

Oracle® Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard

Copyright © 2008, 2010, 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 software 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 RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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 which 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.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd.

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

Preface	7
1 Replicating Data With Oracle Data Guard Software	13
Replicating Data in an Oracle Data Guard Protection Group (Task Map)	14
Overview of Oracle Data Guard Data Replication	15
Oracle Data Guard Shadow Resource Groups	15
Oracle Data Guard Replication Resource Groups	16
Initially Configuring Oracle Data Guard Software	16
Oracle Data Guard Broker Configurations	17
▼ How to Set Up Your Primary Database	19
▼ How to Configure the Primary Database Listener and Naming Service	22
▼ How to Prepare Your Standby Database	24
▼ How to Configure the Standby Database Listener and Naming Service	27
▼ How to Start and Recover Your Standby Database	30
▼ How to Verify That Your Configuration Is Working Correctly	31
▼ How to Complete Configuring and Integrating Your Standby Database	31
▼ How to Create and Enable an Oracle Data Guard Broker Configuration	32
2 Administering Oracle Data Guard Protection Groups	35
Working With Oracle Data Guard Protection Groups	35
Overview of Administering Protection Groups	35
▼ How to Administer an Oracle Data Guard Protection Group (Example)	36
Creating, Modifying, Validating, and Deleting an Oracle Data Guard Protection Group	42
▼ How to Create and Configure an Oracle Data Guard Protection Group	42
▼ How to Modify an Oracle Data Guard Protection Group	44
▼ How to Validate an Oracle Data Guard Protection Group	45
How the Data Replication Layer Validates the Application Resource Groups and Data Replication Entities	46

▼ How to Delete an Oracle Data Guard Protection Group	47
Administering Oracle Data Guard Application Resource Groups	48
▼ How to Add an Application Resource Group to an Oracle Data Guard Protection Group	49
▼ How to Delete an Application Resource Group From an Oracle Data Guard Protection Group	51
Administering Oracle Data Guard Broker Configurations	52
▼ How to Add an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group	53
How the Data Replication Subsystem Verifies the Oracle Data Guard Broker Configuration	55
▼ How to Modify an Oracle Data Guard Broker Configuration	56
▼ How to Delete an Oracle Data Guard Broker Configuration From an Oracle Data Guard Protection Group	57
Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster	58
▼ How to Replicate the Oracle Data Guard Protection Group Configuration to a Partner Cluster	58
Activating and Deactivating a Protection Group	60
▼ How to Activate an Oracle Data Guard Protection Group	60
▼ How to Deactivate an Oracle Data Guard Protection Group	63
Resynchronizing an Oracle Data Guard Protection Group	65
▼ How to Resynchronize an Oracle Data Guard Protection Group	65
Checking the Runtime Status of Oracle Data Guard Data Replication	66
Displaying an Oracle Data Guard Runtime Status Overview	66
Displaying a Detailed Oracle Data Guard Runtime Status	67
3 Migrating Services That Use Oracle Data Guard Data Replication	71
Detecting Cluster Failure on a System That Uses Oracle Data Guard Data Replication	71
Detecting Primary Cluster Failure	71
Detecting Failure of the Standby Cluster	72
Migrating Services That Use Oracle Data Guard With a Switchover	72
▼ How to Switch Over an Oracle Data Guard Protection Group From the Primary to the Standby Cluster	73
Actions Performed by the Geographic Edition Software During a Switchover	74
Forcing a Takeover on Systems That Use Oracle Data Guard	75
▼ How to Force Immediate Takeover of Oracle Data Guard Services by a Standby Cluster ..	76
Actions Performed by the Geographic Edition Software During a Takeover	77

Recovering Oracle Data Guard Data After a Takeover	78
▼ How to Resynchronize and Revalidate the Protection Group Configuration	79
▼ How to Perform a Failback Switchover on a System That Uses Oracle Data Guard Replication	81
▼ How to Perform a Failback Takeover on a System That Uses Oracle Data Guard Replication	85
Recovering From an Oracle Data Guard Data Replication Error	88
▼ How to Recover From a Data Replication Error	88
A Geographic Edition Properties for Oracle Data Guard Broker Configurations	91
Oracle Data Guard Broker Configuration Properties	91
Index	95

Preface

The *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard* provides procedures for administering Oracle Data Guard data replication with Oracle Solaris Cluster Geographic Edition software on both SPARC and x86 based systems.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

Who Should Use This Book

This document is intended for system administrators, support personnel, and application developers who work with the Oracle Solaris Cluster Geographic Edition (Geographic Edition) product, Oracle RAC, and Oracle Data Guard software.

To understand the concepts that are described in this book, you need to be familiar with the Solaris Operating System (Solaris OS) and also have expertise with Oracle Solaris Cluster software and with the Oracle RAC software that is supported for use with Oracle Solaris Cluster software.

How This Book Is Organized

This guide contains the following chapters and appendix:

[Chapter 1, “Replicating Data With Oracle Data Guard Software,”](#) describes how to configure data replication with Oracle Data Guard software.

[Chapter 2, “Administering Oracle Data Guard Protection Groups,”](#) describes how to administer data replication with Oracle Data Guard software.

[Chapter 3, “Migrating Services That Use Oracle Data Guard Data Replication,”](#) describes how to migrate services for maintenance or in the event that your cluster fails.

[Appendix A, “Geographic Edition Properties for Oracle Data Guard Broker Configurations,”](#) describes the properties for Geographic Edition data replications that use Oracle Data Guard.

Related Documentation

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

Topic	Documentation
Overview	<i>Oracle Solaris Cluster Geographic Edition Overview</i>
	<i>Oracle Solaris Cluster Geographic Edition 3.3 Documentation Center</i>
Installation	<i>Oracle Solaris Cluster Geographic Edition Installation Guide</i>
Data Replication	<i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>
	<i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy and Universal Replicator</i>
	<i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard</i>
	<i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i>
System administration	<i>Oracle Solaris Cluster Geographic Edition System Administration Guide</i>

Topic	Documentation
Overview	<i>Oracle Solaris Cluster Overview</i>
	<i>Oracle Solaris Cluster 3.3 Documentation Center</i>
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Hardware installation and administration	<i>Oracle Solaris Cluster 3.3 Hardware Administration Manual</i>
	Individual hardware administration guides
Software installation	<i>Oracle Solaris Cluster Software Installation Guide</i>
Data service installation and administration	<i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i>
	Individual data service guides
Data service development	<i>Oracle Solaris Cluster Data Services Developer's Guide</i>
System administration	<i>Oracle Solaris Cluster System Administration Guide</i>
	<i>Oracle Solaris Cluster Quick Reference</i>
Software upgrade	<i>Oracle Solaris Cluster Upgrade Guide</i>

Topic	Documentation
Error messages	<i>Oracle Solaris Cluster Error Messages Guide</i>
Command and function references	<i>Oracle Solaris Cluster Reference Manual</i> <i>Oracle Solaris Cluster Data Services Reference Manual</i> <i>Oracle Solaris Cluster Quorum Server Reference Manual</i>

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your release of Oracle Solaris Cluster software at <http://wikis.sun.com/display/SunCluster/Home/>.

Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model and serial numbers of your systems
- The release number of the operating system (for example, the Solaris 10 11/06 OS)
- The release number of Oracle Solaris Cluster software (for example, 3.3)
- The contents of the `/var/adm/messages` file

Use the following commands to gather information about your systems for your service provider.

Command	Function
<code>prtconf -v</code>	Displays the size of system memory and reports information about peripheral devices.
<code>psrinfo -v</code>	Displays information about processors.
<code>showrev -p</code>	Reports which patches are installed.
SPARC: <code>prtdiag -v</code>	Displays system diagnostic information.
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Oracle Solaris Cluster release and package version information.

Documentation, Support, and Training

See the following web sites for additional resources:

- [Documentation \(http://docs.sun.com\)](http://docs.sun.com)
- [Support \(http://www.oracle.com/us/support/systems/index.html\)](http://www.oracle.com/us/support/systems/index.html)
- [Training \(http://education.oracle.com\)](http://education.oracle.com) – Click the Sun link in the left navigation bar.

Oracle Welcomes Your Comments

Oracle welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions for improvement, go to <http://docs.sun.com> and click Feedback. Indicate the title and part number of the documentation along with the chapter, section, and page number, if available. Please let us know if you want a reply.

[Oracle Technology Network \(http://www.oracle.com/technetwork/index.html\)](http://www.oracle.com/technetwork/index.html) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the [Discussion Forums \(http://forums.oracle.com\)](http://forums.oracle.com).
- Get hands-on step-by-step tutorials with [Oracle By Example \(http://www.oracle.com/technology/obe/start/index.html\)](http://www.oracle.com/technology/obe/start/index.html).
- Download [Sample Code \(http://www.oracle.com/technology/sample_code/index.html\)](http://www.oracle.com/technology/sample_code/index.html).

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename.</code>

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Replicating Data With Oracle Data Guard Software

This chapter describes how to configure data replication with Oracle Data Guard software in an Oracle Solaris Cluster Geographic Edition (Geographic Edition) environment.

This chapter covers the following topics:

- “Replicating Data in an Oracle Data Guard Protection Group (Task Map)” on page 14
- “Overview of Oracle Data Guard Data Replication” on page 15
- “Initially Configuring Oracle Data Guard Software” on page 16

This release of Geographic Edition supports the Oracle Data Guard Physical standby type. The Logical standby type is not supported.

Geographic Edition software supports the use of Oracle Data Guard for data replication when used with Oracle Real Application Clusters (Oracle RAC) software. Before you can replicate data with Oracle Data Guard, you must be familiar with the Oracle Data Guard documentation. For information about installing and configuring the Oracle Data Guard software and its latest patches, see the Oracle Data Guard documentation.

Note – During data replication, data from a primary cluster is copied to a backup, or standby cluster. The standby cluster can be located at a site that is geographically separated from the primary cluster. The distance between the primary and standby clusters depends on the distance that your data replication product supports.

The example procedures in this chapter show how to configure Oracle Data Guard to replicate data between a primary and a standby database.

Replicating Data in an Oracle Data Guard Protection Group (Task Map)

The following table summarizes the steps for configuring Oracle Data Guard data replication in a protection group.

TABLE 1-1 Administration Tasks for Oracle Data Guard Data Replication

Task	Description
Perform an initial configuration of the Oracle Data Guard software.	See “Initially Configuring Oracle Data Guard Software” on page 16.
Create a protection group that is configured for Oracle Data Guard data replication.	See “How to Create and Configure an Oracle Data Guard Protection Group” on page 42.
Add a configuration that is controlled by Oracle Data Guard.	See “How to Add an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group” on page 53.
Add an application resource group to the protection group.	See “How to Add an Application Resource Group to an Oracle Data Guard Protection Group” on page 49.
Replicate the protection group configuration to a standby cluster.	See “How to Replicate the Oracle Data Guard Protection Group Configuration to a Partner Cluster” on page 58.
Activate the protection group.	See “How to Activate an Oracle Data Guard Protection Group” on page 60.
Check the runtime status of replication.	See “Checking the Runtime Status of Oracle Data Guard Data Replication” on page 66.
Detect failure.	See “Detecting Cluster Failure on a System That Uses Oracle Data Guard Data Replication” on page 71.
Migrate services by using a switchover.	See “Migrating Services That Use Oracle Data Guard With a Switchover” on page 72.
Migrate services by using a takeover.	See “Forcing a Takeover on Systems That Use Oracle Data Guard” on page 75.
Recover data after forcing a takeover.	See “Recovering Oracle Data Guard Data After a Takeover” on page 78.

Overview of Oracle Data Guard Data Replication

This section provides an overview of the integration of Oracle Data Guard with Geographic Edition and highlights the differences between support for Oracle Data Guard and other data replication products, such as Sun StorageTek Availability Suite software, Hitachi TrueCopy or Universal Replicator, and EMC SRDF.

Oracle Data Guard Shadow Resource Groups

You can add an Oracle Data Guard Broker configuration that is controlled by the Oracle Data Guard software to a protection group. The Geographic Edition software creates a shadow RAC server proxy resource group for each Oracle Data Guard Broker configuration. The name of a shadow resource group conforms to the following format:

ODGconfigurationname-rac-proxy-svr-shadow-rg

For example, an Oracle Data Guard Broker configuration named `sales` that is controlled by the Oracle Data Guard software has a shadow RAC server proxy resource group named `sales-rac-proxy-svr-shadow-rg`. If, however, the configuration name contains one or more periods (`.`), the periods are converted to underscore characters (`_`) to construct the resource group name. Consequently, the configuration name `mysales.com` has a shadow resource group named `mysales_com-rac-proxy-svr-shadow-rg`.

The shadow RAC server proxy resource group “shadows” the real Oracle RAC server proxy resource group that you created to manage and monitor the Oracle RAC databases that are under the control of Oracle Solaris Cluster software.

Each shadow resource group contains a single resource: a `SUNW.gds` resource whose probe script reflects the status of the Oracle RAC server proxy resource group. The name of this resource conforms to the following format:

ODGconfigurationname-rac-proxy-svr-shadow-rs

For more information about Oracle RAC server proxy resource groups, see [Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide](#).

A shadow RAC server proxy resource group is required because, unlike other Geographic Edition replication products, the Oracle Data Guard software is an integral part of the Oracle RAC software. Oracle Data Guard requires the Oracle RAC software to be running and the databases started to replicate its data.

Consequently, putting the real Oracle RAC server proxy resource group under Geographic Edition control would result in the Oracle RAC database's being shut down on the standby cluster. In contrast, the shadow RAC server proxy resource group can be placed under the control of Geographic Edition. You can do so without disrupting the data replication process while still allowing the configuration to conform to the usual Geographic Edition structure for managing application resource groups.

The state of the shadow RAC server proxy resource group indicates whether the database that is monitored and controlled by the Oracle RAC server proxy resource group is the primary or the standby cluster. In other words, this state indicates whether the database is online on the primary cluster and unmanaged on the standby cluster. Furthermore, the status of the shadow RAC server proxy resource reflects both the status of the Oracle RAC server proxy resource and whether the database is the primary or the standby.

Oracle Data Guard Replication Resource Groups

When an Oracle Data Guard Broker configuration that is controlling the Oracle Data Guard software is added to a protection group, the Geographic Edition software creates a special replication resource for the specific Oracle Data Guard Broker configuration in the replication resource group. By monitoring these replication resource groups, the Geographic Edition software is able to monitor the overall status of replication. One replication resource group with one replication resource for each Oracle Data Guard Broker configuration is created for each protection group.

The name of the replication resource group conforms to the following format:

ODGProtectiongroupName-odg-rep-rg.

The replication resource in the replication resource group monitors the replication status of the Oracle Data Guard Broker configuration on the local cluster, which is reported by the Oracle Data Guard Broker software.

The name of the replication resource conforms to the following format:

ODGBrokerConfigurationName-odg-rep-rs.

Note – In Oracle Data Guard, a data replication resource is enabled when the protection group is activated in the cluster. Consequently, in Oracle Data Guard, in a cluster in which the protection group is deactivated, the data replication status appears as unknown.

Initially Configuring Oracle Data Guard Software

This section describes the initial steps that you need to perform to configure Oracle Data Guard replication in the Geographic Edition product.

