable at runtime

`External_Dependency_Allowed` (boolean)

Specifies whether to allow any dependencies between resource groups and resources that belong to this protection group and resource groups and resources that do not belong to this protection group. The system sets this property on the local cluster, then propagates the value to the partner cluster.

**Category:**  
Optional

**Default:**  
False

**Tunable:**  
Any time

**RoleChange\_ActionArgs (string)**

Specifies the arguments that follow the system-defined arguments that are paired with the role change callback script during a protection group switchover or takeover. The system sets this property on the local cluster, then propagates the value to the partner cluster.

**Category:**  
Optional

**Default:**  
None

**Tunable:**  
Assigned at creation and tunable at runtime

**RoleChange\_ActionCmd (string)**

Specifies the absolute path to the executable file to run when the primary cluster of the protection group changes. This file should exist on all nodes of both partner clusters that host the protection group. This script is run before the application resource group is brought online on the new primary cluster. The system sets this property on the local cluster, then propagates the value to the partner cluster.

**Category:**  
Optional

**Default:**  
None

**Tunable:**  
Assigned at creation and tunable at runtime

**Timeout (integer, but you must specify a value equal to or greater than 20)**

Specifies, in seconds, the longest time that the Geographic Edition software waits for a response after running a geopg command. Examples of responses include `start`, `stop`, `switchover`, and `takeover`. If the command does not return within the `timeout` period, the Geographic Edition software reports the operation as timed out, even if the underlying command that was run eventually completes successfully.

The `timeout` period applies to operations on a per-cluster basis. An operation with a local scope times out if the operation is not completed after the specified `timeout` period.

An operation with a global scope consists of an action on the local cluster and an action on the remote cluster. The local and remote 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.

For example, the following command is started with a local scope:

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

If you set the `timeout` property 3000 seconds, the `geopg start` command times out if the operation does not complete after 3000 seconds.

You can start the same command with a global scope as follows:

```
# geopg start -e global protectiongroupname
```

If the `timeout` property is set to 3000 seconds, the `geopg start` 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.

The protection group `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.

You can specify a maximum of 1,000,000 (one million) seconds for the `timeout` property.

**Category:**

Optional

**Default:**

3600

**Tunable:**

Assigned at creation and tunable at runtime

EMC Symmetrix  
Remote Data Facility  
Data Replication Type  
Properties

Data replication property: `Cluster_dgs` (stringarray)

Lists the disk device groups where the data is written. The disk device groups you specify must be of the same type. You can specify more than one disk device group by separating each disk device group with a comma. The system sets this property on the local cluster, then propagates the value to the partner cluster.

To specify Solaris Volume Manager disk sets when a disk set name is not the same on both clusters, use the following syntax which includes the cluster names:

```
Cluster_dgs=local-cluster[diskset]:remote-cluster[diskset]
```

Include the brackets ([ ]) around all disk set names for each cluster. Separate multiple disk set names with a comma. For example:

```
Cluster_dgs=paris[dg1, dg2]:newyork[dg10, dg11]
```

**Category:**  
Optional

**Default:**  
Empty

**Tunable:**  
Only when the protection group is inactive on both partner clusters

Data replication property: `DG_or_CG` (string)

Specifies the device group identifier as “SRDF Device Group.” You must set this property to DG. The system sets this property on the local cluster, then propagates the value to the partner cluster.

**Category:**  
Optional

**Default:**  
DG

**Tunable:**  
Never

Data replication property: `NodeList` (stringarray)

Lists the host names of the machines that can be primary for the replication mechanism. This list is comma delimited. The system sets this property value on only the local cluster.

**Category:**  
Optional

**Default:**  
Empty

**Tunable:**  
Any time

Data replication property: `R1_SID` (integer)

Specifies the Symmetrix Remote Data Facility (SRDF) identification of the primary devices. The value also set on the partner cluster as the `R2_SID` value.

**Note** – If you do not specify a value for this property when you add a device group, the Geographic Edition software searches for the value and sets the property automatically.

**Category:**  
Required

**Default:**  
None, until you add a Symmetrix Remote Data Facility (SRDF) device group

**Tunable:**  
Any time

Data replication property: `R2_SID` (integer)

Specifies the Symmetrix Remote Data Facility (SRDF) identification of the secondary devices. The value is also set on the partner cluster as the `R2_SID` value.

