

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

ORACLE®

Part No: E39669
September 2014, E39669-02

Copyright © 2008, 2014, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

Copyright © 2008, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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

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

Contents

Using This Documentation	13
1 Replicating Data With Oracle Data Guard Software	15
Replicating Data in an Oracle Data Guard Protection Group (Task Maps)	16
Overview of Oracle Data Guard Data Replication	17
Oracle Data Guard Shadow Resource Groups	17
Oracle Data Guard Replication Resource Groups	19
Initially Configuring Oracle Data Guard Replication	19
Oracle Data Guard Broker Configurations	20
▼ How to Set Up Your Primary Database	21
▼ How to Configure the Primary Database Listener and Naming Service	24
▼ How to Prepare Your Standby Database	27
▼ How to Configure the Standby Database Listener and Naming Service	30
▼ How to Start and Recover Your Standby Database	33
▼ How to Verify That Your Configuration Is Working Correctly	34
▼ How to Complete Configuring and Integrating Your Standby Oracle Database	34
▼ How to Create and Enable an Oracle Data Guard Broker Configuration	35
▼ How to Configure Oracle Solaris Cluster Manageability Resources for the Primary and Standby Databases	38
2 Creating Oracle Data Guard Protection Groups	39
Creating and Validating an Oracle Data Guard Protection Group	39
▼ How to Create and Configure an Oracle Data Guard Protection Group	39
Adding an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group	41
Adding an Application Resource Group to an Oracle Data Guard Protection Group	45
Validating an Oracle Data Guard Protection Group	48
3 Configuring Oracle Data Guard Protection Groups	53

Working With Oracle Data Guard Protection Groups	53
Overview of Administering Protection Groups	54
▼ How to Administer an Oracle Data Guard Protection Group (Example)	54
Administering Oracle Data Guard Application Resource Groups	61
▼ How to Delete an Application Resource Group From an Oracle Data Guard Protection Group	61
Administering Oracle Data Guard Broker Configurations	63
Oracle Data Guard Broker Configuration Properties	63
▼ How to Modify Geographic Edition Properties of an Oracle Data Guard Broker Configuration	64
▼ How to Delete an Oracle Data Guard Broker Configuration From an Oracle Data Guard Protection Group	65
Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster	66
▼ How to Replicate the Oracle Data Guard Protection Group Configuration to a Partner Cluster	66
Checking the Runtime Status of Oracle Data Guard Data Replication	69
Displaying an Oracle Data Guard Runtime Status Overview	69
Displaying a Detailed Oracle Data Guard Runtime Status	70
Enabling Debugging for Runtime Messages	73
4 Migrating Services That Use Oracle Data Guard Data Replication	75
Detecting Cluster Failure on a System That Uses Oracle Data Guard Data Replication	75
Detecting Primary Cluster Failure	75
Detecting Failure of the Standby Cluster	76
Migrating Services That Use Oracle Data Guard With a Switchover	76
Actions Performed by the Geographic Edition Software During a Switchover	77
Forcing a Takeover on Systems That Use Oracle Data Guard	78
Actions Performed by the Geographic Edition Software During a Takeover	79
Recovering Oracle Data Guard Data After a Takeover	80
▼ How to Perform a Failback Switchover or Failback Takeover	81
Recovering From an Oracle Data Guard Data Replication Error	85
▼ How to Recover From a Data Replication Error	85
A Geographic Edition Properties for Oracle Data Guard Broker Configurations	87
Oracle Data Guard Broker Replication Component Properties	87

Index 91

Tables

TABLE 1-1	Configuration Tasks for Oracle Data Guard Data Replication	16
TABLE 1-2	Administration Tasks for Oracle Data Guard Data Replication	17
TABLE 3-1	Status and Status Messages of an Online Oracle Data Guard Replication Resource Group	71

Examples

- EXAMPLE 2-1** Adding an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group 45
- EXAMPLE 2-2** Adding an Application Resource Group to an Oracle Data Guard Protection Group 48
- EXAMPLE 2-3** Validating the Configuration of a Protection Group 51
- EXAMPLE 3-1** Deleting an Application Resource Group From a Protection Group 63
- EXAMPLE 3-2** Deleting an Oracle Data Guard Broker Configuration From an Oracle Data Guard Protection Group 66
- EXAMPLE 3-3** Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster 69

Using This Documentation

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

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

Product Documentation Library

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

Access to Oracle Support

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

Documentation Accessibility

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

Feedback

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

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 \(Task Maps\)” on page 16](#)
- [“Overview of Oracle Data Guard Data Replication” on page 17](#)
- [“Initially Configuring Oracle Data Guard Replication” on page 19](#)

Geographic Edition software supports the use of Oracle Data Guard for data replication when used with HA for Oracle, HA for Oracle External Proxy, or Oracle Real Application Clusters software.