Note – The steps in this document that describe how to use Oracle tools and commands, such as `dgmgrl`, are intended for illustration only. Consult your Oracle documentation to determine the detailed procedures that you need to follow to satisfy the particular needs of your environment.

The example protection group, `sales-pg`, in this section has been configured in a partnership that consists of two (partner) clusters, `cluster-paris` and `cluster-newyork`. An Oracle RAC database, which is managed and monitored by an individual Oracle RAC server proxy resource group on each cluster, is shadowed by the `mysales_com-rac-proxy-svr-shadow-rg` shadow RAC server proxy resource group. The application data is contained in the `sales` database and replicated by Oracle Data Guard as part of the `mysales.com` Oracle Data Guard Broker configuration.

The shadow RAC server proxy resource group, `mysales_com-rac-proxy-svr-shadow-rg`, and the Oracle Data Guard Broker configuration, `mysales.com`, are present on both the `cluster-paris` and the `cluster-newyork` clusters. However, the names for the Oracle RAC server proxy resource group they shadow might be different on both the `cluster-paris` and the `cluster-newyork` clusters. The `sales-pg` protection group protects the application data by managing the replication of data between the `cluster-paris` and the `cluster-newyork` clusters.

This section provides the following information:

- [“Oracle Data Guard Broker Configurations” on page 17](#)
- [“How to Set Up Your Primary Database” on page 19](#)
- [“How to Configure the Primary Database Listener and Naming Service” on page 22](#)
- [“How to Prepare Your Standby Database” on page 24](#)
- [“How to Configure the Standby Database Listener and Naming Service” on page 27](#)
- [“How to Start and Recover Your Standby Database” on page 30](#)
- [“How to Verify That Your Configuration Is Working Correctly” on page 31](#)
- [“How to Complete Configuring and Integrating Your Standby Database” on page 31](#)
- [“How to Create and Enable an Oracle Data Guard Broker Configuration” on page 32](#)

Oracle Data Guard Broker Configurations

To define Oracle Data Guard Broker configurations, you need to determine the following information:

- **The name of the Oracle Data Guard Broker configuration**, such as `mysales.com`, being replicated between the `cluster-paris` and `cluster-newyork` clusters.
- **The unique database names that are taking part in the replication**, such as `sales` on the `cluster-paris` cluster, and `salesdr` on the `cluster-newyork` cluster.

- **The Oracle service names for these databases**, such as `sales - svc` on the `cluster-paris` cluster and `salesdr - svc` on the `tcluster-newyork` cluster. These names are held in the `tnsnames.ora` files in the `${ORACLE_HOME}/network/admin` directory of the nodes that are hosting the Oracle database that is being replicated, or in the Oracle naming service directory.
- **The database standby type for the Oracle Data Guard Broker configuration**, which you set to either `physical` or `logical`.
- **The replication mode for the Oracle Data Guard Broker configuration**, which you set to `MaxPerformance`, `MaxAvailability`, or `MaxProtection`.

After you configure Oracle Data Guard between a pair of primary and standby databases, you create an Oracle Data Guard Broker configuration by using the `${ORACLE_HOME}/bin/dgmgctl` command to define the properties of the named replication. You can use this command to set and to retrieve the previously listed Oracle Data Guard Broker properties.

You also need to determine the names of the Oracle RAC server proxy resource groups that manage the Oracle RAC databases on each cluster. You configure these names by using the Data Service configuration wizard through the `clsetup` command, or by following the instructions in [Appendix D, “Command-Line Alternatives,” in *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide*](#).

Of the Oracle Data Guard Broker configuration properties that are listed in the following table, you can change only the Protection Mode property with the Geographic Edition software. You cannot use the Geographic Edition software to modify other Oracle Data Guard Broker properties in the configuration, such as the `DelayMins`, `MaxFailure`, `MaxConnections`, and `NetTimeout` properties. You need to adjust these properties manually by using the Oracle Data Guard Broker command, or by modifying the appropriate database parameters that are held in the `spfile` server parameter file or the `init${SID}.ora` file through SQL*Plus.

Property	Allowed Values	Description
Protection Mode	MaxPerformance, MaxAvailability or MaxProtection	The data replication mode that is being used by Oracle, ranging from asynchronous (MaxPerformance) to synchronous (MaxProtection)
Standby type	physical or logical	The type of replication that is being performed, either Redo Apply (physical) or SQL Apply (logical) held as part of the primary database definition
Configuration name		The name for the Oracle Data Guard Broker configuration, which consists of a primary and a standby database
Primary database		The name of the primary database, its net service name, and its standby type

Property	Allowed Values	Description
Secondary database		The name of the standby database and its net service name

The Geographic Edition software modifies the Oracle Data Guard Broker configuration role changes during switchover and takeover operations.

For more information about the Oracle Data Guard Broker configuration, refer to the [Oracle Data Guard Broker documentation \(http://download.oracle.com/docs/cd/B19306_01/server.102/b14230/toc.htm\)](http://download.oracle.com/docs/cd/B19306_01/server.102/b14230/toc.htm).

▼ How to Set Up Your Primary Database

In the following steps, the primary cluster is called `cluster-paris` (nodes `phys-paris-1` and `phys-paris-2`), and the standby cluster is called `cluster-newyork` (`phys-newyork-1` and `phys-newyork-2`). The suffix `-crs` is appended to the Oracle Clusterware virtual IP host names.

The primary database on `cluster-paris` is called `sales` and has instances `sales1` and `sales2`. The standby database on `cluster-newyork` is called `salesdr` and has instances `salesdr1` and `salesdr2`. The suffix `-svc` is appended to each net naming service name for each of the databases and individual instances, for example, `sales-svc` or `sales1-svc`.

Before You Begin Ensure that you have edited your Oracle user `.profile` or `.cshrc` file to set the correct Oracle SID, `ORACLE_HOME`, and `PATH` environment variables for the local Oracle RAC database instance. Unless otherwise stated, you only need to run the commands from a node in the primary cluster that hosts a protected database instance.

- 1 **Verify that you can resolve the Oracle virtual IP addresses that are used by Oracle Clusterware on all primary and standby nodes.**

```
phys-paris-1# getent hosts phys-paris-1-crs
10.11.112.41    phys-paris-1-crs
...
```

- 2 **Create a database on the primary cluster.**

Use either the Oracle Database Configuration Assistant (dbca) or the SQL*Plus utility.

- 3 **Verify that an Oracle password file exists for the primary database.**

```
oracle (phys-paris-1)$ cd ${ORACLE_HOME}/dbs
oracle (phys-paris-1)$ ls -l orapwsales1
lrwxrwxrwx 1 oracle oinstall 25 November 2 02:06 orapwsales1
-> /oradata/SALES/orapwsales
```

Oracle Data Guard needs a consistent Oracle password file on all participating nodes in the primary and standby clusters.

If a password file does not exist, create one as follows:

```
oracle (phys-paris-1)$ orapwd file=${ORACLE_HOME}/dbs/orapwsales1 \
    password=sysdba_password
```

You can then move this file to a location on shared storage and create a symbolic link to that file from each node. Change the file name to reflect the local SID on each node. Later, you will copy this file to the standby cluster (cluster-newyork).

4 Ensure that the database is in logging mode by using the sqlplus command.

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> alter database force logging;
Database altered.
```

5 Configure the Oracle Data Guard Broker configuration file locations.

Run the sqlplus command as follows, substituting the two file names with ones that suit your configuration. Ensure that these files are located on shared storage that is visible to all cluster-paris nodes.

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> alter system set dg_broker_config_file1='/oradata/SALES/dr1sales.dat'
    2 scope=both sid='*';
System altered.
SQL> alter system set dg_broker_config_file2='/oradata/SALES/dr2sales.dat'
    2 scope=both sid='*';
System altered.
```

6 Shut down all database instances.

7 On the primary database, mount a single database instance and enable the Oracle database flashback capability.

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> startup mount;
ORACLE instance started.
```

```
Total System Global Area  532676608 bytes
Fixed Size                  2031416 bytes
Variable Size               276824264 bytes
Database Buffers           247463936 bytes
Redo Buffers                6356992 bytes
Database mounted.
System altered.
SQL> alter database archivelog;
Database altered.
SQL> alter database flashback on;
Database altered.
SQL> alter database open;
Database altered.
```

8 Restart the other database instances.

9 Create database standby redo logs.

Depending on your configuration, you might need to add a number of standby redo logs. The name, number, and size of these logs depend on a number of factors, including whether you use the Optimal Flexible Architecture (OFA), how many online redo log files you have, and the size of those log files.

The following example shows how to configure a single 50-Mbyte standby redo log file, where the OFA naming scheme is being used. A default, two-node Oracle RAC database normally requires that you add six log files.

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> alter database add standby logfile size 50m;
Database altered.
```

10 Configure the Oracle log archiving destinations.

Depending on your configuration, you might need to alter or add one or more of the Oracle log archive destination parameters. These parameters have a number of tunable properties. Consult the Oracle documentation for details.

The following example shows two log archive destinations being set, one for the local cluster and one for the standby cluster, where OFA naming is used.

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> alter system set log_archive_dest_1='location=use_db_recovery_file_dest
  2 arch mandatory valid_for=(all_logfiles,all_roles)
  3 db_unique_name=sales' scope=both sid='*';
System altered.
```

```
SQL> alter system set log_archive_dest_2='service=salesdr-svc
  2 lgwr sync affirm valid_for=(online_logfiles,primary_role)
  3 db_unique_name=salesdr' scope=both sid='*';
System altered.
```

```
SQL> alter system set log_archive_dest_10='location=use_db_recovery_file_dest'
  2 scope=both sid='*';
System altered.
```

```
SQL> alter system set standby_file_management='AUTO' scope=both sid='*';
System altered.
```

11 Configure the Fetch Archive Log (FAL) parameters.

For the database to know where to get missing archive redo logs on the server and where to send them on the client, you need to set the FAL system properties. These properties use the net service names of the source and destination databases. You run the following `sqlplus` command to set the parameters to the correct values for your configuration.

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> alter system set fal_server='salesdr-svc' scope=both sid='*';
System altered.
```

```
SQL> alter system set fal_client='sales-svc' scope=both sid='*';
System altered.
```

▼ How to Configure the Primary Database Listener and Naming Service

1 Create a static listener for Oracle Data Guard.

Note – Perform this step on all `cluster-paris` nodes.

Oracle Data Guard requires that you configure a static listener. The following example uses `${ORACLE_HOME}=/oracle/oracle/product/10.2.0/db_1` and shows where to add the entry for the static listener in the `${ORACLE_HOME}/network/admin/listener.ora` file. The `SID_LIST_LISTENER_PHYS-PARIS-1` and `(SID_NAME = sales1)` lines vary from node to node, while the `(GLOBAL_DBNAME=sales_DGMGRL)` differs on `cluster-newyork`. Later, you will add these entries on the `cluster-newyork` nodes.

```
oracle (phys-paris-1)$ cat ${ORACLE_HOME}/network/admin/listener.ora
SID_LIST_LISTENER_PHYS-PARIS-1 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = PLSExtProc)
      (ORACLE_HOME = /oracle/oracle/product/10.2.0/db_1)
      (PROGRAM = extproc)
    )
    (SID_DESC =
      (SID_NAME = sales1)
      (GLOBAL_DBNAME=sales_DGMGRL)
      (ORACLE_HOME = /oracle/oracle/product/10.2.0/db_1)
    )
  )
oracle (phys-paris-1)$
```

2 Restart the listener.

To enable the static entries, restart the Oracle listener processes on each of the nodes on `cluster-paris`.

```
oracle (phys-paris-1)$ lsnrctl stop LISTENER_PHYS_PHYS-PARIS-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:04:56

Copyright (c) 1991, 2006, Oracle. All rights reserved.

Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))
The command completed successfully
oracle$ lsnrctl start LISTENER_PHYS_PHYS-PARIS-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:05:04
...Services Summary...
Service "PLSExtProc" has 1 instance(s).
  Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
Service "sales_DGMGRL" has 1 instance(s).
  Instance "sales1", status UNKNOWN, has 1 handler(s) for this service...
The command completed successfully
```

Wait while databases register with listener

```

oracle (phys-paris-1)$ lsnrctl status LISTENER_PHYS_PHYS-PARIS-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:04:56

Copyright (c) 1991, 2006, Oracle. All rights reserved.
...
Services Summary...
Service "PLSExtProc" has 1 instance(s).
  Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
Service "sales" has 2 instance(s).
  Instance "sales1", status READY, has 2 handler(s) for this service...
  Instance "sales2", status READY, has 1 handler(s) for this service...
Service "salesXDB" has 2 instance(s).
  Instance "sales1", status READY, has 1 handler(s) for this service...
  Instance "sales2", status READY, has 1 handler(s) for this service...
Service "sales_DGB" has 2 instance(s).
  Instance "sales1", status READY, has 2 handler(s) for this service...
  Instance "sales2", status READY, has 1 handler(s) for this service...
Service "sales_DGMGRL" has 1 instance(s).
  Instance "sales1", status UNKNOWN, has 1 handler(s) for this service...
Service "sales_XPT" has 2 instance(s).
  Instance "sales1", status READY, has 2 handler(s) for this service...
  Instance "sales2", status READY, has 1 handler(s) for this service...
The command completed successfully

```

3 Verify the net service naming entries for all database instances.

Ensure that the naming service method that you are using, either `tnsnames.ora` or the directory service, has entries defined for all the Oracle database instances in both clusters.