**Note** – If you do not specify a value for this property when you add a device group, the Geographic Edition software searches for the value and sets the property automatically.

**Category:**

Required

**Default:**

None, until you add a Symmetrix Remote Data Facility (SRDF) device group

**Tunable:**

Any time

Oracle Data Guard Data  
Replication Type  
Properties

Data replication property: `local_database_name` (string)

Name of the local Oracle database in the Oracle Data Guard Broker configuration that is being replicated to the remote cluster. This name is the Oracle `db_unique_name` initialization parameter for the Oracle RAC database on the local cluster. The name is also set on the partner cluster as the `remote_database_name` value.

**Category:**

Required

**Default:**

None

**Tunable:**

At creation

Data replication property: `local_db_service_name` (string)

Oracle net service name that is used to connect to the local Oracle database. The name is also set on the partner cluster as the `remote_db_service_name` value.

**Category:**

Required

**Default:**

None

**Tunable:**

Any time

Data replication property: `local_rac_proxy_svr_rg_name` (string)

Name of the local Oracle RAC server proxy resource group that manages the local database in the Oracle Data Guard Broker configuration. The name is also set on the partner cluster as the `remote_rac_proxy_svr_rg_name` value.

**Category:**

Required

**Default:**

None

**Tunable:**

At creation

Data replication property: `remote_database_name` (string)

Name of the remote database in the Oracle Data Guard Broker configuration that is being replicated from the local cluster. This name is the Oracle `db_unique_name` initialization parameter for the Oracle RAC database on the remote cluster. The name is also set on the local cluster as the `local_database_name` value.

**Category:**

Required

**Default:**

None

**Tunable:**

At creation

Data replication property: `remote_db_service_name` (string)

Oracle net service name that is used to connect to the remote Oracle database. The name is also set on the local cluster as the `local_db_service_name` value.

**Category:**

Required

**Default:**

None

**Tunable:**

Any time

Data replication property: `remote_rac_proxy_svr_rg_name` (string)

Name of the remote Oracle RAC server proxy resource group on the partner cluster that manages the remote database in the Oracle Data Guard Broker configuration. The name is also set on the local cluster as the `local_rac_proxy_svr_rg_name` value.

**Category:**

Required

**Default:**

None

**Tunable:**

At creation

Data replication property: `replication_mode` (string)

The Oracle Data Guard replication mode between the primary database and the standby database. The system sets this property on the local cluster, then propagates the value to the partner cluster.

Valid values to which you set this property include `maxAvailability`, `maxPerformance`, and `maxProtection`.

**Category:**

Required

**Default:**

None

**Tunable:**

Any time

Data replication property: `standby_type` (string)

Type of Oracle standby database that is used in the Oracle Data Guard Broker configuration. The system sets this property on the local cluster, then propagates the value to the partner cluster.

Valid values to which you set this property include `logical` and `physical`.

**Category:**

Required

**Default:**

None

**Tunable:**

At creation

Data replication property: `sysdba_password` (string)

Password for the Oracle SYSDBA privileged database user. The system sets this property on the local cluster, then propagates the value to the partner cluster.

Do not specify a password on the command line. If you specify only

`-p sysdba_password=`, the `geopg` command prompts you to type an actual password, which is not displayed as you type it.

**Category:**

Required

**Default:**

None

**Tunable:**

Any time

Data replication property: `sysdba_username` (string)

Name of an Oracle SYSDBA privileged database user who can perform the Oracle Data Guard Broker switchover and takeover operations on both the primary and standby clusters. The Oracle Solaris Cluster Geographic Edition software uses this property to monitor and manage the Oracle Data Guard Broker configurations. The system sets this property on the local cluster, then propagates the value to the partner cluster.

**Category:**

Required

**Default:**

None

**Tunable:**

Any time

Oracle Solaris  
Availability Suite Data  
Replication Type  
Properties

Device group property: `Enable_volume_set` (boolean)

Determines whether the volume sets that are defined in the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file, and the fallback snapshots that are defined in the `/var/cluster/geo/avs/AVS-devicegroup-snapshot.ini` file are to be enabled by the Oracle Solaris Cluster Geographic Edition software. You can set this property to `True` or to `False`. The system sets this property on the local cluster, then propagates the value to the partner cluster.