Beginning with Oracle Solaris Cluster 4.2 SRU 1, Geographic Edition software also supports the use of Oracle Data Guard instances running on a system that is not under Oracle Solaris Cluster control, called a *remote node*. Replication can be performed between two databases that are both running on remote nodes, or performed between a database on a remote node and another database on a cluster node that is configured with Geographic Edition software.

Note - To manage an Oracle Data Guard database on a remote node, you must configure the HA for Oracle External Proxy resource for the remote database instance on the same cluster where you create the protection group.

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. If you plan to configure an Oracle Data Guard database on a remote node, also see [“Oracle Solaris Cluster Data Service for Oracle External Proxy Guide”](#).

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.



Caution - When using HA for Oracle for an Oracle Data Guard database instance, take special care to coordinate the switchover with HA for Oracle. HA for Oracle monitors the database and will take action when the database's state changes unexpectedly. Likewise, when Geographic Edition is used to remotely manage an Oracle Data Guard database instance, the remote node must not have HA for Oracle configured for this Oracle Data Guard instance.

When both HA for Oracle and Geographic Edition are used for the same Oracle Data Guard database instance on the same Oracle Solaris Cluster node, if you must switch over the Oracle Data Guard state directly, first unmonitor the HA for Oracle resource. After the switchover is done, ensure that the `Dataguard_role` and `Standby_mode` resource properties exactly match the current database settings locally, before you re-enable monitoring of the HA for Oracle resource.

Failure to take these precautions might result in HA for Oracle issuing database restarts and failing to stay online.

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 Maps)

The following tables summarizes the tasks for configuring and administering Oracle Data Guard data replication in a protection group. These procedures apply whether you configure Oracle Data Guard replication on a cluster node under Oracle Solaris Cluster control or on a remote node that is not under direct Oracle Solaris Cluster control.

TABLE 1-1 Configuration 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 Replication” on page 19.
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 39.
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 41.
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 46.
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 66.
Activate the protection group.	See “How to Activate a Protection Group” in “Oracle Solaris Cluster Geographic Edition System Administration Guide” .

Task	Description
Check the runtime status of replication.	See “ Checking the Runtime Status of Oracle Data Guard Data Replication ” on page 69.

TABLE 1-2 Administration Tasks for Oracle Data Guard Data Replication

Task	Description
Detect failure.	See “ Detecting Cluster Failure on a System That Uses Oracle Data Guard Data Replication ” on page 75.
Migrate services by using a switchover.	See “ Migrating Services That Use Oracle Data Guard With a Switchover ” on page 76.
Migrate services by using a takeover.	See “ Forcing a Takeover on Systems That Use Oracle Data Guard ” on page 78.
Recover data after forcing a takeover.	See “ Recovering Oracle Data Guard Data After a Takeover ” on page 80.

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 Availability Suite, *Hitachi TrueCopy* or *Universal Replicator*, and EMC SRDF.

Oracle Data Guard Shadow 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 shadow Oracle database-server resource group and resource. The shadow Oracle database-server resource group “shadows” the real Oracle database-server resource group that you created to manage and monitor the Oracle databases that are under the control of Oracle Solaris Cluster software.

The name of a shadow resource group conforms to the following format:

ODG-configuration-name-rac-proxy-svr-shadow-rg

This format applies regardless of which data service is configured for the Oracle Database software — Support for Oracle RAC, HA for Oracle External Proxy, or HA for Oracle.

For example, an Oracle RAC database in an Oracle Data Guard Broker configuration named `sales` has a shadow Oracle database-server 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`.