The following example shows the type of entries that you include for the `cluster-paris` cluster only. Entries for the `cluster-newyork` cluster are added in [“How to Configure the Standby Database Listener and Naming Service” on page 27](#). Also, add entries for the standby (`salesdr`) database instances that you create later when you modify the `pf` file parameter file. In the example, the `sales` database dynamically registers a service name of `sales` with the listeners (see the database `service_names` initialization parameter).

```

oracle (phys-paris-1)$ cat ${ORACLE_HOME}/network/admin/tnsnames.ora
SALES1-SVC =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = phys-paris-1-cr1)(PORT = 1521)
        (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = sales)
      (INSTANCE_NAME = sales1)
    )
  )
)
SALES2-SVC =
  (DESCRIPTION =

```

```

        (ADDRESS_LIST =
          (ADDRESS = (PROTOCOL = TCP)(HOST = phys-paris-2-crs)(PORT = 1521)
            (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
        )
        (CONNECT_DATA =
          (SERVER = DEDICATED)
          (SERVICE_NAME = sales)
          (INSTANCE_NAME = sales2)
        )
      )
)

SALES-SVC =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = phys-paris-1-crs)(PORT = 1521)
        (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
      (ADDRESS = (PROTOCOL = TCP)(HOST = phys-paris-2-crs)(PORT = 1521)
        (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
      (LOAD_BALANCE = yes)
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = sales)
    )
  )
)

LISTENERS_SALES =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = phys-paris-1-crs)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = phys-paris-2-crs)(PORT = 1521))
  )
)

```

▼ How to Prepare Your Standby Database

1 Create a backup of the primary database.

The following example shows how to use the Oracle Recovery Manager (RMAN) utility to create a copy of the primary database that you can restore on the standby cluster-newyork cluster. The example also shows how to avoid performing a separate step to create a control file for the standby database. For more information about the options for completing this step, see your Oracle documentation.

```

oracle (phys-paris-1)$ rman
RMAN> connect target sys/DBA_password@sales-svc;
RMAN> connect auxiliary /;
RMAN> backup device type disk tag 'mybkup' database include current
2> controlfile for standby;
RMAN> backup device type disk tag 'mybkup' archivelog all not backed up;

```

2 Copy the backup files to the standby system.

Create the appropriate directory hierarchies on the `cluster-newyork` cluster and copy the database backup to this cluster. The actual locations that you specify for the files that are shown in the example depend on the specific choices that you made when you configured the database.

```
oracle (phys-newyork-1)$ mkdir -p $ORACLE_BASE/admin/salesdr
oracle (phys-newyork-1)$ cd $ORACLE_BASE/admin/salesdr
oracle (phys-newyork-1)$ mkdir adump bdump cdump dpdump hdump pfile udump
    Make the directory for the database backup
oracle (phys-newyork-1)$ mkdir -p /oradata/flash_recovery_area/SALES/backupset/date
    Copy over the files
oracle (phys-newyork-1)$ cd /oradata/flash_recovery_area/SALES/backupset/date
oracle (phys-newyork-1)$ scp oracle@phys-paris-1:'pwd'/* .
    Make the base directory for new database files
oracle (phys-newyork-1)$ mkdir -p /oradata/SALESDR
```

3 Create a pfile parameter file.

Create a suitable server initialization file for the standby (`salesdr`) database. The easiest way to create this file is to copy the parameters for the primary database and modify them. The following example shows how to create a pfile parameter file:

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> CREATE PFILE='/tmp/initpfile_for_salesdr.ora' FROM SPFILE;
File created.
SQL> quit
```

4 Modify the pfile parameter file.

Change all entries that are particular to the primary cluster to entries that are suitable for the standby cluster, as shown in the following example. Modify entries that are prefixed by an Oracle SID, that is, `sales1` or `sales2`, to use standby database instance SID names, that is, `salesdr1` and `salesdr2`. Depending on your configuration, you might need to make additional changes.

Note – Do not change the `db_name` parameter, as it must remain `sales` on both clusters.

You created these directories previously

```
*.audit_file_dest='/oracle/oracle/product/10.2.0/db_1/admin/salesdr/adump'
*.background_dump_dest='/oracle/oracle/product/10.2.0/db_1/admin/salesdr/bdump'
*.user_dump_dest='/oracle/oracle/product/10.2.0/db_1/admin/salesdr/udump'
*.core_dump_dest='/oracle/oracle/product/10.2.0/db_1/admin/salesdr/cdump'
```

Remove the following entry

```
*.control_files='...list primary control files...'
```

Add this entry

```
*.db_unique_name='salesdr'
```

```
*.dg_broker_config_file1='/oradata/SALESDR/dr1salesdr.dat'
*.dg_broker_config_file2='/oradata/SALESDR/dr2salesdr.dat'

*.dispatchers='(PROTOCOL=TCP) (SERVICE=salesdrXDB)'
```

Switch the client and server entries around, as shown in the following entries

```
*.fal_client='salesdr-svc'
*.fal_server='sales-svc'

*.remote_listener='LISTENERS_SALESDR'
```

Switch the log archive destinations

```
*.log_archive_dest_1='location=use_db_recovery_file_dest arch
mandatory valid_for=(all_logfiles,all_roles) db_unique_name=salesdr'
*.log_archive_dest_2='service=sales-svc lgwr sync affirm
valid_for=(online_logfiles,primary_role) db_unique_name=sales'
```

- 5 Copy the pfile parameter file to the standby system.
- 6 Start the standby database and convert the pfile parameter file to an spfile server parameter file.

- a. As the Oracle user, log in to one of the cluster-newyork nodes and convert the pfile parameter file to an spfile server parameter file.

```
oracle (phys-newyork-1)$ ORACLE_SID=salesdr1 export ORACLE_SID
oracle (phys-newyork-1)$ sqlplus '/ as sysdba'
SQL> startup nomount pfile='/tmp/initpfile_for_salesdr.ora';
SQL> create spfile='/oradata/SALESDR/spfilesalesdr.ora'
2> from pfile='/tmp/initpfile_for_salesdr.ora';
SQL> shutdown
```

- b. Create a \${ORACLE_HOME}/dbs/initsalesdr1.ora file on all cluster-newyork nodes and, in that file, insert the following entry:

```
oracle (phys-newyork-1) cat ${ORACLE_HOME}/dbs/initsalesdr1.ora
SPFILE='/oradata/SALESDR/spfilesalesdr.ora'
```

- c. Restart the database, on one node only, to prepare for restoring the backed-up primary database.

```
oracle (phys-newyork-1) sqlplus '/ as sysdba'
You are now starting from the spfile
```

```
SQL> startup nomount
ORACLE instance started.
```

```
Total System Global Area 532676608 bytes
Fixed Size                  2031416 bytes
Variable Size               289407176 bytes
Database Buffers            234881024 bytes
Redo Buffers                 6356992 bytes
```

- 7 **Copy the Oracle password file for the primary database for use by the standby database.**
 - a. **Copy the Oracle password file that you created on the cluster-paris cluster.**
Place the file on shared storage on the cluster-newyork cluster.
 - b. **Create links to this file from each of the cluster-newyork nodes.**
Again change the name of the symbolic link to reflect the Oracle SID on the local standby node.

▼ How to Configure the Standby Database Listener and Naming Service

1 Create a static listener for Oracle Data Guard.

Note – Perform this step on all cluster-newyork nodes.

Oracle Data Guard requires that you configure a static listener.

The following example uses `${ORACLE_HOME}=/oracle/oracle/product/10.2.0/db_1` and shows where to add the entry for the static listener in the `${ORACLE_HOME}/network/admin/listener.ora` file. The `SID_LIST_LISTENER_PHYS-NEWYORK-1` and `(SID_NAME = salesdr1)` lines vary from node to node, while the `(GLOBAL_DBNAME=salesdr_DGMGRL)` differs on cluster-paris.

```
oracle (phys-newyork-1)$ cat ${ORACLE_HOME}/network/admin/listener.ora
SID_LIST_LISTENER_PHYS-NEWYORK-1 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = PLSExtProc)
      (ORACLE_HOME = /oracle/oracle/product/10.2.0/db_1)
      (PROGRAM = extproc)
    )
    (SID_DESC =
      (SID_NAME = salesdr1)
      (GLOBAL_DBNAME=salesdr_DGMGRL)
      (ORACLE_HOME = /oracle/oracle/product/10.2.0/db_1)
    )
  )
oracle (phys-newyork-1)$
```

2 Restart the listener.

To enable the static entries, restart the Oracle listener processes on each of the nodes on cluster-newyork.

```
oracle (phys-newyork-1)$ lsnrctl stop LISTENER_PHYS_PHYS-NEWYORK-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:04:56
```

```

Copyright (c) 1991, 2006, Oracle. All rights reserved.

Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))
The command completed successfully
oracle$ lsnrctl start LISTENER_PHYS_PHYS-NEWYORK-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:05:04

Copyright (c) 1991, 2006, Oracle. All rights reserved.

Starting /oracle/oracle/product/10.2.0/db_1/bin/tnslsnr: please wait...

TNSLSNR for Solaris: Version 10.2.0.4.0 - Production

Services Summary...
Service "PLSExtProc" has 1 instance(s).
  Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
Service "salesdr_DGMGRL" has 1 instance(s).
  Instance "salesdr1", status UNKNOWN, has 1 handler(s) for this service...
The command completed successfully

```

Wait while databases register with listener

```

oracle (phys-newyork-1)$ lsnrctl status LISTENER_PHYS_PHYS-NEWYORK-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:04:56

Copyright (c) 1991, 2006, Oracle. All rights reserved....
Services Summary...
Service "PLSExtProc" has 1 instance(s).
  Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
Service "salesdr" has 2 instance(s).
  Instance "salesdr1", status READY, has 2 handler(s) for this service...
  Instance "salesdr2", status READY, has 1 handler(s) for this service...
Service "salesdrXDB" has 2 instance(s).
  Instance "salesdr1", status READY, has 1 handler(s) for this service...
  Instance "salesdr2", status READY, has 1 handler(s) for this service...
Service "salesdr_DGB" has 2 instance(s).
  Instance "salesdr1", status READY, has 2 handler(s) for this service...
  Instance "salesdr2", status READY, has 1 handler(s) for this service...
Service "salesdr_DGMGRL" has 1 instance(s).
  Instance "salesdr1", status UNKNOWN, has 1 handler(s) for this service...
Service "salesdr_XPT" has 2 instance(s).
  Instance "salesdr1", status READY, has 2 handler(s) for this service...
  Instance "salesdr2", status READY, has 1 handler(s) for this service...
The command completed successfully

```

3 Verify the net service naming entries for all database instances.

Ensure that the naming service method that you are using, either `tnsnames.ora` or the directory service, has entries defined for all the Oracle database instances in both clusters.

The following example shows the type of entries that you include for the `cluster-newyork` cluster only. Entries for the `cluster-paris` cluster are added in [“How to Configure the Primary](#)

[Database Listener and Naming Service” on page 22](#). In the example, the salesdr database dynamically registers a service name of salesdr with the listeners (see the database service_names initialization parameter).

```
oracle (phys-newyork-1)$ cat ${ORACLE_HOME}/network/admin/tnsnames.ora
SALESDR1-SVC =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = phys-newyork-1-crs)(PORT = 1521)
        (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = salesdr)
      (INSTANCE_NAME = salesdr1)
    )
  )

SALESDR2-SVC =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = phys-newyork-2-crs)(PORT = 1521)
        (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = salesdr)
      (INSTANCE_NAME = salesdr2)
    )
  )

SALESDR-SVC =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = phys-newyork-1-crs)(PORT = 1521)
        (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
      (ADDRESS = (PROTOCOL = TCP)(HOST = phys-newyork-2-crs)(PORT = 1521)
        (SEND_BUF_SIZE = 65535)(RECV_BUF_SIZE = 65535))
      (LOAD_BALANCE = yes)
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = salesdr)
    )
  )

LISTENERS_SALESDR =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = phys-newyork-1-crs)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = phys-newyork-2-crs)(PORT = 1521))
  )
)
```

4 Verify that the standby listener `listener.ora` and `tnsnames.ora` files have the correct entries, and restart the listener process.

Ensure that these files include the static Oracle Data Guard listener entry and the naming service entries for the primary and standby cluster database service. If you are not using the Oracle directory naming service lookup, you need to include the entries in `tnsnames.ora`.

```
oracle (phys-newyork-1)$ lsnrctl stop LISTENER_PHYS-NEWYORK-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:04:56
```

Copyright (c) 1991, 2006, Oracle. All rights reserved.

```
Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))
The command completed successfully
```

```
oracle$ lsnrctl start LISTENER_PHYS-NEWYORK-1
LSNRCTL for Solaris: Version 10.2.0.4.0 - Production on 29-OCT-2008 02:05:04
```

Copyright (c) 1991, 2006, Oracle. All rights reserved.

```
Starting /oracle/oracle/product/10.2.0/db_1/bin/tnslsnr: please wait...
```

```
TNSLSNR for Solaris: Version 10.2.0.4.0 - Production
```

```
...
```

```
Services Summary...
```

```
Service "PLSExtProc" has 1 instance(s).
```

```
Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
```

```
Service "salesdr_DGMGR" has 1 instance(s).
```

```
Instance "salesdr1", status UNKNOWN, has 1 handler(s) for this service...
```

```
The command completed successfully
```

▼ How to Start and Recover Your Standby Database

1 Restore the database backup.

Continuing to work on the `cluster-newyork` cluster, you can now restore the data from the backup of the primary database to the standby database.

The following example shows how to use the Oracle Recovery Manager (RMAN) utility.

```
oracle (phys-newyork-1) rman
RMAN> connect target sys/oracle@sales-svc;
RMAN> connect auxiliary /;
RMAN> duplicate target database for standby nofilenamecheck;
...
```

2 Add standby redo logs to the standby database.

The exact requirements that you must meet depend on your configuration. The steps you follow are identical to those that you followed for the primary cluster.

3 Enable flashback on the standby database.

```
oracle (phys-newyork-1)$ sqlplus '/ as sysdba'
SQL> alter database flashback on;
Database altered.
```

```
SQL> shutdown immediate;
SQL> startup mount;
ORACLE instance started.
...
```

4 Recover the standby database.

```
oracle (phys-newyork-1) sqlplus '/ as sysdba'
SQL> alter database recover managed standby database using current logfile disconnect;
```

▼ How to Verify That Your Configuration Is Working Correctly

1 Verify that the log file transmission is working.

When the SQL> prompt is displayed, log in to one of the database instances on the cluster-paris cluster and perform a couple log switches.

```
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> alter system switch logfile;
SQL> alter system switch logfile;
```

2 Check the \${ORACLE_HOME}/admin/sales/bdump/alert_sales1.log for any problems that might have prevented the logs from being archived.

If there are errors, correct them. This process might take time. You can check that the network connectivity is correct by using the following command:

```
oracle (phys-paris-1)$ tnsping salesdr-svc
oracle (phys-newyork-1)$ tnsping sales-svc
```

▼ How to Complete Configuring and Integrating Your Standby Database

1 Register the new database and instances with Oracle Clusterware.

Place the standby database under Oracle Clusterware control and configure it to mount when Oracle Clusterware starts.

```
oracle (phys-newyork-1)$ srvctl add database -d salesdr \
-r PHYSICAL_STANDBY -o $ORACLE_HOME -s mount;
oracle (phys-newyork-1)$ srvctl add instance -d salesdr \
-i salesdr1 -n $phys-newyork-1;
oracle (phys-newyork-1)$ srvctl add instance -d salesdr \
-i salesdr2 -n $phys-newyork-2;
```

2 Configure the Oracle Solaris Cluster Oracle RAC manageability resources.

Integrate the standby database with Oracle Solaris Cluster. You can use either the Data Service configuration wizard that is available through the clsetup utility or the browser-based Oracle

Solaris Cluster Manager. By integrating the standby database, you allow the standby to be managed as the primary database is, should a failover or takeover be necessary.

Note – The resource and resource group that you create are used by the Geographic Edition Oracle Data Guard integration.

3 Enable Oracle Data Guard on both the primary and standby databases.

You need to perform the following steps on only one node in each cluster (cluster-paris and cluster-newyork).

```
oracle (phys-newyork-1)$ sqlplus '/ as sysdba'
SQL> alter system set dg_broker_start=true scope=both sid='*';
SQL> quit
oracle (phys-paris-1)$ sqlplus '/ as sysdba'
SQL> alter system set dg_broker_start=true scope=both sid='*';
SQL> quit
```

▼ How to Create and Enable an Oracle Data Guard Broker Configuration

To use Oracle Data Guard with Geographic Edition, you need to create an Oracle Data Guard Broker configuration.

In the following example procedure, the Oracle Data Guard Broker configuration is called `mysales.com`. The `salesdr` database is a physical copy of the `sales` database.

1 Create an Oracle Data Guard Broker configuration for the primary database.

You use the `dgmgrl` command to create the Oracle Data Guard Broker configuration. You need to know the name of the Oracle Data Guard Broker configuration that you want to create, the name of the primary database, and the net service name through which to connect. You will need to know these properties again, when you specify the configuration to Geographic Edition.