If the `enable_volume_set` property is set to `True`, volume sets that are defined in the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file are enabled when the device group is added to the protection group or when the protection group that contains the device group is replicated from the partner cluster. Any volume sets of that device group that are not defined in the file are disabled. The `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file must exist on every node of the primary and standby cluster. In addition, if the `enable_volume_set` property is set to `True`, the fallback snapshots defined in the `/var/cluster/geo/avs/AVS-devicegroup-snapshot.ini` file are enabled. A fallback snapshot is a compact dependent shadow volume that is created immediately prior to the resynchronization of a secondary replicated volume in the Oracle Solaris Cluster Geographic Edition software. If resynchronization fails, the secondary replicated volume can be reconstructed by using the fallback snapshot. The `/var/cluster/geo/avs/AVS-devicegroup-snapshot.ini` file must exist on any node of a cluster on which a fallback snapshot will be automatically enabled. Fallback snapshot configuration is optional.

If the `Enable_volume_set` property is set to `False`, you must manually enable volume sets and fallback snapshots by using the Availability Suite commands.

The parameters in the volume set file that are handled by the Geographic Edition software are described in the following table.

**Note** – The Geographic Edition software does not handle other parameters of the volume set file, such as, size of memory queue, and number of asynchronous threads. You must adjust these parameters manually by using Availability Suite commands.

Parameter	Description
-----------	-------------

---

phost	Primary host. The logical host of the server on which the primary volume resides.
pdev	Primary device. Primary volume partition. Specify full path names only.
pbitmap	Primary bitmap. Volume partition in which the bitmap of the primary partition is stored. Specify full path names only.
shost	Secondary host. The logical host of the server on which the secondary volume resides.
sdev	Secondary device. Secondary volume partition. Specify full path names only.
sbitmap	Secondary bitmap. Volume partition in which the bitmap of the secondary partition is stored. Specify full path names only.
ip	Network transfer protocol. IP address.
sync   async	Operating mode. <ul style="list-style-type: none"> <li>■ In sync mode the I/O operation is confirmed as complete only when the volume on the standby cluster has been updated.</li> <li>■ In async mode the primary host I/O operation is confirmed as complete before updating the volumes on the standby clusters.</li> </ul>
g iogroupname	An I/O group name. The set must be configured in the same I/O group on both the primary and standby clusters. This parameter is optional and need only be configured if you have an I/O group.
q qdev	Disk queue volume. Full path name for the volume to be used as a disk-based I/O queue for an asynchronous disk set.
C devicegroupname	Device group name. The device group name or resource tag of the local data and bitmap volumes in cases where this information is not implied by the name of the volume. For example, <code>/dev/md/avsset/rdisk/vol</code> indicates a device group named <code>avsset</code> . It can be Solaris Volume Manager volumes or raw device volumes.

---

The parameters required to enable a fallback snapshot for a replicated volume are shown in the following example line from a *AVS-devicegroup-snapshot.ini* configuration file. This line defines one compact dependent shadow volume that is to be created immediately prior to a resynchronization of the listed secondary volume. Each line of the *AVS-devicegroup-snapshot.ini* file configures a fallback snapshot for one secondary volume in the device group in the format shown in the following example:

```
/dev/md/avsset/rdisk/d100 /dev/md/avsset/rdisk/d102 /dev/md/avsset/rdisk/d103
```

Each line in the *AVS\_devicegroup-snapshot.ini* file contains the following types of entries:

- `/dev/md/avsset/rdisk/d100` – Secondary replicated volume
- `/dev/md/avsset/rdisk/d102` – Fallback snapshot volume

- /dev/md/avset/rdisk/d102 – Fallback snapshot bitmap

**Category:**

Optional

**Default:**

False

**Tunable:**

Until you add the device group to a protection group

Device group property: `Local_logical_host` (string)

Defines the local logical host name that is used for the replication of the disk device group.