Similarly, an HA for Oracle database in an Oracle Data Guard Broker configuration named `inventory` that is controlled by the Oracle Data Guard software has a shadow Oracle database-server resource group named `inventory-rac-proxy-svr-shadow-rg`.

The requirements for consistent name construction are two-fold. First, this allows a shadow resource group to be added to a protection group even when one cluster uses Oracle RAC and the other uses HA for Oracle. The second reason is that this format is required for backward compatibility.

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

- **HA for Oracle or failover HA for Oracle External Proxy** – *ODG-configuration-name-oracle-svr-shadow-rs*
- **Oracle RAC or scalable HA for Oracle External Proxy** – *ODG-configuration-name-rac-proxy-svr-shadow-rs*

For more information about Oracle database-server resource groups, see [“Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide”](#), [“Oracle Solaris Cluster Data Service for Oracle External Proxy Guide”](#), and [“Oracle Solaris Cluster Data Service for Oracle Database Guide”](#).

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

Consequently, putting the real Oracle database-server resource group under Geographic Edition control would result in the Oracle database being shut down on the standby cluster. In contrast, the shadow Oracle database-server 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. In addition, putting the shadow resource groups under Geographic Edition control enables you to declare on the shadow resource-group affinities and other relationships with other resource groups. These other resource groups can then also be controlled by Geographic Edition.

The state of the shadow Oracle database-server resource group indicates whether the database that is monitored and controlled by the Oracle database-server 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:

- If the shadow Oracle database-server resource group is online, its cluster is the primary cluster for that Oracle database, which is represented by the resource in the actual Oracle database-server resource group.
- If the shadow Oracle database-server resource group is offline and unmanaged, its cluster is the standby for that Oracle database.

Furthermore, the status of the shadow Oracle database-server resource reflects both the status of the Oracle database-server 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:

ODG-protection-group-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 replication resource also checks the accessibility of the configuration through the remote service name, to ensure that the username and password information, or the Oracle wallet if used, is correctly configured.

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

ODG-Broker-configuration-name-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 Replication

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

Guard replication on a system that is not under Oracle Solaris Cluster control, referred to as a *remote node*.

This section provides the following information:

- [“Oracle Data Guard Broker Configurations” on page 20](#)
- [“How to Set Up Your Primary Database” on page 21](#)
- [“How to Configure the Primary Database Listener and Naming Service” on page 24](#)
- [“How to Prepare Your Standby Database” on page 27](#)
- [“How to Configure the Standby Database Listener and Naming Service” on page 30](#)
- [“How to Start and Recover Your Standby Database” on page 33](#)
- [“How to Verify That Your Configuration Is Working Correctly” on page 34](#)
- [“How to Complete Configuring and Integrating Your Standby Oracle Database” on page 34](#)
- [“How to Create and Enable an Oracle Data Guard Broker Configuration” on page 35](#)
- [“How to Configure Oracle Solaris Cluster Manageability Resources for the Primary and Standby Databases” on page 38](#)

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 `cluster-newyork` cluster. These names are held in one of the following locations:
 - The `tnsnames.ora` files in the `${ORACLE_HOME}/network/admin` directory of the nodes that host the replicated Oracle database
 - The Oracle naming service directory
- **The database standby type for the Oracle Data Guard Broker configuration**, which you set to either `logical`, `physical`, or `snapshot`.
- **The replication mode for the Oracle Data Guard Broker configuration**, which you set to `MaxPerformance`, `MaxAvailability`, or `MaxProtection`.
- **The names of the Oracle database-server resource groups that manage the Oracle databases on each cluster**. You configure these names by using the data service configuration wizard through the `clsetup` command. Alternatively, follow the instructions in [“Registering and Configuring HA for Oracle Database”](#) in [“Oracle Solaris Cluster Data](#)

Service for Oracle Database Guide” or Appendix D, “Command-Line Alternatives,” in “Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide”.

- If running the primary or secondary replicated Oracle database on a remote node, **the name of the HA for Oracle External Proxy resource and resource group on that remote node**. For configuration instructions, see “Registering and Configuring HA for Oracle External Proxy” in “Oracle Solaris Cluster Data Service for Oracle External Proxy Guide”.

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.

▼ How to Set Up Your Primary Database

Perform this task from cluster nodes where you want Oracle Data Guard to run. These can be nodes of an Oracle Solaris Cluster configuration or a system that is not under Oracle Solaris Cluster control.