```
oracle (phys-paris-1)$ dgmgrl sys/sysdba_password@sales-svc
DGMGR> create configuration mysales.com as primary
DGMGR> database is sales connect identifier is sales-svc;
```

If you find errors when you connect to the Oracle Data Guard Broker, check the `#{ORACLE_HOME}/admin/sales/bdump/alert_prim_sid.log` file. You can check that the configuration has been created by using the following command:

```
oracle (phys-paris-1)$ dgmgrl sys/sysdba_password@sales-svc
DGMGR> show configuration;
Configuration
  Name:                mysales.com
  Enabled:             NO
  Protection Mode:    MaxPerformance
  Fast-Start Failover: DISABLED
  Databases:
```

```
sales - Primary database
```

```
Current status for "mysales.com":
DISABLED
```

2 Add the standby database to the Oracle Data Guard Broker configuration.

You need to know the name of the standby database, the net service name through which to connect, and the type of standby (physical or logical).

```
oracle (phys-paris-1)$ dgmgrrl sys/sysdba_password@sales-svc
DGMGRL> add database salesdr as connect identifier is
salesdr-svc maintained as physical;
```

3 Configure the apply instance for the standby database.

If the standby database is also a multi-instance Oracle RAC database, you can specify the instance on which you would prefer the transmitted archive redo logs to be applied. Before you enable the configuration, issue the following command:

```
oracle$ dgmgrrl sys/sysdba_password@sales-svc
DGMGRL> edit database salesdr set property PreferredApplyInstance='salesdr1';
```

4 To verify that the Oracle Data Guard Broker configuration is working correctly, enable the configuration.

```
oracle (phys-paris-1)$ dgmgrrl sys/sysdba_password@sales-svc
DGMGRL> enable configuration;
```

If you have successfully performed all steps, you can check the status of the configuration by using the following command:

```
oracle$ dgmgrrl sys/sysdba_password@sales-svc
DGMGRL> show configuration;
Configuration
Name: mysales.com
Enabled: YES
Protection Mode: MaxPerformance
Fast-Start Failover: DISABLED
Databases:
sales - Primary database
salesdr - Physical standby database
```

```
Current status for "mysales.com":
SUCCESS
```

5 Verify that the Oracle Data Guard Broker configuration can switch over.

Before you add the Oracle Data Guard Broker configuration to Geographic Edition, you need to verify that you can perform a switchover of the database from the primary to the standby and back again. If this switchover does not work, Geographic Edition will not be able to perform this operation either.

```
oracle (phys-paris-1)$ dgmgrrl sys/sysdba_password@sales-svc
DGMGRL> switchover to salesdr
Performing switchover NOW, please wait...
Operation requires shutdown of instance "sales1" on database "sales"
```

```
Shutting down instance "sales1"...
ORA-01109: database not open

Database dismounted.
ORACLE instance shut down.
Operation requires shutdown of instance "salesdr1" on database "salesdr"
Shutting down instance "salesdr1"...
ORA-01109: database not open

Database dismounted.
ORACLE instance shut down.
Operation requires startup of instance "sales1" on database "sales"
Starting instance "sales1"...
ORACLE instance started.
Database mounted.
Operation requires startup of instance "salesdr1" on database "salesdr"
Starting instance "salesdr1"...
ORACLE instance started.
Database mounted.
Switchover succeeded, new primary is "salesdr"

DGMGR> switchover to sales;
Performing switchover NOW, please wait...
Operation requires shutdown of instance "salesdr1" on database "salesdr"
Shutting down instance "salesdr1"...
ORA-01109: database not open

Database dismounted.
ORACLE instance shut down.
Operation requires shutdown of instance "sales1" on database "sales"
Shutting down instance "sales1"...
ORA-01109: database not open

Database dismounted.
ORACLE instance shut down.
Operation requires startup of instance "salesdr1" on database "salesdr"
Starting instance "salesdr1"...
ORACLE instance started.
Database mounted.
Operation requires startup of instance "sales1" on database "sales"
Starting instance "sales1"...
ORACLE instance started.
Database mounted.
Switchover succeeded, new primary is "sales"
```

Administering Oracle Data Guard Protection Groups

This chapter describes how to administer data replication with Oracle Data Guard software.

This chapter covers the following topics:

- “Working With Oracle Data Guard Protection Groups” on page 35
- “Creating, Modifying, Validating, and Deleting an Oracle Data Guard Protection Group” on page 42
- “Administering Oracle Data Guard Application Resource Groups” on page 48
- “Administering Oracle Data Guard Broker Configurations” on page 52
- “Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster” on page 58
- “Activating and Deactivating a Protection Group” on page 60
- “Resynchronizing an Oracle Data Guard Protection Group” on page 65
- “Checking the Runtime Status of Oracle Data Guard Data Replication” on page 66

Working With Oracle Data Guard Protection Groups

Unlike other data replication mechanisms, such as Sun StorageTek Availability Suite, Hitachi TrueCopy or Universal Replicator, and EMC SRDF, Oracle Data Guard is an integral part of Oracle RAC software. Consequently, you do not place Oracle RAC server proxy resource groups under Geographic Edition control as you do when you are using one of these host or storage-based data replication mechanisms. You can add Oracle Data Guard Broker configurations for databases that are being replicated by Oracle Data Guard to Geographic Edition without stopping the databases or the replication.

Overview of Administering Protection Groups

To add an existing Oracle Data Guard Broker configuration that contains an Oracle Data Guard replicated database to a new protection group, you will complete the following general procedures.

1. On a node in either cluster, create the protection group.
This procedure is covered in “[How to Create and Configure an Oracle Data Guard Protection Group](#)” on page 42.
2. On the same node, add the Oracle Data Guard Broker configuration to the protection group.
This procedure is covered in “[How to Add an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group](#)” on page 53.
3. On a node in the *other* cluster, retrieve the protection group configuration.
This procedure is covered in “[How to Replicate the Oracle Data Guard Protection Group Configuration to a Partner Cluster](#)” on page 58.
4. On the same node, add the Oracle shadow RAC server proxy resource group and application resource group to the protection group.
This procedure is covered in “[How to Add an Application Resource Group to an Oracle Data Guard Protection Group](#)” on page 49.
5. Activate the protection group, either globally from either cluster or locally from the primary.
This procedure is covered in “[How to Activate an Oracle Data Guard Protection Group](#)” on page 60.

▼ How to Administer an Oracle Data Guard Protection Group (Example)

Note – The following example shows all the steps that are involved in administering Oracle Data Guard protection groups, as described in more detail in procedures that are included later in this chapter.

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

```
phys-paris-1# geopg create -d odg -o primary -s paris-newyork-ps sales-pg
Protection group "sales-pg" has been successfully created
```

The `cluster-paris` cluster is the primary cluster. You do not need to set any additional Oracle Data Guard protection group properties.

2 Add the Oracle Data Guard Broker configuration, `mysales.com`, to the protection group.

This command creates the `mysales_com-rac-proxy-svr-shadow-rg` shadow resource group.



Caution – To ensure security, do *not* supply a password when you specify the `sysdba_password` property. 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. You can pipe the password to the command if you want to drive the `geopg` command from another shell script.

Also, to run the following command successfully, you must already be able to connect to both a local and a remote database service.

```
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"
```

3 Confirm that the shadow Oracle RAC and replication resource groups and resources that you added to the protection group in the preceding step were added.

```
phys-paris-1# clresourcegroup status
=== Cluster Resource Groups ===
```

Group Name	Node Name	Suspended	Status
rac-framework-rg	phys-paris-1	No	Online
	phys-paris-2	No	Online
scal-oradata-dg-rg	phys-paris-1	No	Online
	phys-paris-2	No	Online
qfs-oradata-mds-rg	phys-paris-1	No	Online
	phys-paris-2	No	Offline
scal-oradata-mp-rg	phys-paris-1	No	Online
	phys-paris-2	No	Online
rac_server_proxy-rg	phys-paris-1	No	Online
	phys-paris-2	No	Online
geo-clusterstate	phys-paris-1	No	Online
	phys-paris-2	No	Online
geo-infrastructure	phys-paris-1	No	Offline

	phys-paris-2	No	Online
sales-pg-odg-rep-rg	phys-paris-1	No	Online
	phys-paris-2	No	Offline
mysales_com-rac-proxy-svr-shadow-rg	phys-paris-1	No	Unmanaged
	phys-paris-2	No	Unmanaged

phys-paris-1# clresource status

Resource Name	Node Name	State	Status Message
-----	-----	----	-----
rac-framework-rs	phys-paris-1	Online	Online
	phys-paris-2	Online	Online
rac-udlm-rs	phys-paris-1	Online	Online
	phys-paris-2	Online	Online
rac-svm-rs	phys-paris-1	Online	Online
	phys-paris-2	Online	Online
crs_framework-rs	phys-paris-1	Online	Online
	phys-paris-2	Online	Online
scal-oradata-dg-rs	phys-paris-1	Online	Online - Diskgroup online
	phys-paris-2	Online	Online - Diskgroup online
qfs-oradata-mds-rs	phys-paris-1	Online	Online - Service is online.
	phys-paris-2	Offline	Offline
scal-oradata-mp-rs	phys-paris-1	Online	Online
	phys-paris-2	Online	Online
rac_server_proxy-rs	phys-paris-1	Online	Online - Oracle instance UP
	phys-paris-2	Online	Online - Oracle instance UP
geo-servicetag	phys-paris-1	Online but not monitored	Online
	phys-paris-2	Online but not monitored	Online
geo-clustername	phys-paris-1	Offline	Offline
	phys-paris-2	Online	Online - LogicalHostname online.
geo-hbmonitor	phys-paris-1	Offline	Offline
	phys-paris-2	Online	Online - Daemon OK
geo-failovercontrol	phys-paris-1	Offline	Offline
	phys-paris-2	Online	Online - Service is online.
mysales_com-odg-rep-rs	phys-paris-1	Offline	Offline
	phys-paris-2	Offline	Offline
mysales_com-rac-proxy-svr-shadow-rs	phys-paris-1	Offline	Offline
	phys-paris-2	Offline	Offline

4 Locally activate the protection group.

```
phys-paris-1# geopg start -e local sales-pg
Processing operation... The timeout period for this operation on
  each cluster is 3600 seconds (3600000 milliseconds)...
Protection group "sales-pg" successfully started.
```

If your `mysales.com` Oracle Data Guard Broker configuration is not already enabled, this process might take a few minutes or more. The actual time that the process takes depends on the configuration of your primary and standby databases as well as the distance between the clusters.

5 Verify that the data replication is successfully started.

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

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

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

Protection group "sales-pg"         : Error
Partnership                         : paris-newyork-ps
Synchronization                     : Error

Cluster cluster-paris               : OK
Role                                 : Primary
Activation State                     : Activated
Configuration                        : OK
Data replication                     : OK
Resource groups                      : None

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

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

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

7 Confirm that the shadow Oracle RAC and replication resource groups and resources for the protection group that you retrieved in the preceding step were retrieved.

```
phys-newyork-1# clresourcegroup status
```

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

Group Name	Node Name	Suspended	Status
------------	-----------	-----------	--------

Resource Name	Node Name	State	Status Message
rac-framework-rg	phys-newyork-1	No	Online
	phys-newyork-2	No	Online
scal-oradata-dg-rg	phys-newyork-1	No	Online
	phys-newyork-2	No	Online
qfs-oradata-mds-rg	phys-newyork-1	No	Online
	phys-newyork-2	No	Offline
scal-oradata-mp-rg	phys-newyork-1	No	Online
	phys-newyork-2	No	Online
rac_server_proxy-rg	phys-newyork-1	No	Online
	phys-newyork-2	No	Online
geo-clusterstate	phys-newyork-1	No	Online
	phys-newyork-2	No	Online
geo-infrastructure	phys-newyork-1	No	Offline
	phys-newyork-2	No	Online
sales-pg-odg-rep-rg	phys-newyork-1	No	Online
	phys-newyork-2	No	Offline
mysales_com-rac-proxy-svr-shadow-rg	phys-newyork-1	No	Unmanaged
	phys-newyork-2	No	Unmanaged

phys-newyork-1# clresource status

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
rac-framework-rs	phys-newyork-1	Online	Online
	phys-newyork-2	Online	Online
rac-udlm-rs	phys-newyork-1	Online	Online
	phys-newyork-2	Online	Online
rac-svm-rs	phys-newyork-1	Online	Online
	phys-newyork-2	Online	Online
crs_framework-rs	phys-newyork-1	Online	Online
	phys-newyork-2	Online	Online
scal-oradata-dg-rs	phys-newyork-1	Online	Online - Diskgroup online
	phys-newyork-2	Online	Online - Diskgroup online
qfs-oradata-mds-rs	phys-newyork-1	Online	Online - Service is online.
	phys-newyork-2	Offline	Offline
scal-oradata-mp-rs	phys-newyork-1	Online	Online
	phys-newyork-2	Online	Online
rac_server_proxy-rs	phys-newyork-1	Online	Online - Oracle instance UP
	phys-newyork-2	Online	Online - Oracle instance UP
geo-servicetag	phys-newyork-1	Online but not monitored	Online

	phys-newyork-2	Online but not monitored	Online
geo-clustername	phys-newyork-1 phys-newyork-2	Offline Online	Offline Online - LogicalHostname online.
geo-hbmonitor	phys-newyork-1 phys-newyork-2	Offline Online	Offline Online - Daemon OK
geo-failovercontrol	phys-newyork-1 phys-newyork-2	Offline Online	Offline Online - Service is online.
mysales_com-odg-rep-rs	phys-newyork-1 phys-newyork-2	Offline Offline	Offline Offline
mysales_com-rac-proxy-svr-shadow-rs	phys-newyork-1 phys-newyork-2	Offline Offline	Offline Offline

8 From any node in a partner cluster, add the shadow Oracle RAC server proxy resource group to the protection group.

```
# geopg add-resource-group mysales_com-rac-proxy-svr-shadow-rg sales-pg
Following resource groups were successfully added:
"mysales_com-rac-proxy-svr-shadow-rg"
```

Adding the shadow Oracle RAC server proxy resource group to the protection group is not critical to the operation of the replication. The resource contained within it simply reflects the status of the real Oracle RAC server proxy resource group and highlights whether the cluster is the Oracle Data Guard primary cluster.

9 From any node in a partner cluster, globally activate the protection group on both clusters.

```
# geopg start -e global sales-pg
Processing operation... The timeout period for this operation on
each cluster is 3600 seconds (3600000 milliseconds)...
Protection group "sales-pg" successfully started.
```

10 Verify that the protection group is successfully created and activated.

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

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

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

Protection group "sales-pg" : OK
Partnership              : "paris-newyork-ps"
Synchronization          : OK

Cluster cluster-newyork   : OK
Role                      : Primary
```

```
PG activation state : Activated
Configuration      : OK
Data replication   : OK
Resource groups    : OK

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

Creating, Modifying, Validating, and Deleting an Oracle Data Guard Protection Group

This section covers the following topics:

Note – You can create protection groups that are not configured to use data replication. To create a protection group that does not use a data replication subsystem, omit the `-d datareplicationtype` option when you use the `geopg` command. If you omit this option, the `geoadm status` command shows that the state of data replication is `NONE`.

▼ How to Create and Configure an Oracle Data Guard Protection Group

The following example builds on the example configuration that was described in [Chapter 1, “Replicating Data With Oracle Data Guard Software.”](#)

In this example, the `sales` database is online on the `cluster-paris` cluster and is protected by Oracle Data Guard.

Ensure that the `mysales.com` Oracle Data Guard Broker configuration exists before you proceed, as Geographic Edition does *not* create the configuration for you.

Before You Begin Ensure that the following conditions are met:

- Your clusters are members of a partnership.
- The protection group that you are creating does not already exist.

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

You can also replicate the existing configuration of a protection group from a remote cluster to the local cluster. For more information, see “[Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster](#)” on page 58.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 On all nodes of the local cluster, create a new protection group.