Do not use an underscore character (`_`) in the logical host name. The name is also set on the partner cluster as the `Remote_logical_host` value.**Category:**

Required

**Default:**

None

**Tunable:**

Until you add the device group to a protection group

Data replication property: `NodeList` (stringarray)

Lists the host names of the machines that can be primary for the replication mechanism.

You can specify more than one host name by separating each host name with a comma.

The system sets this property value on only the local cluster.

**Category:**

Optional

**Default:**

All nodes in the cluster

**Tunable:**

Only when the protection group is inactive

Device group property: `Remote_logical_host` (string)Specifies the remote logical host name that is used for the replication of the disk device group. The name is also set on the partner cluster as the `Local_logical_host` value.**Category:**

Required

**Default:**

None

**Tunable:**

Until you add the device group to a protection group



**Exit Status** The following exit values are returned:

0	The command completed successfully.
nonzero	An error occurred.

**Examples** **EXAMPLE 1** Adding an Oracle Data Guard Broker Configuration to a Protection Group

The following `geopg` command adds the Oracle Data Guard Broker configuration `mysales.com` to the protection group `sales-pg`.

To ensure security, do *not* supply a password when you specify the `sysdba_password` property. If you do not provide a password, the `geopg` command prompts for one.

```
phys-paris-1# geopg add-replication-component \
  -p local_database_name=sales \
  -p remote_database_name=salesdr \
  -p local_db_service_name=sales-svc \
  -p remote_db_service_name=salesdr-svc \
  -p standby_type=physical \
  -p replication_mode=MaxPerformance \
  -p sysdba_username=sys \
  -p sysdba_password= \
  -p local_rac_proxy_svr_rg_name=sales-rac-proxy-svr-rg \
  -p remote_rac_proxy_svr_rg_name=salesdr-rac-proxy-svr-rg \
  mysales.com sales-pg
```

Oracle Data Guard configuration "mysales.com" successfully added to the protection group "sales-pg"

**EXAMPLE 2** Creating a Protection Group That Uses Data Replication

The following `geopg` command creates a Availability Suite protection group `avspg` on the primary cluster `cluster-paris`.

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

**EXAMPLE 3** Creating a Protection Group That Does Not Use Data Replication

The following `geopg` command creates a protection group `example-pg` that is not configured to use data replication.

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

**EXAMPLE 4** Deleting a Protection Group While Keeping Application Resource Groups Online

The following `geopg` command keeps two application resource groups, `apprg1` and `apprg2`, online while deleting their protection group, `avspg`. Remove the application resource groups from the protection group, then delete the protection group.

**EXAMPLE 4** Deleting a Protection Group While Keeping Application Resource Groups Online  
(Continued)

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

**EXAMPLE 5** Modifying a Protection Group

The following `geopg` command modifies the timeout property of the protection group `avspg`.

```
# geopg set-prop -p Timeout=3000 avspg
```

**EXAMPLE 6** Adding a Data Replication Device Group to an Availability Suite Protection Group

The following `geopg` command creates an Availability Suite data replication device group in the `avspg` protection group.

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

If the command returns successfully, volume sets that are defined in the `/var/cluster/geo/avs/avset-volset.ini` file are enabled. The file might look like this for a Solaris Volume Manager device group.

```
lh-paris-1 /dev/md/avset/rdisk/d100 /dev/md/avset/rdisk/d101
lh-newyork-1 /dev/md/avset/rdisk/d100 /dev/md/avset/rdisk/d101
ip async q /dev/md/avset/rdisk/d102 C avset
```

The description of the parameters is as follows:

- `lhparis-1` — Primary host
- `/dev/md/avset/rdisk/d100` — Primary data
- `/dev/md/avset/rdisk/d101` — Primary bitmap
- `lh-newyork-1` — Secondary host
- `/dev/md/avset/rdisk/d100` — Secondary data
- `/dev/md/avset/rdisk/d101` — Secondary bitmap
- `ip` — Protocol
- `async` — Mode
- `q /dev/md/avset/rdisk/d102` — disk queue
- `C` — C tag
- `avset` — Device group

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes.