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.

Note - Oracle Clusterware is a component of Oracle Grid Infrastructure.

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`.

Note - Although the following example procedure uses an Oracle 10g RAC database, the principles are the same for an Oracle 11g RAC database. This release of Geographic Edition software supports a minimum of Oracle version 11.2.0.3.

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 archive log;
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 database 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

Perform this task on each primary cluster node where you want Oracle Data Guard to run. These can be nodes of an Oracle Solaris Cluster configuration or a system that is not under Oracle Solaris Cluster control.

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 network 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 30](#). Also, add entries for the standby (`salesdr`) database instances that you create later when you modify the `pfile` 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-crs)(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

Perform this task from cluster nodes where you want Oracle Data Guard to run. These can be nodes of an Oracle Solaris Cluster configuration or a remote server that is not under Oracle Solaris Cluster control.

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 Database 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 an \${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

Perform this task on each node of the standby cluster where you want Oracle Data Guard to run. These can be nodes of an Oracle Solaris Cluster configuration or a system that is not under Oracle Solaris Cluster control.

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 24](#). 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_DGMGRL" 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 Oracle Database

1. **If you are using Oracle Clusterware, register the new database.**

Note - Oracle Clusterware is a component of Oracle Grid Infrastructure.

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

■ **For a single-instance database, run the following command:**

```
oracle (phys-newyork-1)$ srvctl add database -d salesdr \  
-r PHYSICAL_STANDBY -o $ORACLE_HOME -s open;
```

■ **For Oracle RAC, run the following commands:**

```
oracle (phys-newyork-1)$ srvctl add database -d salesdr \  
-r PHYSICAL_STANDBY -o $ORACLE_HOME -s open;
```

```

-r PHYSICAL_STANDBY -o $ORACLE_HOME -s open;
oracle (phys-newyork-1)$ 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. Enable Oracle Data Guard on both the primary and standby databases.

Perform the following commands 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
DGMGRL> create configuration mysales.com as primary
DGMGRL> 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
DGMGRL> 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)$ dgmgrl 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$ dgmgrl 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)$ dgmgrl 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$ dgmgrl 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)$ dgmgrl 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"

DGMGRL> 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"

```

Next Steps Go to [“How to Configure Oracle Solaris Cluster Manageability Resources for the Primary and Standby Databases”](#) on page 38.

▼ How to Configure Oracle Solaris Cluster Manageability Resources for the Primary and Standby Databases

This procedure integrates the primary and standby Oracle Data Guard databases with Oracle Solaris Cluster software. By configuring these resource groups and resources, you enable Oracle Solaris Cluster software to monitor and manage the primary and standby databases. These resource groups are then used by the Geographic Edition framework to monitor and manage the Oracle Data Guard configuration.

Perform this procedure on nodes of the Oracle Solaris Cluster configuration.

● **Configure the Oracle Solaris Cluster manageability resources.**

The Oracle Solaris Cluster resource groups that you create are then entered as properties when the Oracle Data Guard configuration is added to the protection group. The Geographic Edition framework uses these resource groups to monitor and manage the Oracle Data Guard configuration during protection group operations such as switchover and takeover.

You can use either the data service configuration wizard that is available through the `clsetup` utility or use Oracle Solaris Cluster maintenance commands. Follow procedures for the data service you use:

- **HA for Oracle** – [“How to Register and Configure HA for Oracle Database With Oracle Grid Infrastructure for a Cluster \(CLI\)”](#) in [“Oracle Solaris Cluster Data Service for Oracle Database Guide ”](#)
- **Oracle RAC** – [“How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate”](#) in [“Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide ”](#) or [“Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands”](#) in [“Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide ”](#).
- **HA for Oracle External Proxy** – [“Registering and Configuring HA for Oracle External Proxy”](#) in [“Oracle Solaris Cluster Data Service for Oracle External Proxy Guide ”](#)

Creating Oracle Data Guard Protection Groups

This chapter describes how to create an Oracle Data Guard protection group.