```
phys-node-n# geopg create -s partnershipname -d odg \
-o localrole [-p property [-p...]] protectiongroupname
```

- s *partnershipname* Specifies the name of the partnership.
- d odg Specifies that the protection group data is replicated by Oracle Data Guard software.
- o *localrole* Specifies the role of this protection group on the local cluster as either primary or secondary.
- p *propertysetting* Specifies the properties of the protection group.

You can specify the following properties:

- **Description** – Describes the protection group.
- **Timeout** – Specifies the timeout period for the protection group, in seconds.

protectiongroupname 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,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

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

- If the validation is successful, the local Configuration status is set to OK and the Synchronization status is set to Error.
- If the validation is unsuccessful, the protection group is not created.

▼ How to Modify an Oracle Data Guard Protection Group

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

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Modify the configuration of the protection group.

```
phys-node-n# geopg set-prop -p property[-p...] protectiongroupname
-p property
```

Specifies the properties of the protection group.

For more information about the properties that you can set, see [Appendix A, “Standard Geographic Edition Properties,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

```
protectiongroupname
```

Specifies the name of the protection group.

If the partner cluster contains a protection group of the same name, the `geopg set-prop` command also propagates the new configuration information to the partner cluster.

The `geopg set-prop` command revalidates the protection group with the new configuration information. If the validation is unsuccessful on the local cluster, the configuration of the protection group is not modified. Otherwise, the Configuration is modified and its status is set to OK on the local cluster.

If the Configuration status is set to OK on the local cluster, but the validation is unsuccessful on the partner cluster, the Configuration is modified on the partner cluster and the configuration status is set to Error on the partner cluster.

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,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

Example 2-1 Modifying the Configuration of a Protection Group

This example shows how to modify the `timeout` property of a protection group.

```
phys-paris-1# geopg set-prop -p Timeout=300 sales-pg
```

▼ How to Validate an Oracle Data Guard Protection Group

Before You Begin When the `Configuration` status of a protection group is displayed as `Error` in the output of the `geoadm status` command, you can validate the configuration by using the `geopg validate` command. This command checks the current state of the protection group and its entities.

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

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

Before validating the configuration of a protection group, ensure that the protection group you want to validate exists locally and that the common agent container is online on all nodes of both clusters in the partnership.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Validate the configuration of the protection group.

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

```
phys-node-n# geopg validate protectiongroupname
```

Example 2–2 Validating the Configuration of a Protection Group

This example shows how to validate a protection group.

```
phys-node-n# geopg validate sales-pg
```

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

During protection group validation, the Oracle Data Guard data replication layer validates the application resource groups and the data replication entities. The Oracle Data Guard data replication layer verifies the following conditions:

- **That the resource group in the protection group that is being validated does not contain an Oracle Oracle RAC server proxy resource group that contains an Oracle Oracle RAC server proxy resource**

You cannot add these resource groups to an Oracle Data Guard protection group because the Oracle RAC database that is managed by the Oracle RAC server proxy resource is shut down on the standby cluster when the protection group is started globally, thus disabling the Oracle Data Guard data replication.

- **That the `Auto_start_on_new_cluster` property in an application resource group in the protection group is set to `False`**

When you bring a protection group online on the primary cluster, the data replication layer brings the application resources groups that are participating in that protection group online only on the same primary cluster. Setting the `Auto_start_on_new_cluster` property to `False` prevents the Oracle Solaris Cluster resource group manager from automatically starting the application resource groups. In this case, the startup of resource groups is reserved for the Geographic Edition software.

When the protection group is activated, application resource groups need to be online only on the primary cluster. Ensure that the following conditions are met:

- **That the Oracle `dgmgrl` command is showing a `SUCCESS` status for the each of the Oracle Data Guard Broker configurations**

The presence of Oracle `ORA-` messages in the output from the `dgmgrl` command might indicate that the `sysdba_username` password is incorrect or that the cluster has been disabled. This information is reflected in the status of the replication resource for the Oracle Data Guard Broker configuration.

- **That the Oracle Data Guard Broker configuration details match those 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 cluster), the replication mode, and the standby type.

- **That the `sysdba_username` password is valid for the standby cluster, to ensure that switchovers are possible**

When validation is complete, the Geographic Edition software creates the shadow RAC server proxy resource group and resource, the replication resource group, and the resources for this replication resource group, if they do not exist, and brings them online. If a resource group or a resource with the same name already exists, the Geographic Edition operations might modify their properties. Geographic Edition software cannot create a new resource group or resource of the same name if one already exists.

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

▼ How to Delete an Oracle Data Guard Protection Group

Before You Begin To delete a protection group on all clusters, run the `geopg delete` command on each cluster where the protection group exists.

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

- The protection group exists locally
- The protection group is offline on the local cluster

Note – To keep the application resource groups online while deleting a protection group, remove the application resource groups from the protection group before deleting the protection group. You do not need to do anything to shadow RAC server proxy resource groups, as deleting the protection group removes these resource groups without affecting the Oracle RAC server proxy resource groups that they shadow.

1 Log in to a node in the cluster where you want to delete the protection group, for example, `cluster-paris`.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

`cluster-paris` is the primary cluster. See “[Example Geographic Edition Cluster Configuration](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide* for a sample cluster configuration.

2 Delete the protection group.

```
phys-node-n# geopg delete protectiongroupname
```

This command deletes the configuration of the protection group from the local cluster. The command also removes the Oracle Oracle RAC server proxy resource groups and the replication resource group for the Oracle Data Guard Broker configuration in the protection group.

If the protection group is not deleted, the Configuration status is set to Error. Resolve the error and rerun the `geopg delete` command.

Example 2-3 Deleting a Protection Group

This example shows how to delete a protection group from both partner clusters.

```
# rlogin cluster-paris -l root
phys-paris-1# geopg delete sales-pg
# rlogin cluster-newyork -l root
phys-newyork-1# geopg delete sales-pg
```

Example 2-4 Deleting a Protection Group While Keeping Application Resource Groups Online

This example shows how to keep two application resource groups, `apprg1` and `apprg2`, online, while deleting the protection group that they share, `sales-pg`.

Remove the application resource groups from the protection group and delete the protection group.

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

Unlike other data replication modules, the Oracle RAC server proxy resource group is not added to the protection group. Instead, a shadow RAC server proxy resource group is added to represent this resource group. You can add and remove the shadow RAC server proxy 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 can have no data to replicate, as only Oracle Data Guard data replication is supported in this particular 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.

Administering Oracle Data Guard Application Resource Groups

To make an application highly available, you must ensure that the application is managed as a resource in an application resource group. Unlike other data replication modules, the Oracle RAC server proxy resource group is not added to the protection group. Instead, a shadow RAC server proxy resource group is added to represent this resource group.

You can add and remove the shadow RAC server proxy resource group to and from the protection group at any time without affecting the Oracle Data Guard data replication. This fact does not prevent you from adding other, non-Oracle RAC server proxy resource groups to the protection group if necessary. However, these applications cannot use any data that requires replication to the standby cluster as only Oracle Data Guard is supported in this type of protection group.

You need to replicate, on the standby cluster, all entities that you configure for the primary cluster's application resource group. Examples of entities that you need to replicate are application data resources, configuration files, and resource groups. Resource group names must also match on both clusters. In addition, the data that the application resource uses needs to be replicated on the standby cluster.

This section shows you how to perform the following procedures:

- “How to Add an Application Resource Group to an Oracle Data Guard Protection Group” on page 49
- “How to Delete an Application Resource Group From an Oracle Data Guard Protection Group” on page 51

▼ How to Add an Application Resource Group to an Oracle Data Guard Protection Group

Before You Begin You can add an existing resource group, other than an Oracle RAC server proxy resource group containing an Oracle RAC server proxy resource, to the list of application resource groups for a protection group. If you do try to add an Oracle RAC server proxy resource group, the `geopg` command returns an error.

Before you add an application resource group (of any other type) to a protection group, ensure that the following conditions are met:

- The protection group is defined.
- The application resource group does not need any data replicating. You are not prevented from adding such resource groups, but the Oracle Data Guard module does not coordinate the switchover of other types of data replication.
- The resource group to add already exists on both clusters and is in an appropriate state.
- The `Auto_start_on_new_cluster` property of the resource group is set to `False`. You can determine the setting of this property by using the `clresourcegroup show` command.

```
phys-node-n# clresourcegroup show -p auto_start_on_new_cluster apprg1
```

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

```
phys-node-n# clresourcegroup set -p Auto_start_on_new_cluster=False apprg1
```

Setting the `Auto_start_on_new_cluster` property to `False` prevents the Oracle Solaris Cluster resource group manager from automatically starting the resource groups in the protection group.

When the protection group is activated, application resource groups need to be online only on the primary cluster.

The application resource group does not have dependencies on resource groups and resources outside of this protection group unless the `External_Dependency_Allowed` protection group property is set to `TRUE`. To add several application resource groups that share dependencies while the `External_Dependency_Allowed` protection group property is set to `FALSE`, you need to add all the application resource groups that share dependencies to the protection group in a single operation. If you add the application resource groups separately, the operation fails.

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

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

If the resource group to add is `Online` and the protection group is deactivated, the request is rejected. Before you add an online resource group, you need to activate the protection group.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Add an application resource group to the protection group.

```
phys-node-n# geogg add-resource-group resourcegrouplist protectiongroup
```

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

protectiongroup Specifies the name of the protection group.

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

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,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

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

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

Example 2-5 Adding an Application Resource Group to an Oracle Data Guard Protection Group

This example shows how to add two application resource groups, `apprg1` and `apprg2`, to `sales-pg`.

```
phys-paris-1# geopg add-resource-group apprg1,apprg2 sales-pg
```

▼ How to Delete an Application Resource Group From an Oracle Data Guard Protection Group

You can remove an application resource group from a protection group without altering the state or contents of the application resource group. You can remove shadow RAC server proxy resource groups at any time, without affecting the Oracle RAC server proxy resource groups or Oracle RAC databases that they represent. You can remove these resource groups because the shadow RAC server proxy resource groups simply reflect the status of the real Oracle RAC server proxy resource groups and do not control the Oracle RAC databases.

Before You Begin Ensure that the following conditions are met:

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

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see [“Geographic Edition Software and RBAC”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Remove the application resource group from the protection group.

```
phys-node-n# geopg remove-resource-group resourcegrouplist protectiongroup
resourcegrouplist
```

Specifies the name of the application resource group.

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

```
protectiongroup
```

Specifies the name of the protection group.

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

If the resource group that is being removed shares dependencies with other resource groups in the protection group and the `External_Dependency_Allowed` protection group property is set to `FALSE`, you also need to remove all other resource groups that share dependencies with the resource group that is being removed.

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

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

Example 2-6 Deleting an Application Resource Group From a Protection Group

This example shows how to remove two application resource groups, `apprg1` and `apprg2`, from `sales-pg`.

```
phys-paris-1# geopg remove-resource-group apprg1,apprg2 sales-pg
```

Administering Oracle Data Guard Broker Configurations

The following procedures describe how to administer Oracle Data Guard Broker data replication configurations in an Oracle Data Guard protection group.

- [“How to Add an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group” on page 53](#)
- [“How the Data Replication Subsystem Verifies the Oracle Data Guard Broker Configuration” on page 55](#)
- [“How to Modify an Oracle Data Guard Broker Configuration” on page 56](#)
- [“How to Delete an Oracle Data Guard Broker Configuration From an Oracle Data Guard Protection Group” on page 57](#)

For details about configuring an Oracle Data Guard protection group, see [“How to Create and Configure an Oracle Data Guard Protection Group”](#) on page 42.

▼ How to Add an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group

Before You Begin

A protection group is the container for the application resource groups, which contain data for services that are protected from disaster. Geographic Edition software protects the data by replicating it from the primary cluster to the standby cluster. By adding an Oracle Data Guard Broker configuration to a protection group, Geographic Edition software monitors the replication status of the Oracle RAC database that belongs to that Oracle Data Guard Broker configuration.

Geographic Edition software also controls the role and state of the Oracle Data Guard Broker configuration during protection group operations, such as start, stop, switchover, and takeover.

Before you add an Oracle Data Guard Broker configuration to a protection group, ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- If the partner cluster can be reached, the protection group is offline on the local cluster and the partner cluster.
- The Oracle Data Guard Broker configuration exists on both the local cluster and the partner cluster.
- The Oracle RAC server proxy resource group and Oracle RAC server proxy resources that manage the Oracle RAC databases that are replicated by Oracle Data Guard exist on both the local and the partner cluster.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see [“Geographic Edition Software and RBAC”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Add an Oracle Data Guard Broker configuration to the protection group.

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

```
phys-node-n# geopg add-replication-component \  
-p property [-p...] ODGConfigurationName protectiongroupname
```

-p *property*

Specifies the properties of either the Oracle Data Guard Broker configuration, the Oracle RAC server proxy resource group, or the Oracle database user name and the associated password.

You can specify the following properties:

- `local_database_name` – Name of the local database in the Oracle Data Guard Broker configuration.
- `local_db_service_name` – Oracle net service name for the local database.
- `local_rac_proxy_svr_rg_name` – Name of the local Oracle RAC server proxy resource group that manages the local database in the Oracle Data Guard Broker configuration.
- `remote_database_name` – Name of the remote database in the Oracle Data Guard Broker configuration.
- `remote_db_service_name` – Oracle net service name for the remote database.
- `remote_rac_proxy_svr_rg_name` – Name of the Oracle RAC server proxy resource group on the partner cluster that manages the remote database in the Oracle Data Guard Broker configuration.
- `replication_mode` – Replication mode for the database in the Oracle Data Guard Broker configuration.
- `standby_type` – Standby type for the database in the Oracle Data Guard Broker configuration.
- `sysdba_password` – Password for the Oracle SYSDBA privileged database user. Do not specify the actual 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.
- `sysdba_username` – Name of an Oracle SYSDBA privileged database user who can perform the Oracle Data Guard Broker switchover and takeover operations.

For more information about the properties that you can set, see [Appendix A, “Standard Geographic Edition Properties,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

ODGConfigurationName

Specifies the name of the new Oracle Data Guard Broker configuration.

protectiongroupname

Specifies the name of the protection group that contains the new Oracle Data Guard Broker configuration.

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,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

Example 2-7 Adding an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group

This example shows how to add an Oracle Data Guard Broker configuration to the sales-pg protection group.

To run the following command successfully, you must already be able to connect to both a local and a remote database service.