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	ha-cluster/geo/framework

---

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

**See Also** [geoadm\(1M\)](#), [geohb\(1M\)](#), [geops\(1M\)](#), [rbac\(5\)](#)

**Name** geops – create, configure, and manage partnerships between clusters

**Synopsis** /usr/cluster/bin/geops [*subcommand*] -?  
/usr/cluster/bin/geops -V  
/usr/cluster/bin/geops add-trust -c *remotepartnerclustername*  
/usr/cluster/bin/geops create -c *remotepartnerclustername* [-h *heartbeatname*]  
[-p *property* [-p...]] *partnershipname*  
/usr/cluster/bin/geops join-partnership *remoteclustername* [-h *heartbeatname*]  
*partnershipname*  
/usr/cluster/bin/geops leave-partnership *partnershipname*  
/usr/cluster/bin/geops list *partnershipname*  
/usr/cluster/bin/geops remove-trust -c *remotepartnerclustername*  
/usr/cluster/bin/geops set-prop -p *property* [-p...] *partnershipname*  
/usr/cluster/bin/geops update *partnershipname*  
/usr/cluster/bin/geops verify-trust [-c *remotepartnerclustername*]

**Description** The geops command enables you to create, configure, and manage the partnerships that are defined between clusters. A partnership is a pair of clusters that define a cluster infrastructure on which an application might be protected against disaster.

Two clusters must be configured to trust each other before you create a partnership between two clusters.

A partnership requires a running heartbeat between clusters. Partner clusters monitor each other with heartbeats.

The Oracle Solaris Cluster Geographic Edition software notifies all remote partners of changes in the partnership configuration.

Run the geops command on a cluster that has been enabled for partnership.

To run the geops command to create, configure, and manage the partnerships, you must be assigned the proper role-based access control (RBAC) rights profile.

If you have root access, you have permissions to perform any operation. If you do not have root access, the following RBAC rights apply:

- **Basic Solaris User.** You can read information about Oracle Solaris Cluster Geographic Edition entities by running commands such as `geopg list`, `geohb list`, and `geops list`.
- **Geo Management.** You can perform all the read operations that someone with Basic Solaris User access can perform. You can also perform administrative and configuration tasks such as `geohb create`, `geopg switchover`, `geoadm start`, and `geoadm stop`.

**Note** – You must have root access to run the following subcommands:

- `geops add-trust`
- `geops remove-trust`
- `geops verify-trust`

For more information, see the [rbac\(5\)](#) man page and “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

The general form of this command is as follows:

```
geops [subcommand] [options] [operands]
```

You can omit *subcommand* only if *options* specifies the option `-?`, `-o`, `-V`, or `-v`.

Each option of this command has a long form and a short form. Both forms of each option are given with the description of the option in the OPTIONS section of this man page.

**Subcommands** The following subcommands are supported:

`add-trust`

Retrieve public keys of a remote cluster and add the keys to the `truststore` file on every node of the local cluster. The `truststore` file is at `/etc/caao/instances/default/security/jsse/truststore`. The system performs this action on only the local cluster.

The new key is not added to the `truststore` file if an entry with the same alias already exists in the `truststore` file on a node. If the existing entry contains a key that is different from the new key, you must delete the entry from the `truststore` file before you add the new key.

`create`

Create a partnership between the local cluster and a remote cluster. The two clusters must be configured to trust each other. The system performs this action on only the local cluster.

`join-partnership`

Enable the local cluster to join an existing partnership. The system performs this action on the local cluster, then propagates the action to the partner cluster.

`leave-partnership`

Remove the local cluster from a partnership. The system performs this action on the local cluster and notifies the partner cluster that the local cluster has left the partnership. This subcommand deletes the partnership when the last participating cluster in the partnership leaves.

`list`

Display partnership configuration information. The system performs this action on only the local cluster.

**remove -t trust**

Remove all keys for the remote cluster from the `truststore` file on the local node. You must remove all keys for the remote cluster from the local host before you can add a new key. The system performs this action on only the local cluster.

To remove a key from a cluster, dissolve the partnership between the two clusters (by running `geops leave` on each cluster) and run the `remove -t trust` subcommand from each node of the cluster.

**set -prop**

Modify the properties of a partnership. This subcommand updates the local cluster configuration and notifies remote partner clusters of the configuration change.