Creating and Validating an Oracle Data Guard Protection Group

This section covers the following topics:

- [“How to Create and Configure an Oracle Data Guard Protection Group” on page 39](#)
- [“Adding an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group” on page 41](#)
- [“Adding an Application Resource Group to an Oracle Data Guard Protection Group” on page 45](#)
- [“Validating an Oracle Data Guard Protection Group” on page 48](#)

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

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

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

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

Before You Begin Ensure that the following conditions are met:

- The Oracle Data Guard Broker configuration exists. Geographic Edition does *not* create the configuration for you.
- 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 66.

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

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

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

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

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

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

```
phys-node-n# geopg create -s partnership -d odg \  
-o local-role [-p property [-p...]] protection-group
```

`-s partnership`

Specifies the name of the partnership.

`-d odg`

Specifies that the protection group data is replicated by Oracle Data Guard software.

`-o local-role`

Specifies the role of this protection group on the local cluster as either primary or secondary.

`-p property-setting`

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.

protection-group

Specifies the name of the protection group.

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

Adding an Oracle Data Guard Broker Configuration to an Oracle Data Guard Protection Group

A protection group is the container for the replication component and 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 status of the data replication that corresponds to the database in the 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.

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

Before You Begin 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 primary and standby database systems.
- The Oracle database-server resource group and Oracle database-server resources that manage the Oracle database that is replicated by Oracle Data Guard exist on both the local and the partner cluster.
- If the cluster will remotely manage an Oracle database, each cluster node is installed and configured with HA for Oracle External Proxy.
- If you are using HA for Oracle External Proxy to manage replication of a remote Oracle database, all cluster nodes where HA for Oracle External Proxy is installed are also installed with the Oracle Database software client administrative kit. This software installs `$ORACLE_HOME/bin/sqlplus` and `$ORACLE_HOME/bin/dgmgrl`.
- If you are using HA for Oracle External Proxy, ensure that one of the following is true on all cluster nodes:
 - The `TNS_ADMIN` property of the HA for Oracle External Proxy resource is set to `$ORACLE_HOME/network/admin`.
 - The `/cluster/geo/odg/odg_configuration_name_config` file contains the entry `ORACLE_HOME=$ORACLE_HOME`.
- If you are using HA for Oracle External Proxy, ensure that you have granted the `sysdba` privilege to the user that the agent utilizes to monitor the database. The following SQL command grants this privilege:

```
$ sqlplus '/ as sysdba'  
SQL> grant sysdba to hauser;  
Grant succeeded.
```

The `hauser` is the name you set in the `dbuser` property of the corresponding `SUNW.oracle_external_proxy` resource. For more information, see [“Remote Database User”](#) in [“Oracle Solaris Cluster Data Service for Oracle External Proxy Guide”](#).

1. Ensure that the Oracle Data Guard Broker properties `BystandersFollowRoleChange` and `FAST_START FAILOVER` are set properly.

a. Set `BystandersFollowRoleChange` to `NONE`.

```
DGMGRL> edit configuration set property BystandersFollowRoleChange=NONE;
```

b. Ensure that `FAST_START FAILOVER` is disable.

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

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

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

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

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

3. For HA for Oracle, ensure that the `Standby_mode` and `Dataguard_role` extension properties of the `SUNW.oracle_server` resource match the current standby mode of the Oracle Data Guard Broker database.

Perform this step on one node of **each** partner cluster that runs HA for Oracle.

```
phys-newyork-n# clresource set -p Standby_mode=mode \  
-p Dataguard_role=role ora-db-rs  
phys-paris-n# clresource set -p Standby_mode=mode \  
-p Dataguard_role=role ora-db-rs
```

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

Note - You can also accomplish this step by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, highlight the protection group name, and in the Data Replication Components section click Add. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

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...] ODG-configuration-name protection-group
```

`-p property`

Specifies the properties of either the Oracle Data Guard Broker configuration, the Oracle database-server 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_oracle_svr_rg_name` – Name of the local Oracle database-server resource group that manages the local database in the Oracle Data Guard Broker configuration. This resource group can be configured for Oracle RAC, HA for Oracle, or HA for Oracle External Proxy.
- `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_oracle_svr_rg_name` – Name of the Oracle database-server resource group on the partner cluster that manages the remote database in the Oracle Data Guard Broker configuration. This resource group can be configured for Oracle RAC, HA for Oracle, or HA for Oracle External Proxy.
- `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.
If you use an Oracle wallet, you do not need to specify this password.
- `sysdba_username` – Name of an Oracle SYSDBA privileged database user who can perform the Oracle Data Guard Broker switchover and takeover operations.
If you use an Oracle wallet, you do not need to specify this password.

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”.

ODG-configuration-name

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

protection-group

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 2-1 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
```

Adding an Application Resource Group to an Oracle Data Guard Protection Group

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 database-server resource group is not added to the protection group. Instead, a shadow Oracle database-server resource group is added to represent this resource group.