```
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
```

How the Data Replication Subsystem Verifies the Oracle Data Guard Broker Configuration

When you add an Oracle Data Guard Broker configuration to a protection group, the data replication layer verifies that the Oracle Data Guard Broker configuration exists.

When you run the `geopg add-replication-component` command, if they do not already exist, an Oracle shadow RAC server proxy resource group and a replication resource group for the Oracle Data Guard Broker configuration are created. In addition, the configuration is successfully validated.

The Oracle shadow RAC server proxy resource group contains a Oracle Solaris Cluster resource. This resource is based on the generic data service `SUNW.gds` resource type. The Oracle shadow RAC server proxy resource shadows the real Oracle RAC server proxy resource that manages and monitors the Oracle RAC database in the Oracle Data Guard Broker configuration.

For more information about the shadow RAC server proxy resource group, see [“Oracle Data Guard Shadow Resource Groups”](#) on page 15.

The replication resource group contains a Oracle Solaris Cluster resource that is based on the generic data service `SUNW.gds` resource type. The replication resource monitors the state of the database replication as reported by Oracle Data Guard Broker.

For more information about replication resources, see [“Oracle Data Guard Replication Resource Groups”](#) on page 16.

For the validation to be successful, ensure that the following conditions are met:

- The resource group that is named in the `local_rac_proxy_svr_rg_name` property contains a resource of type `SUNW.scaLable_rac_server_proxy`. This resource is used to determine the values for `{ORACLE_HOME}` and the local Oracle RAC SID values.
- The Oracle `dgmgrl` command shows a `SUCCESS` status for the Oracle Data Guard Broker configuration. The presence of Oracle `ORA-` messages in the output from the `dgmgrl` command might indicate that the `sysdba_username` password is incorrect or that the cluster has been disabled. Oracle errors are returned as part of the messages that are generated by the `validate` command.
- The `sysdba_username` password is valid for the standby cluster to ensure that switchovers are possible.
- The Oracle Data Guard Broker configuration details match those 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 cluster), the replication mode, and the standby type.



Caution – Do not use Oracle Solaris Cluster commands to change, remove, or bring offline these resources or resource groups. Use only Geographic Edition commands to administer shadow RAC server proxy resource groups, replication resource groups, and resources that are internal entities that are managed by Geographic Edition software. Altering the configuration or state of these entities directly with Oracle Solaris Cluster commands could result in an unrecoverable failure.

▼ How to Modify an Oracle Data Guard Broker Configuration

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Modify the Oracle Data Guard Broker configuration.

This command modifies the properties of an Oracle Data Guard Broker configuration 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 of the same name.

```
phys-node-n# geopg modify-replication-component -p property [-p...] \  
ODGConfigurationName protectiongroupname
```

`-p property`

Specifies the properties of the data replication Oracle Data Guard Broker configuration.

For more information about the properties that you can set, see [Appendix A, “Standard Geographic Edition Properties,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

ODGConfigurationName

Specifies the name of the Oracle Data Guard Broker configuration.

protectiongroupname

Specifies the name of the protection group that contains the Oracle Data Guard Broker configuration.

▼ How to Delete an Oracle Data Guard Broker Configuration From an Oracle Data Guard Protection Group

Before You Begin Before you remove an Oracle Data Guard Broker configuration from a protection group, ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- If the partner cluster can be reached, the protection group is offline on the local cluster and the partner cluster.
- The Oracle Data Guard Broker configuration is managed by the protection group.

For information about deleting protection groups, refer to [“How to Delete an Oracle Data Guard Protection Group”](#) on page 47.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see [“Geographic Edition Software and RBAC”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Remove the Oracle Data Guard Broker configuration.

This command removes an Oracle Data Guard Broker configuration from 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 of the same name.

This command removes the Oracle Data Guard Broker configuration from the protection group. This command also deletes the Oracle shadow RAC server proxy resource group and replication resource group for this Oracle Data Guard Broker configuration.

`phys - node - n# geopg remove-replication-component ODGConfigurationName protectiongroupname`

ODGConfigurationName Specifies the name of the Oracle Data Guard Broker configuration.

protectiongroupname Specifies the name of the protection group.

Example 2-8 Deleting an Oracle Data Guard Broker Configuration From an Oracle Data Guard Protection Group

This example shows how to delete an Oracle Data Guard Broker configuration from an Oracle Data Guard protection group.

```
phys-paris-1# geopg remove-replication-component mysales.com sales-pg
```

Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster

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

▼ How to Replicate the Oracle Data Guard Protection Group Configuration to a Partner Cluster

Before You Begin Before you replicate the configuration of an Oracle Data Guard protection group to a partner cluster, ensure that the following conditions are met:

- The protection group is defined on the remote cluster, not on the local cluster.
- The Oracle Data Guard Broker configuration in the protection group on the remote cluster exists on the local cluster.
- The application resource groups in the protection group on the remote cluster exist on the local cluster.
- The `Auto_start_on_new_cluster` property of the resource groups is set to `False`. You can view this property by using the `clresourcegroup show` command.

```
phys-node-n# clresourcegroup show -p Auto_start_on_new_cluster apprg
```

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

```
phys-node-n# clresourcegroup set -y Auto_start_on_new_cluster=False apprg1
```

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

When the protection group is activated, application resource groups need to be online only on the primary cluster.

- You have *not* added the shadow RAC server proxy resource group for an Oracle Data Guard Broker configuration to a protection group application resource group list before that resource group exists on all clusters.

Note – You must replicate the protection group configuration to a partner cluster *before* you can add a shadow RAC server proxy resource group to a protection group.

When you successfully add the Oracle Data Guard configuration to the protection group on the clusters on which the protection group exists, Oracle Data Guard creates the shadow RAC server proxy resource group on the clusters. The means by which you can successfully add a shadow RAC server proxy resource group to a protection group include the following:

If an Oracle Data Guard protection group does not contain an Oracle Data Guard Broker configuration, once you replicate the protection group on the partner cluster and add the Oracle Data Guard Broker configuration to it, Oracle Data Guard adds the shadow RAC server proxy resource group on both clusters.

If an Oracle Data Guard protection group contains an Oracle Data Guard Broker configuration, does not contain a shadow RAC server proxy resource group on one cluster, and is not replicated on the partner cluster, when you replicate the protection group to the partner cluster, Oracle Data Guard creates the shadow RAC server proxy resource group on the partner cluster.

Once a shadow RAC server proxy resource group exists on both clusters, you can add that resource group to the protection group.

1 Log in to phys-newyork-1.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Replicate the protection group configuration to the partner cluster.

```
phys-newyork-1# geopg get -s partnershipname ODGprotectiongroup
```

-s partnershipname Specifies the name of the partnership from which the protection group configuration information is gathered.

ODGprotectiongroup Specifies the name of the protection group.

The `geopg get` command retrieves the configuration information of the protection group from the remote cluster and creates the protection group on the local cluster.

Note – The `geogg get` command replicates Geographic Edition related entities. For information about how to replicate Oracle Solaris Cluster entities, see [“Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

Example 2–9 Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster

This example shows how to replicate the configuration of `sales-pg` to `cluster-newyork`.

```
# rlogin phys-newyork-1 -l root
phys-newyork-1# geogg get -s paris-newyork-ps sales-pg
```

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

- If the validation is successful, the Configuration status is set to OK and the protection group is created on the local cluster.
- If the validation fails, the protection group is not created on the local cluster. Resolve the error and replicate the protection group again.

Activating and Deactivating a Protection Group

This section describes how to perform the following procedures:

- [“How to Activate an Oracle Data Guard Protection Group”](#) on page 60
- [“How to Deactivate an Oracle Data Guard Protection Group”](#) on page 63

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

For more information about configuring protection groups, see [“How to Create and Configure an Oracle Data Guard Protection Group”](#) on page 42.

▼ How to Activate an Oracle Data Guard Protection Group

You can activate a protection group in the following ways:

- Globally, which activates a protection group on both clusters where the protection group has been configured
- On the primary cluster only
- On a standby cluster only

When you activate a protection group, the data replication product that you are using determines the clusters on which data replication can start. For example, 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 a protection group locally from the standby cluster, data replication does not start. However, if you activate a protection group globally from the standby cluster, data replication starts.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Activate the protection group on the local cluster.

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

```
phys-node-n# geopg start -e scope [-n] ODGprotectiongroup
```

`-e scope`

Specifies the scope of the command.

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

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

`-n`

Prevents the start of data replication when the protection group starts.

If you omit this option, the data replication subsystem starts at the same time as the protection group, and the command performs the following operations on each Oracle Data Guard Broker configuration in the protection group:

- Verifies that the resource group that is named in the `local_rac_proxy_svr_rg_name` property contains a resource of type `SUNW.scalable_rac_server_proxy`.

- Verifies that the Oracle `dgmgrl` command can connect using the values that are given for `sysdba_username`, `sysdba_password`, and `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, and the standby type.

ODGprotectiongroup

Specifies the name of the protection group.

The `geopg start` command uses the `clrs enable resources` and `clrg online resourcegroups` command to bring resource groups and resources online. For more information about using this command, see the `clresource(1CL)` and `clresourcegroup(1CL)` man pages.

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, including the shadow RAC server proxy resource groups, in the protection group online on the local cluster

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 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 an Oracle Data Guard Protection Group” on page 45](#).

Example 2–10 Globally Activating an Oracle Data Guard Protection Group

This example shows how to activate a protection group globally.

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

Example 2–11 Activating an Oracle Data Guard Protection Group Locally

This example shows how to activate 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 an Oracle Data Guard Protection Group

You can deactivate a protection group in the following ways:

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

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.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Deactivate the protection group on all nodes of the local cluster.

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

```
phys-node-n# geopg stop -e scope [-D] protectiongroupname
```

-e scope

Specifies the scope of the command.

If the scope is `local`, the command operates on the local cluster only. If the scope is `global`, the command operates on both clusters where the protection group is located.

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

-D

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

If you omit this option, the data replication subsystem and the protection group are both stopped. If the role of the protection group on the local cluster is set to `primary` and you omit the `-D` option, the application resource groups are taken offline and put in an Unmanaged state.

protectiongroupname

Specifies the name of the protection group.

If the role of the protection group is primary on the local cluster, the `geopg stop` command disables the Oracle Data Guard Broker configuration.

If the `geopg stop` command fails, run the `geoadm status` command to see the status of each 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 an Oracle Data Guard Protection Group”](#) on page 45.

Example 2-12 Deactivating an Oracle Data Guard Protection Group on All Clusters

This example shows how to deactivate a protection group on all clusters.

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

Example 2-13 Deactivating an Oracle Data Guard Protection Group on a Local Cluster

This example shows how to deactivate a protection group on the local cluster.

```
phys-paris-1# geopg stop -e local sales-pg
```

Example 2-14 Stopping Oracle Data Guard Data Replication While Leaving the Protection Group Online

This example shows how to stop only data replication on a local cluster.

```
phys-paris-1# geopg stop -e local -D sales-pg
```

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

```
phys-paris-1# geopg stop -e local sales-pg
```

Example 2-15 Deactivating an Oracle Data Guard Protection Group While Keeping Application Resource Groups Online

This example shows how to keep online two application resource groups, `apprg1` and `apprg2`, while deactivating their protection group, `sales-pg`.

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

```
phys-paris-1# geopg remove-resource-group apprg1,apprg2 sales-pg
```

2. Deactivate the protection group.

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

Resynchronizing an Oracle Data Guard Protection Group

You can resynchronize the configuration information of the local protection group with the configuration information that you retrieve from the partner cluster. The cluster on which you run the command to resynchronize forfeits its own protection group configuration of the partner cluster. To determine if you need to resynchronize a protection group, you use the `geoadm status` command. If the value of the `Synchronization` parameter for a protection group is listed as `Error`, you need to resynchronize the protection group.

For example, you might need to resynchronize protection groups after booting the cluster. For more information, see “[Booting a Cluster](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Resynchronizing a protection group updates only entities that are related to Geographic Edition. For information about how to update Oracle Solaris Cluster entities, see “[Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

▼ How to Resynchronize an Oracle Data Guard Protection Group

Before You Begin You need to deactivate the protection group on the cluster where you run the `geopg update` command.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Resynchronize the protection group.

```
phys-node-n# geopg update protectiongroupname
```

Example 2–16 Resynchronizing an Oracle Data Guard Protection Group

This example shows how to resynchronize a protection group.

```
phys-paris-1# geopg update sales-pg
```

Checking the Runtime Status of Oracle Data Guard Data Replication

You can obtain an overall view of the status of replication, as well as a more detailed runtime status of the Oracle Data Guard software from the status of the replication resource groups. The following sections describe how to check the runtime status of replication:

- “Displaying an Oracle Data Guard Runtime Status Overview” on page 66
- “Displaying a Detailed Oracle Data Guard Runtime Status” on page 67

Displaying an Oracle Data Guard Runtime Status Overview

The status of each Oracle Data Guard data replication resource indicates the status of replication on a particular Oracle Data Guard Broker configuration. The status of all the resources under a protection group are aggregated in the replication status.

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

▼ How to Check the Overall Runtime Status of Replication

1 Log in to a node of a cluster where the protection group is defined.

To complete this step, you need to be assigned the Basic Solaris User RBAC rights profile. For more information about RBAC, see “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Check the runtime status of replication.

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

Refer to the Protection Group section of the output for replication information. The output of this command includes the following information:

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

3 Check the runtime status of data replication for each Oracle Data Guard protection group.