**update**

Synchronize information with the partner cluster. The `geops update` command overwrites the configuration on the local cluster with the configuration from the remote partner cluster.

This subcommand cannot update a partnership while the cluster is disconnected from the partner cluster.

If the remote partner cluster is down, modifications to the local cluster might not be propagated to the remote partner cluster.

**verify -t trust**

Retrieve node information on the remote cluster and establish a secure Java Management Extensions (JMX) connection from the local host to every node of the remote cluster. The `verify -t trust` subcommand verifies that the certificates have been exchanged properly between the remote and the local clusters.

If you do not specify a remote cluster, the `verify -t trust` subcommand verifies the connections between the local host and other nodes of the local cluster.

**Options** The following options are supported:

`-?`

`--help`

Displays help information. When this option is used, no other processing is performed.

You can specify this option without a subcommand or with a subcommand.

If you specify this option without a subcommand, the list of subcommands for this command is displayed.

If you specify this option with a subcommand, the usage options for the subcommand are displayed.

The question mark might be interpreted as a special character by some shells. Use quotes (`-"?"`) or an escape character to avoid pattern matching.

*-c remoteclustername*

*--cluster remoteclustername*

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

```
cluster-paris
```

When you use this option with the *add-trust*, *remote-trust*, or *verify-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 must have the following pattern:

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

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

**Note** – The *remove-trust* subcommand removes only keys you have added with the *add-trust* subcommand, and keys with an alias that matches the following pattern:

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

*-h heartbeatname*

*--heartbeat-name heartbeatname*

Specifies an identifier for the heartbeat on a partner cluster that the local cluster can use to monitor partner availability. You must create the heartbeat by running the *geohb* command before you specify the heartbeat in the *geops* command.

If this option is omitted, the Oracle Solaris Cluster Geographic Edition software uses the default heartbeat mechanism between member clusters.

*-p property*

*--property property*

Specifies the value of partnership properties. You can specify multiple properties at one time by using multiple statements.

See the EXTENDED DESCRIPTION section for a description of the properties.

*-V*

*--version*

Displays the version of the command.

Do not specify this option with subcommands, operands, or other options. The subcommands, operands, or other options are ignored. The *-V* option only displays the version of the command. No other operations are performed.

**Operands** The following operand is supported:

*partnershipname* Specifies the name of the partnership

**Extended Description** The following section lists the partnership properties. These properties are set at creation and tunable at runtime.

Partnership Properties Description

Describes the partnership. The system sets this property on the local cluster, then propagates the value to the partner cluster.

Optional property.

Type: String.

Default value: Empty string.

**Notification\_ActionCmd**

Specifies the path to the script or command that is triggered when a heartbeat-loss notification is issued. The system sets this property on the local cluster, then propagates the value to the partner cluster.

Optional property.

Type: String.

Default value: Empty string.

The Oracle Solaris Cluster Geographic Edition software enables you to specify a command to run when a heartbeat-loss notification is issued. You can specify the path to the command by using the `Notification_ActionCmd` property. The command runs with root permissions, so the file must have root ownership and execution permissions. If both `Notification_ActionCmd` and `Notification_EmailAddr` properties have been configured, an email is sent after the command that is specified in the `Notification_ActionCmd` property is run.

**Caution** – A heartbeat loss indicates only a lost connection to a remote cluster and not the cause for the lost connection. A remote cluster that has lost a heartbeat connection might still be functioning even though the heartbeat monitoring has lost its connection.

Therefore, you must exercise caution in the script you specify for the `Notification_actioncmd` parameter. For example, if you specify a forced takeover or a switchover in the `Notification_actioncmd` script, a takeover or a switchover might be run even though one might not be necessary.

**Notification\_EmailAddr**

Specifies the email address to send messages to when a heartbeat-loss notification is issued. You can specify multiple email addresses by separating each email address with a comma. The system sets this property on the local cluster, then propagates the value to the partner cluster.



Optional property.

Type: String array.

Default value: Empty string.

The `Notification_EmailAddrs` and the `Notification_ActionCmd` properties enable notification in heartbeat-loss events. Heartbeat-loss events are detected locally on each cluster of the partnership, and the notification is triggered locally on the cluster where the event is detected. The email addresses and the notification action path should be valid on each cluster in the partnership.