You can add and remove the shadow Oracle database-server resource group to and from the protection group at any time without affecting the Oracle Data Guard data replication. This fact does not prevent you from adding other, non-Oracle database-server 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.

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

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

Before You Begin You can add an existing resource group, other than an Oracle database-server resource group containing an Oracle database-server resource, to the list of application resource groups for a protection group. If you do try to add an Oracle database-server 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.

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

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

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

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

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

2. Ensure that the `Auto_start_on_new_cluster` property of the resource group is set to `False`.

```
# clresourcegroup show -p Auto_start_on_new_cluster resource-group
```

If necessary, change the property value to `False`.

```
# clresourcegroup set -p Auto_start_on_new_cluster=False resource-group
```

3. **If the application resource group must have dependencies on resource groups and resources that are not managed by this protection group, ensure that the `External_dependencies_allowed` property of the protection group is set to `TRUE`.**

```
# geopg list protection-group | grep -i external_dependencies_allowed
```

If necessary, change the property value to `True`.

```
# geopg set-prop -p External_dependencies_allowed=TRUE protection-group
```

4. **(Optional) If the protection group is offline, take offline the application resource group.**

If the protection group is offline, the application resource group must also be offline before it can successfully be added to the protection group.

```
# clresourcegroup offline resource-group
```

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

```
phys-node-n# geopg add-resource-group resource-group-list protection-group
```

resource-group-list

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

protection-group

Specifies the name of the protection group.

This command adds an application resource group to a protection group on the local cluster. 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 2-2 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# geogg add-resource-group apprg1,apprg2 sales-pg
```

Validating an Oracle Data Guard Protection Group

This section provides the following information:

- [“How the Data Replication Layer Validates the Application Resource Groups” on page 48](#)
- [“How the Data Replication Subsystem Verifies the Oracle Data Guard Broker Configuration” on page 49](#)
- [“How to Validate an Oracle Data Guard Protection Group” on page 50](#)

How the Data Replication Layer Validates the Application Resource Groups

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

- **The resource group under the control of the protection group that is being validated does not contain a resource group that contains an Oracle database-server resource.**
 - If you add a failover resource group, it must not contain a `SUNW.oracle_server` resource.
 - If you add a scalable resource group, it must not contain a `SUNW.scalable_rac_server_proxy` resource.
 - Neither a failover nor a scalable resource group can contain an `ORCL.oracle_external_proxy` resource.

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

- **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 in the protection group need to be online only on the primary cluster.

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

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, a shadow Oracle database-server resource group and a replication resource group for the Oracle Data Guard Broker configuration are created. In addition, the configuration is successfully validated on the local cluster. However, the configuration might not be valid on the remote cluster. You can use the `geopg validate protection-group` command on the remote cluster to troubleshoot an invalid configuration.

Note - To avoid possible configuration errors, do not create these resource groups separately, before running the `geopg add-replication-component` command.

The shadow Oracle database-server resource group contains an Oracle Solaris Cluster resource. This resource is based on the Generic Data Service `SUNW.gds` resource type. The shadow Oracle database-server resource shadows the real Oracle database-server resource that manages and monitors the Oracle database in the Oracle Data Guard Broker configuration.

For more information about the shadow Oracle database-server resource group, see [“Oracle Data Guard Shadow Resource Groups” on page 17](#).

The replication resource group contains an Oracle Solaris Cluster resource that is based on the `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 19](#).

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

- The resource group that is named in the `local_oracle_svr_rg_name` property contains a resource of the appropriate resource type:

- For a scalable resource group, the property specifies a resource group that contains a resource of the `SUNW.scalable_rac_server_proxy` or `SUNW.oracle_external_proxy` resource type.
- For a failover resource group, the property specifies a resource group that contains a resource of the `SUNW.oracle_server` or `SUNW.oracle_external_proxy` resource type.

When the resource group contains a resource of the `SUNW.scalable_rac_server_proxy` or `SUNW.oracle_server` resource type, that resource is used to determine the values for `${ORACLE_HOME}` and the local Oracle database SID values.

- The Oracle Data Guard `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. Or, the Oracle wallet connection mechanism, `dgmgrl /@service_name`, can successfully connect to the broker.
- 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, the standby type, that `FAST_START FAILOVER` is disabled, and that `BystandersFollowRoleChange` is set to `NONE`.



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 Oracle database-server 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 Validate an Oracle Data Guard Protection Group

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

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

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

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

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

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

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 protection-group
```