```
phys-paris-1 clresource status -v ODGConfigurationName-odg-rep-rs
```

Refer to the `Status` and `StatusMessage` fields that are presented for the Oracle Data Guard Broker configuration data replications that you want to check. For more information about these fields, see [Table 2-1](#).

Displaying a Detailed Oracle Data Guard Runtime Status

One replication resource group exists for each protection group. The name of the replication resource group conforms to the following format:

```
ODGprotectiongroupname-odg-rep-rg
```

If you add an Oracle Data Guard Broker configuration to a protection group, the Geographic Edition software creates a resource for that configuration. This resource monitors and displays the status of replication for the Oracle Data Guard Broker configuration. The name of each resource conforms to the following format:

```
ODGConfigurationName-odg-rep-rs
```

You can monitor the state of the replication resource to give you the overall status of replication. Use the `clresource status` command as follows to obtain the State and Status Message values for the replication status of the Oracle Data Guard Broker configuration:

```
phys-node-n# clresource status -v ODGConfigurationName-odg-rep-rs
```

The State is `Online` while the resource is online.

The following table describes the Status and Status Message values that are returned by the `clresource status` command when the State of the Oracle Data Guard replication resource group is `Online`.

TABLE 2-1 Status and Status Messages of an Online Oracle Data Guard Replication Resource Group

Status	Status Message	Possible Causes
Faulted	Program <i>program-name</i> returned a nonzero exit code	
Faulted	Protection mode " <i>replication-mode</i> " given for local database <i>database</i> does not match configured value " <i>replication-mode</i> "	The Oracle Data Guard Broker configuration has been changed by using the Oracle Data Guard command-line interface (<code>dgmgrl</code>) and has not been updated in Geographic Edition.

TABLE 2-1 Status and Status Messages of an Online Oracle Data Guard Replication Resource Group
 (Continued)

Status	Status Message	Possible Causes
Faulted	Database <i>database</i> does not exist in the configured Oracle Data Guard database list " <i>List-of-databases</i> "	The database has been deleted from the Oracle Data Guard Broker configuration using the Oracle Data Guard command-line interface (<i>dgmgrl</i>).
Faulted	Oracle errors " <i>List-of-ORA-xxxxx-errors</i> " were found in the Oracle Data Guard broker (<i>dgmgrl</i>) output when connecting by using " <i>connect-string</i> "	
Faulted	Role " <i>role</i> " given for database <i>database</i> does not match role " <i>role</i> " configured for Oracle Data Guard	The database might have been changed from a physical standby to a logical standby, or vice versa.
Unknown	Unexpected error - <i>unexpected-error</i>	
Unknown	Oracle Data Guard broker (<i>dgmgrl connect-string</i>) did not complete a response to the command " <i>command-string</i> " within " <i>number</i> " seconds and was timed out.	The Oracle Data Guard command-line interface (<i>dgmgrl</i>) did not respond to the show configuration command within the specified time, or Oracle Data Guard Broker was busy performing a health check during this period.
Unknown	Password or connect name (<i>connect-string</i>) for remote cluster is incorrect	The <i>sysdba_username</i> , <i>sysdba_password</i> , <i>local_db_service_name</i> , or <i>remote_db_service_name</i> parameter does not match the information that is maintained by the Geographic Edition software.
Unknown	File <i>filename</i> does not exist	A temporary internal file that is used by the Oracle Data Guard module was deleted before it could be read.
Degraded	Program <i>program-name</i> failed to read the Cluster Configuration Repository (CCR)	One of the programs that is used to retrieve information from the CCR failed.
Degraded	Failed to get password for <i>sysdba</i> user name for Oracle Data Guard configuration <i>ODGConfigurationName</i> in protection group <i>ODGprotectiongroupname</i>	The field for the <i>sysdba_password</i> was not found in the Cluster Configuration Repository (CCR) or was longer than expected.

TABLE 2-1 Status and Status Messages of an Online Oracle Data Guard Replication Resource Group
(Continued)

Status	Status Message	Possible Causes
Degraded	Local cluster <i>cluster-name</i> is not primary for Oracle Data Guard configuration <i>ODGConfigurationName</i>	A switchover or failover has been performed in Oracle Data Guard Broker by using a command in the Oracle Data Guard command-line interface (<i>dgmgrl</i>), and the Geographic Edition configuration has not been updated.
Degraded	Oracle Data Guard configuration name <i>ODGConfigurationName</i> found does not match <i>ODGConfigurationName</i>	
Degraded	Database <i>database-name</i> is in the disabled state	A database has been disabled in the Oracle Data Guard Broker using a command in the Oracle Data Guard command-line interface (<i>dgmgrl</i>), and the Geographic Edition configuration has not been updated.
Degraded	Oracle Data Guard configuration <i>ODGConfigurationName</i> is disabled on cluster <i>cluster-name</i>	The standby database in the Oracle Data Guard Broker configuration has been disabled by using a command in the Oracle Data Guard command-line interface (<i>dgmgrl</i>), and the Geographic Edition configuration has not been updated.
Degraded	Oracle Data Guard configuration <i>ODGConfigurationName</i> is disabled	The Oracle Data Guard Broker configuration has been disabled by using a command in the Oracle Data Guard command-line interface (<i>dgmgrl</i>), and the Geographic Edition configuration has not been updated.
Online	Online or replicating in <i>replication-mode</i> mode	

For more information about the `cl resource` command, see the `cl resource(1CL)` man page.

Migrating Services That Use Oracle Data Guard Data Replication

This chapter provides information about migrating services for maintenance or as a result of cluster failure.

This chapter covers the following topics:

- “Detecting Cluster Failure on a System That Uses Oracle Data Guard Data Replication” on page 71
- “Migrating Services That Use Oracle Data Guard With a Switchover” on page 72
- “Forcing a Takeover on Systems That Use Oracle Data Guard” on page 75
- “Recovering Oracle Data Guard Data After a Takeover” on page 78
- “Recovering From an Oracle Data Guard Data Replication Error” on page 88

Detecting Cluster Failure on a System That Uses Oracle Data Guard Data Replication

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

Detecting Primary Cluster Failure

When the primary cluster for a given protection group fails, the standby cluster in the partnership detects the failure. If the cluster that fails is a member of more than one partnership, multiple failure detections might occur.

The following actions occur when the overall state of a protection group changes to the Unknown state:

- Heartbeat failure is detected by a partner cluster.

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

You set this query interval by setting the `Query_interval` property of the heartbeat. If the heartbeat still fails after four attempts due to the `Query_interval` that you configured (three retries and one emergency-mode probing), a `heartbeat-lost` event is generated and logged in the system log. When you specify the default interval, the emergency-mode retry behavior might delay the notification of heartbeat-loss for about nine minutes. Messages are displayed in the GUI and in the output of the `geoadm status` command.

For more information about logging, see “[Viewing the Geographic Edition Log Messages](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Detecting Failure of the Standby Cluster

When a standby cluster for a given protection group fails, a cluster in the same partnership detects the failure. If the cluster that failed is a member of more than one partnership, multiple failure detections might occur.

During failure detection, the following actions occur:

- Heartbeat failure is detected by a partner cluster.
- The heartbeat is activated in emergency mode to verify that the standby cluster failed.
- The cluster notifies the administrator by issuing messages. The system detects all protection groups for which the cluster that failed was acting as standby. The state of these protection groups is set to the Unknown state.

Migrating Services That Use Oracle Data Guard With a Switchover

You perform a switchover of an Oracle Data Guard protection group when you want to migrate services to the partner cluster in an orderly fashion. A switchover includes the following operations:

- Application services are unmanaged on the former primary cluster `cluster-paris`.
For a reminder of which cluster is `cluster-paris`, see “[Example Geographic Edition Cluster Configuration](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
- The data replication role is reversed and now continues to run from the new primary, `cluster-newyork`, to the former primary, `cluster-paris`.

- Application services and the Oracle shadow RAC server proxy resource groups are brought online on the new primary cluster `cluster-newyork`.

This section provides the following information:

- [“How to Switch Over an Oracle Data Guard Protection Group From the Primary to the Standby Cluster” on page 73](#)
- [“Actions Performed by the Geographic Edition Software During a Switchover” on page 74](#)

▼ How to Switch Over an Oracle Data Guard Protection Group From the Primary to the Standby Cluster

Before You Begin For a switchover to occur, data replication must be active between the primary cluster and the standby cluster, that is, the Oracle Data Guard Broker configuration is enabled. Additionally, 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.

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

- Geographic Edition software is running on the both clusters.
- The standby cluster is a member of a partnership.
- Both cluster partners can be reached.
- The overall state of the protection group is set to `OK`.

1 Log in to a cluster node.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see [“Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*](#).

2 Initiate the switchover.

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

```
phys-node-n# geopg switchover [-f] -m newprimarycluster protectiongroupname
```

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

`-m newprimarycluster` Specifies the name of the cluster that is to be the primary cluster for the protection group.

`protectiongroupname` Specifies the name of the protection group.

Example 3-1 Forcing a Switchover From the Primary to the Standby Cluster

This example shows how to perform a switchover to the standby cluster.

```
phys-paris-1# geopg switchover -f -m cluster-newyork sales-pg
```

Actions Performed by the Geographic Edition Software During a Switchover

When you run the `geopg switchover` command, the software confirms that the primary cluster does indeed hold the primary database. 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. If the configuration is healthy, the software performs the following actions on the original primary cluster:

- Takes the application resource groups offline and places them in the Unmanaged state
- Performs a “`switchover to standby-database-name`” command for each Oracle Data Guard Broker configuration in the protection group

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

- Runs the script that is defined in the `RoleChange_ActionCmd` property
- Brings all the Oracle shadow RAC server proxy resource groups and any other application resource groups online

If the command completes successfully, the standby cluster, `cluster-newyork`, becomes the new primary cluster for the protection group. The original primary cluster, `cluster-paris`, becomes the new standby cluster. 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. The Oracle shadow RAC server proxy resource group and any other application resource groups are online on the new primary cluster. Data replication from the new primary cluster to the new standby cluster begins.

This 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 [“How to Validate an Oracle Data Guard Protection Group”](#) on page 45.

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 [“How to Resynchronize an Oracle Data Guard Protection Group”](#) on page 65.

Forcing a Takeover on Systems That Use Oracle Data Guard

You perform a takeover when applications need to be brought online on the standby cluster, regardless of whether the data is completely consistent between the primary database and the standby database. In this section, it is assumed that the protection group has been started.

The following operations occur after you initiate a takeover:

- If the former primary cluster, `cluster-paris`, can be reached and the protection group is not locked for notification handling or some other reason, the protection group is deactivated.

For a reminder of which cluster is `cluster-paris`, see [“Example Geographic Edition Cluster Configuration”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

- Databases replicated in the Oracle Data Guard Broker configurations, which are present in the protection group that is being taken over from the former primary cluster `cluster-paris`, are taken over by the new primary cluster `cluster-newyork`.

Note – This data might not be consistent with the original databases. Data replication from the new primary cluster, `cluster-newyork`, to the former primary cluster, `cluster-paris`, is stopped.

- The protection group is activated without data replication enabled. The former primary databases in each of the Oracle Data Guard Broker configurations that are taken over are placed in a disabled, recovery required, state.

For details about the possible conditions of the primary and standby clusters before and after a takeover, see [Appendix C, “Takeover Postconditions,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

This section provides the following information:

- [“How to Force Immediate Takeover of Oracle Data Guard Services by a Standby Cluster”](#) on page 76

- [“Actions Performed by the Geographic Edition Software During a Takeover”](#) on page 77

▼ How to Force Immediate Takeover of Oracle Data Guard Services by a Standby Cluster

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

- Geographic Edition software is up and running on the cluster.
- The cluster is a member of a partnership.
- The Configuration status of the protection group is set to OK on the standby cluster.

1 Log in to a node in the standby cluster.

To complete this step, you need to be assigned the Geo Management RBAC rights profile. For more information about RBAC, see [“Geographic Edition Software and RBAC”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Initiate the takeover.

```
phys-node-n# geopg takeover [-f] protectiongroupname
```

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

protectiongroupname Specifies the name of the protection group.

Example 3-2 Forcing a Takeover by a Standby Cluster

This example shows how to force the takeover of sales-pg by the standby cluster cluster-newyork.

The node phys-newyork-1 is the first node of the standby cluster. For a reminder of which node is phys-newyork-1, see [“Example Geographic Edition Cluster Configuration”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

```
phys-newyork-1# geopg takeover -f sales-pg
```

Next Steps For information about the state of the primary and the standby clusters after a takeover, see [Appendix C, “Takeover Postconditions,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Actions Performed by the Geographic Edition Software During a Takeover

When you run the `geopg takeover` command, the software confirms that databases in the Oracle Data Guard Broker configuration on the standby cluster, that is, the future primary, are enabled (as you cannot perform a takeover to a disabled database). The software also confirms that the Oracle Data Guard command-line interface `show configuration` command either shows a `SUCCESS` state or is busy performing a health check (`ORA-16610`). If the `show configuration` command returns any other Oracle error code, the takeover fails.

If the original primary cluster, `cluster-paris`, can be reached, the software takes the application resource groups offline and places them in an `Unmanaged` state.

On the original standby cluster, `cluster-newyork`, the software performs the following operations:

- Runs the Oracle Data Guard command line interface `failover to standby-database-name immediate` command
- Runs the script that is specified by the `RoleChange_ActionCmd` property
- If the protection group was active on the original standby cluster before the takeover, brings all Oracle shadow RAC server proxy resource groups and application resource groups online

If the command completes successfully, the standby cluster, `cluster-newyork`, becomes the new primary cluster for the protection group. 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. The Oracle shadow RAC server proxy resource group and any other application resource group are online on the new primary cluster. If the original primary cluster can be reached, it becomes the new standby cluster of the protection group. Replication of all databases that are associated with the Oracle Data Guard Broker configurations of the protection group are stopped.



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

If a previous operation fails, this command returns an error. Use the `geoadm status` command to view the status of each component. For example, the `Configuration` status of the protection group might be set to an `Error` state, 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 the Error state, revalidate the protection group by using the procedures that are described in [“How to Validate an Oracle Data Guard Protection Group”](#) on page 45.

If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures described in [“How to Resynchronize an Oracle Data Guard Protection Group”](#) on page 65.

Recovering Oracle Data Guard Data After a Takeover

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

Geographic Edition software supports the following two kinds of failback:

- **Failback switchover.** During a failback switchover, applications are brought online again on the original primary cluster, `cluster-paris`, after the primary cluster data has been resynchronized with the data on the standby cluster `cluster-newyork`.
For a reminder of which clusters are `cluster-paris` and `cluster-newyork`, see [“Example Geographic Edition Cluster Configuration”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
- **Failback takeover.** During a failback takeover, applications are brought online again on the original primary cluster and use the current data on the primary cluster. Any updates that occurred on the standby cluster are discarded.

If you want to leave the new primary, `cluster-newyork`, as the primary cluster and the original primary cluster, `cluster-paris`, as the standby cluster after the original primary cluster starts again, you can resynchronize and revalidate the protection group configuration. You can resynchronize and revalidate the protection group without performing a switchover or takeover.

This section describes how to perform the following procedures:

- [“How to Resynchronize and Revalidate the Protection Group Configuration”](#) on page 79
- [“How to Perform a Failback Switchover on a System That Uses Oracle Data Guard Replication”](#) on page 81
- [“How to Perform a Failback Takeover on a System That Uses Oracle Data Guard Replication”](#) on page 85

▼ How to Resynchronize and Revalidate the Protection Group Configuration

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

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

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

- 1 **If the original primary cluster, `cluster-paris`, has been down, confirm that the cluster is booted and that the Geographic Edition infrastructure is enabled on the cluster.**

For more information about booting a cluster, see “[Booting a Cluster](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

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

- a. **On `cluster-paris`, deactivate the protection group on the local cluster.**

```
phys-paris-1# geogg stop -e local protectiongroupname
-e local
```

Specifies the scope of the command.

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

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

```
protectiongroupname
```

Specifies the name of the protection group.

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

If you deactivate the protection group, the application resource groups are no longer managed, clearing the `Error` state.

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

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

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

For more information about synchronizing partnerships, see “[Resynchronizing a Partnership](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

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

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

For more information about synchronizing protection groups, see “[Resynchronizing an Oracle Data Guard Protection Group](#)” on page 65.

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

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

For more information, see “[How to Validate an Oracle Data Guard Protection Group](#)” on page 45.

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

When you activate a protection group, the protection group's application resource groups are also brought online.

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

```
-e global
```

Specifies the scope of the command.

By specifying a `global` scope, the command operates on both clusters where the protection group is located.

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

```
protectiongroupname
```

Specifies the name of the protection group.



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

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

For more information about the `geopg start` command, see [“How to Activate an Oracle Data Guard Protection Group” on page 60](#).

5 Confirm that all data is synchronized.

a. Confirm that the state of the protection group on `cluster-newyork` is OK.

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

Refer to the Protection Group section of the output.

b. Confirm that all resources in the replication resource group, `ODGprotectiongroupname-odg-rep-rg`, report a status of OK.