**Exit Status** The following exit values are returned:

0	The command completed successfully.
nonzero	An error has occurred.

**Examples** **EXAMPLE 1** Adding Trust Between Two Clusters

The following geops commands add trust between the clusters `cluster-paris` and `cluster-newyork`.

1. Add cluster `cluster-newyork` to the trust on the cluster `cluster-paris`.

```
cluster-paris# /usr/cluster/bin/geops add-trust -c cluster-newyork
```

Keys are stored on every node of cluster `cluster-paris` with alias `cluster-newyork.certificatenumber`.

2. Add cluster `cluster-paris` to the trust on the cluster `cluster-newyork`.

```
cluster-newyork# /usr/cluster/bin/geops add-trust -c cluster-paris
```

3. Verify trust between the two clusters.

```
cluster-paris# /usr/cluster/bin/geops verify-trust -c cluster-newyork
```

```
cluster-newyork# /usr/cluster/bin/geops verify-trust -c cluster-paris
```

**EXAMPLE 2** Creating a Partnership

The following geops command creates the partnership `paris-newyork-ps`, on the cluster `cluster-paris`.

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

**EXAMPLE 3** Joining a Partnership

The following geops command joins `cluster-paris` and `cluster-newyork` in the `paris-newyork-ps` partnership.

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

**EXAMPLE 4** Modifying a Partnership

The following geops command modifies the notification email address for cluster-paris.

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

**EXAMPLE 5** Leaving a Partnership

The following geops command removes the cluster cluster-paris from the partnership paris-newyork-ps.

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

**EXAMPLE 6** Resynchronizing a Partnership

The following geops command resynchronizes the partnership paris-newyork-ps.

The geops update command overwrites the configuration on the local cluster with the configuration from the remote partner cluster.

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

**Attributes** See [attributes\(5\)](#) for descriptions of the following attributes.

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	ha-cluster/geo/framework
Interface Stability	Evolving

**See Also** [rbac\(5\)](#)

# Index

---

## A

activating protection groups, 24

## C

configuring and managing the heartbeat mechanism,  
geohb, 17

creating, configuring, and managing partnerships  
between clusters, geops, 44

creating or managing protection groups, geopg, 23

## D

disabling the Oracle Solaris Cluster Geographic Edition  
infrastructure, geoadm, 10

displaying the status of the Oracle Solaris Cluster  
Geographic Edition infrastructure, geoadm, 10

## E

enabling the Oracle Solaris Cluster Geographic Edition  
infrastructure, geoadm, 10

## G

geoadm, enabling, disabling, or displaying the status of  
the Oracle Solaris Cluster Geographic Edition  
infrastructure, 10

geohb, configuring and managing the heartbeat  
mechanism., 17

geopg, creating or managing protection groups., 23

geops, creating, configuring, and managing  
partnerships between clusters., 44

## L

local\_database\_name, 35

local\_db\_service\_name, 35, 36

local\_rac\_proxy\_svr\_rg\_name, 35

## O

Oracle Data Guard

properties of

local\_database\_name, 35

local\_db\_service\_name, 35

local\_rac\_proxy\_svr\_rg\_name, 35

remote\_database\_name, 36

remote\_db\_service\_name, 36

remote\_rac\_proxy\_svr\_rg\_name, 36

replication\_mode, 36

standby\_type, 37

sysdba\_password, 37

sysdba\_username, 37

## **P**

### properties

#### Oracle Data Guard

- local\_database\_name, 35
- local\_db\_service\_name, 35
- local\_rac\_proxy\_svr\_rg\_name, 35
- remote\_database\_name, 36
- remote\_db\_service\_name, 36
- remote\_rac\_proxy\_svr\_rg\_name, 36
- replication\_mode, 36
- standby\_type, 37
- sysdba\_password, 37
- sysdba\_username, 37

## **R**

- remote\_database\_name, 36
- remote\_rac\_proxy\_svr\_rg\_name, 36
- replication\_mode, 36

## **S**

- standby\_type, 37
- starting protection groups, 24
- starting resource groups, 24
- sysdba\_password, 37
- sysdba\_username, 37