Example 2-3 Validating the Configuration of a Protection Group

This example shows how to validate a protection group.

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


Configuring 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 53](#)
- [“Administering Oracle Data Guard Application Resource Groups” on page 61](#)
- [“Administering Oracle Data Guard Broker Configurations” on page 63](#)
- [“Replicating the Oracle Data Guard Protection Group Configuration to a Partner Cluster” on page 66](#)
- [“Checking the Runtime Status of Oracle Data Guard Data Replication” on page 69](#)

Working With Oracle Data Guard Protection Groups

Unlike other data replication mechanisms, such as Availability Suite feature of Oracle Solaris, Hitachi TrueCopy or Universal Replicator, and EMC SRDF, Oracle Data Guard is an integral part of Oracle database software. Consequently, you do not place Oracle database-server 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.

You must set the Oracle Data Guard Broker property `BystandersFollowRoleChange` to `NONE` as soon as the broker configuration is created and preferably before the broker is added to the Geographic Edition configuration.

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 39](#).
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 41](#).
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 66](#).
4. On the same node, add the shadow Oracle database-server 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 46](#).
5. Activate the protection group, either globally from either cluster or locally from the primary.
This procedure is covered in [“How to Activate a Protection Group” in “Oracle Solaris Cluster Geographic Edition System Administration Guide ”](#).

▼ 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. **Ensure that the Oracle Data Guard Broker properties `BystandersFollowRoleChange` and `FAST_START FAILOVER` are set properly.**
 - a. **Set `BystandersFollowRoleChange` to `NONE`.**

```
DGMGRL> edit configuration set property BystandersFollowRoleChange=NONE;
```
 - b. **Ensure that `FAST_START FAILOVER` is disable.**

2. 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.

3. 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.

For ease of management and greater security, configure an Oracle wallet to manage public key security credentials on Oracle clients and servers. If you use an Oracle wallet, specify only -p sysdba_username= and the Geographic Edition Data Guard module will use the /@local_service_name or /@remote_service_name connect string syntax. For more information, see “Using Oracle Wallet Manager” in *Oracle Database Advanced Security Administrator's Guide* (http://download.oracle.com/docs/cd/E11882_01/network.112/e10746/asowalet.htm#ASOAG160).

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"
```

4. Confirm that the shadow Oracle database 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 Diskgroup online	phys-paris-1	Online	Online -
	phys-paris-2	Online	Online -
Diskgroup online			

qfs-oradata-mds-rs Service is online.	phys-paris-1	Online	Online -
	phys-paris-2	Offline	Offline
scal-oradata-mp-rs	phys-paris-1	Online	Online
	phys-paris-2	Online	Online
rac_server_proxy-rs Oracle instance UP	phys-paris-1	Online	Online -
	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

5. 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.

6. 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

```

7. 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.

```

8. 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
-----                -
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	Online not
monitored			
	phys-newyork-2	Online but	Online not
monitored			

```

geo-clustername                phys-newyork-1  Offline      Offline
                                phys-newyork-2  Online       Online -
LogicalHostname online.

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

geo-failovercontrol            phys-newyork-1  Offline      Offline
                                phys-newyork-2  Online       Online - Service is
online.

mysales_com-odg-rep-rs         phys-newyork-1  Offline      Offline
                                phys-newyork-2  Offline      Offline

mysales_com-rac-proxy-svr-shadow-rs  phys-newyork-1  Offline      Offline
                                phys-newyork-2  Offline      Offline

```

9. From any node in a partner cluster, add the shadow Oracle database-server 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 database-server 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 database-server resource group and highlights whether the cluster is the Oracle Data Guard primary cluster.

10. 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.