```
phys-newyork-1# clresource status -v ODGprotectiongroupname-odg-rep-rs
```

▼ How to Perform a Failback Switchover on a System That Uses Oracle Data Guard Replication

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

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

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

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

1 If the original primary cluster, `cluster-paris`, failed, confirm that the cluster is restarted and that the Geographic Edition infrastructure is enabled on the cluster.

For more information about restarting a cluster, see [“Booting a Cluster” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*](#).

2 Recover and restore the failed Oracle Data Guard primary database as the new standby.

Refer to the [Oracle documentation \(http://download.oracle.com/docs/cd/B19306_01/server.102/b14239/scenarios.htm#i1049997\)](http://download.oracle.com/docs/cd/B19306_01/server.102/b14239/scenarios.htm#i1049997), which describes how to perform this step.

3 Ensure that the original primary cluster, `cLuster-paris`, is working correctly as part of the Oracle Data Guard configuration.

```
oracle (phys-paris-1)$ dgmgrl sys/sysdba_password@sales-svc
DGMGRL> show configuration;
```

If the original primary cluster, `cLuster-paris`, is working correctly, the show configuration command displays the SUCCESS state.

If the original primary cluster was down at the point of failure, it is marked as a deactivated primary cluster. If the original primary cluster was up at the point of failure, it is marked as a deactivated secondary.

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

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

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

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

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

For more information about synchronizing partnerships, see “[Resynchronizing a Partnership](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

b. Determine whether the protection group on the original primary cluster, `cLuster-paris`, is active.

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

c. If the protection group on the original primary cluster is active, stop the protection group.

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

Specifies the scope of the command.

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

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

protectiongroupname

Specifies the name of the protection group.

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

If you deactivate the protection group, the application resource groups are no longer managed, clearing the `Error` state.

d. Verify that the protection group is stopped.

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

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

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

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

For more information about synchronizing protection groups, see [“Resynchronizing an Oracle Data Guard Protection Group” on page 65](#).

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

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

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

For more information, see [“How to Validate an Oracle Data Guard Protection Group” on page 45](#).

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

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

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

```
-e global
```

Specifies the scope of the command.

By specifying a `global` scope, the command operates on both clusters where the protection group is located.

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

protectiongroupname

Specifies the name of the protection group.

7 Confirm that the data is completely synchronized.

a. Confirm that the state of the protection group on `cluster-newyork` is OK.

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

Refer to the Protection Group section of the output.

b. Confirm that all resources in the replication resource group, `ODGprotectiongroupname-odg-rep-rg`, report a status of OK.

```
phys-newyork-1# clresource status -v ODGprotectiongroupname-odg-rep-rs
```

8 On both partner clusters, ensure that the protection group is activated.

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

```
...
```

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

```
...
```

9 For each protection group, on either cluster, perform a switchover from `cluster-newyork` to `cluster-paris`.

```
phys-node-n# geopg switchover [-f] -m cluster-paris protectiongroupname
```

For more information, see [“How to Switch Over an Oracle Data Guard Protection Group From the Primary to the Standby Cluster”](#) on page 73.

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

10 Ensure that the switchover was performed successfully.

```
phys-node-n# geoadm status
```

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

11 Check the runtime status of the application resource group and data replication for each Oracle Data Guard protection group.

```
phys-node-n# clresourcegroup status -v resourcegroupname
# clresource status -v ODGConfigurationName-odg-rep-rs
```

Refer to the Status and Status Message fields that are presented for the Oracle Data Guard Broker configuration that you want to check. For more information about these fields, see [Table 2–1](#).

For more information about the runtime status of data replication, see “[Checking the Runtime Status of Oracle Data Guard Data Replication](#)” on page 66.

▼ How to Perform a Failback Takeover on a System That Uses Oracle Data Guard Replication

Follow this procedure to restart an application on the original primary cluster, `cluster-paris`, and to use the current data on the original primary cluster.

Note – Any updates that occurred on the standby cluster, `cluster-newyork`, while it was acting as primary are discarded.

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

Note – Conditionally, you can resume using the data on the original primary `cluster-paris`. However, you must not have replicated data from the new primary, `cluster-newyork`, to the original primary cluster, `cluster-paris`, at any point after the takeover operation on `cluster-newyork`.

Before You Begin Before you begin the failback takeover procedure, the clusters must have the following roles:

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

- 1 **If the original primary cluster, `cluster-paris`, failed, confirm that the cluster is restarted and that the Geographic Edition infrastructure is enabled on the cluster.**

For more information about restarting a cluster, see “[Booting a Cluster](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

- 2 **Recover and revert the new Oracle Data Guard primary database to a standby for the original primary to a point in time before the original primary failed.**

Refer to the [Oracle documentation \(http://download.oracle.com/docs/cd/B19306_01/server.102/b14239/scenarios.htm#i1049997\)](http://download.oracle.com/docs/cd/B19306_01/server.102/b14239/scenarios.htm#i1049997), which describes how to perform this step.

Note – You might need to use the `dgmgrl` command to remove and recreate the Oracle Data Guard Broker configuration.

3 Ensure that the original primary cluster, `cluster-paris`, is again working correctly as the primary as part of the Oracle Data Guard configuration.

```
oracle (phys-paris-1)$ dgmgrl sys/sysdba_password@sales-svc
DGMGRL> show configuration;
```

If the original primary cluster, `cluster-paris`, is working correctly, the `show configuration` command displays the `SUCCESS` state.

If the original primary cluster was up at the point of failure, it is marked as a deactivated secondary cluster. Also, the original standby cluster is marked as an activated primary.

If the original primary cluster was down at the point of failure, it is marked as a deactivated primary cluster. Also, the original standby cluster is marked as an activated primary.

4 Determine whether the original primary cluster, `cluster-paris`, was up or down at the point of failure.

- **If the original primary cluster, `cluster-paris`, was down at the point of failure, update the original standby cluster, `cluster-newyork`, to a secondary.**

- a. **On the original standby cluster, that is, the cluster that has become the new primary cluster, stop the protection group.**

```
phys-newyork-1# geopg stop -e local protectiongroupname
```

- b. **On the original standby cluster, that is, the cluster that has become the new primary cluster, update the protection group.**

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

The roles are now correct, but both clusters are marked as deactivated.

For more information about synchronizing protection groups, see [“How to Resynchronize an Oracle Data Guard Protection Group” on page 65](#).

- c. **On `cluster-paris` and on `cluster-newyork`, locally validate the configuration for each protection group.**

Ensure that the protection group is not in an `Error` state. You cannot start a protection group when the protection group is in an `Error` state.

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

For more information, see [“How to Validate an Oracle Data Guard Protection Group” on page 45](#).

- d. From any node in either cluster, globally activate the protection group on both clusters.

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

Once the protection groups are activated on both clusters, you have successfully performed the failback takeover.

- If the original primary cluster, `cluster-paris`, was up at the point of failure, determine the status of the secondary (that is, the original primary) configuration.

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

- If the Configuration status is set to OK, synchronize the configurations.

- a. Initiate a takeover for each protection group on the original primary `cluster-paris`.

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

- b. If the configuration for the original standby cluster, `cluster-newyork`, is marked as Error, validate the configuration for each protection group.

```
cluster-newyork# geopg validate protectiongroupname
```

For more information, see [“How to Validate an Oracle Data Guard Protection Group” on page 45](#).

- c. Globally activate the protection groups on both clusters.

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

Once the protection groups are activated on both clusters, you have successfully performed the failback takeover.

- If the Configuration status is set to Error, resolve the problem.

- a. Deactivate the secondary (that is, the original primary) configuration that is in the Error state.

```
phys-newyork-1# geopg stop -e local protectiongroupname
```

- b. Force a takeover to make the secondary configuration a primary configuration again and to match the underlying Oracle `dgmgrl` configuration.

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

- c. On both the `cluster-paris` and on the `cluster-newyork` clusters, locally validate the configuration for each protection group.

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

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

For more information, see [“How to Validate an Oracle Data Guard Protection Group” on page 45](#).

- d. **From any node in either cluster, globally activate the protection group on both clusters.**

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

Once the protection groups are activated on both clusters, you have successfully performed the failback takeover.

Recovering From an Oracle Data Guard Data Replication Error

When an error occurs at the data replication level, the error is reflected in the status of the resource in the replication resource group of the relevant the Oracle Data Guard Broker configuration.

For example, suppose that Oracle Data Guard Broker configuration sales - pg, which contains the replicated database sales, is changed from protection mode MaxAvailability to MaxPerformance. The state changes for FAULTED are reflected in the following resource status:

```
Resource Status = "FAULTED"
Resource status message = "FAULTED - Protection mode "MaxAvailability" given
for local database sales does not match configured value "MaxPerformance"
```

Note – The Resource State remains On Line because the probe is still running correctly.

Because the resource status has changed, the protection group status also changes. In this case, the local Data Replication state, the Protection Group state on the local cluster, and the overall Protection Group state all become Error.

To recover from an error state, perform the following procedure.

▼ How to Recover From a Data Replication Error

- 1 **Use the procedures in the Oracle Data Guard documentation to determine the causes of the FAULTED state.**
- 2 **Recover from the faulted state by following the Oracle Data Guard procedures.**

If the recovery procedures change the state of the Oracle Data Guard Broker configuration, this state is automatically detected by the resource and is reported as a new protection group state. If the replication mode does not match the Geographic Edition settings, type:

```
phys-paris-1# geopg modify-replication-component \
-p replication_mode=New-protection-mode \
ODGConfigurationName protectiongroupname
```

3 Revalidate the protection group configuration.

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

where *protectiongroupname* specifies the name of the Oracle Data Guard protection group.

4 Review the status of the protection group configuration.

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

where *protectiongroupname* specifies the name of the Oracle Data Guard protection group.

Geographic Edition Properties for Oracle Data Guard Broker Configurations

This appendix describes the properties for Geographic Edition data replications that use Oracle Data Guard.

Oracle Data Guard Broker Configuration Properties

This section describes the Oracle Data Guard Broker configuration properties that the Geographic Edition software defines.

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.

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.

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. An Oracle shadow RAC server proxy resource group shadows the real resource group. If you want, add the shadow to the protection group application resource group list.

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.

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.

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. An Oracle shadow RAC server proxy resource group shadows the real resource group. If you want, add the shadow to the protection group application resource group list.

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.

Valid values to which you set this property include `maximumAvailability`, `maximumPerformance`, and `maximumProtection`.

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.

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.

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. Use this property to monitor and manage the Oracle Data Guard Broker configurations.

Category: Required

Default: None

Tunable: Any time

Index

A

- activating protection groups, 60–62
- administering
 - data replication with Oracle Data Guard, 13–34, 35–69
 - Oracle Data Guard Broker configurations, 52–58
- application resource groups
 - administering, 48–52
 - creating, 49–51
 - removing, 51–52

C

- configuration summary, 14
- configuring
 - Oracle Data Guard Broker configurations, 32–34
 - Oracle Data Guard configuration, 19–21
 - Oracle Data Guard software, 17–19
 - protection groups, 42–43
- creating
 - application resource group, 49–51
 - protection groups, 42–43
 - replication Oracle Data Guard Broker configurations, 53–55

D

- data recovery, 78–88
 - failback switchover, 81–85
 - failback takeover, 85–88

- database standby types, 13
- deactivating protection groups, 63–64
- deleting
 - application resource group, 51–52
 - protection groups, 47–48
 - replication Oracle Data Guard Broker configuration, 57–58
- detecting failure, 71–72

F

- failback switchover, 81–85
- failback takeover, 85–88
- failure
 - detecting, 71–72
 - primary cluster, 71–72
 - standby cluster, 72

L

- local_database_name, 54, 91
- local_db_service_name, 54, 91, 92
- local_rac_proxy_svr_rg_name, 54, 91
- Logical standby, restriction, 13

M

- migrating services, 71–89
 - data recovery after, 78–88
 - with a switchover, 72–75

migrating services (*Continued*)

- with a takeover, 75–78

modifying

- protection groups, 44
- replication Oracle Data Guard Broker configurations, 56–57

O

Oracle Data Guard

- administering data replication with, 13–34, 35–69
- configuring software, 17–19
- detecting failure, 71–72
- initial software configuration, 16–34
- migrating services that use, 71–89
- properties of, 91–93
- properties of
 - local_database_name, 54, 91
 - local_db_service_name, 54, 91
 - local_rac_proxy_svr_rg_name, 54, 91
 - remote_database_name, 54, 92
 - remote_db_service_name, 54, 92
 - remote_rac_proxy_svr_rg_name, 54, 92
 - replication_mode, 54, 92
 - standby_type, 54, 92
 - sysdba_password, 54, 93
 - sysdba_username, 54, 93
- replication resource groups, 16
- runtime status, 66–69
 - overall, 66–67
- shadow resource groups, 15–16

Oracle Data Guard Broker configurations

- adding to protection group, 53–55
- administering, 52–58
- configuring, 32–34
- modifying, 56–57
- removing, 57–58

Oracle Data Guard configuration

- configuring, 19–21
- setting up primary database, 19–21

P

- partner clusters, 17

- partnerships, 17

- Physical standby, 13

primary cluster

- data recovery, 78–88
- failure detection, 71–72
- switchover, 72–75
- takeover, 75–78

properties

- Oracle Data Guard, 91–93

- local_database_name, 54, 91
- local_db_service_name, 54, 91
- local_rac_proxy_svr_rg_name, 54, 91
- remote_database_name, 54, 92
- remote_db_service_name, 54, 92
- remote_rac_proxy_svr_rg_name, 54, 92
- replication_mode, 54, 92
- standby_type, 54, 92
- sysdba_password, 54, 93
- sysdba_username, 54, 93

protection groups

- activating, 60–62
- adding application resource group to, 49–51
- adding Oracle Data Guard Broker configurations to, 53–55
- adding shadow RAC server proxy resource group to, 49–51
- configuring, 42–43
- creating, 42–43
- creation strategies, 35–42
- deactivating, 63–64
- deleting, 47–48
- modifying, 44
- modifying Oracle Data Guard Broker configurations for, 56–57
- removing application resource group, 51–52
- removing Oracle Data Guard Broker configuration from, 57–58
- removing shadow RAC server proxy resource group, 51–52
- replicating configuration of, 58–60
- resynchronizing, 65
- validating, 45

R

recovery

See data recovery

 from replication error, 88–89

remote_database_name, 54, 92

remote_db_service_name, 54

remote_rac_proxy_svr_rg_name, 54, 92

replication

 adding replication component, 53–55

 initial configuration of, 16–34

 migrating services, 71–89

 modifying Oracle Data Guard Broker

 configurations, 56–57

 Oracle Data Guard, 13–34, 35–69

 protection group configuration, 58–60

 recovering from errors, 88–89

 removing Oracle Data Guard Broker

 configuration, 57–58

 resource groups, 16

 runtime status details, 67–69

 runtime status overview, 66–67

replication_mode, 54, 92

replication resource groups and status, 67–69

resource groups

 application, 48–52

 replication, 16

 shadow, 15–16

resynchronizing protection groups, 65

runtime status

 replication, 66–69

 state and status messages, 67–69

S

shadow resource groups, 15–16

standby cluster

 failure detection, 72

 switchover, 72–75

 takeover, 75–78

standby_type, 54, 92

switchover, 72–75

 actions performed during, 74–75

 primary to standby, 73–74

sysdba_password, 54, 93

sysdba_username, 54, 93

T

takeover, 75–78

 actions performed during, 77–78

 data recovery after, 78–88

 failback switchover, 81–85

 failback takeover, 85–88

 how to force, 76

V

validating protection groups, 45