```

11. 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

```

Administering Oracle Data Guard Application Resource Groups

This section shows you how to delete an application resource group from an Oracle Data Guard protection group.

▼ 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 Oracle database-server resource groups at any time, without affecting the Oracle database-server resource groups or Oracle databases that they represent. You can remove these resource groups because the shadow Oracle database-server resource groups simply reflect the status of the real Oracle database-server resource groups and do not control the Oracle databases.

Note - You can also accomplish this task by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, click the protection group name, and in the Data Replication Components section click Remove. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Before You Begin Ensure that the following conditions are met:

- The protection group is defined on the local cluster.

- The resource group to remove is part of the application resource groups of the protection group.

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

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

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

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

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

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

```
phys-node-n# geopg remove-resource-group resource-group-list protection-group
```

resource-group-list

Specifies the name of the application resource group.

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

protection-group

Specifies the name of the protection group.

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 3-1 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.

- [“Oracle Data Guard Broker Configuration Properties” on page 63](#)
- [“How to Modify Geographic Edition Properties of an Oracle Data Guard Broker Configuration” on page 64](#)
- [“How to Delete an Oracle Data Guard Broker Configuration From an Oracle Data Guard Protection Group” on page 65](#)

For details about configuring an Oracle Data Guard protection group, see [“How to Create and Configure an Oracle Data Guard Protection Group” on page 39](#).

Oracle Data Guard Broker Configuration Properties

The `Protection Mode` property is the only Oracle Data Guard Broker configuration property that you can change by using the Geographic Edition maintenance commands.

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 must 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`.

For example, if you change the `standby_type` property with the Geographic Edition software, this change does not convert the configuration between the physical and snapshot standby states. Therefore, you must always set the value of the `standby_type` property to a value that matches the standby state that the database currently has configured. Otherwise, the configuration will experience probe and validate errors.

▼ How to Modify Geographic Edition Properties of an Oracle Data Guard Broker Configuration

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

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

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

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

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

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

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...] \  
ODG-configuration-name protection-group
```

-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](#)”.

ODG-configuration-name

Specifies the name of the Oracle Data Guard Broker configuration.

protection-group

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

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

Before You Begin 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 [“Deleting Protection Groups and Data Replication Components” in “Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

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

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

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

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

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

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 shadow Oracle database-server resource group and replication resource group for this Oracle Data Guard Broker configuration.

```
phys-node-n# geopg remove-replication-component ODG-configuration-name protection-group
```

ODG-configuration-name

Specifies the name of the Oracle Data Guard Broker configuration.

protection-group

Specifies the name of the protection group.

Example 3-2 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

Perform this procedure from a node of the cluster to which you want information replicated, for example, `phys-newyork-1`.

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

Before You Begin 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 Oracle database-server 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 Oracle database-server 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 Oracle database-server resource group on the clusters. The means by which you can successfully add a shadow Oracle database-server resource group to a protection group include the following:

- If an Oracle Data Guard protection group contains an Oracle Data Guard Broker configuration, when you replicate the protection group to the partner cluster, the Geographic Edition module for Oracle Data Guard creates any missing shadow Oracle database-server resource group on the partner cluster.
- 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, the Geographic Edition module for Oracle Data Guard adds the shadow Oracle database-server resource group on both clusters.

Once a shadow Oracle database-server resource group exists on both clusters, you can add that resource group to the protection group.

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

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

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

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

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

2. Replicate the protection group configuration to the partner cluster.

```
phys-newyork-1# geopg get -s partnership ODG-protection-group
```

```
-s partnership
```

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

```
ODG-protection-group
```

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. If the corresponding Oracle Data Guard configuration is enabled, the `geopg get` command also disables the Oracle Data Guard configuration.

The `geopg get` command replicates Geographic Edition related entities. To replicate Oracle Solaris Cluster resource groups, resource types, and resources, use the `cluster export -t rg,rt,rs` command to generate an XML cluster configuration file, modify the XML file for the expected configuration on the secondary cluster. Then run the `clresource create` command

with the `-a` option to apply the configuration updates. For more information, see [“How to Configure Oracle Solaris Cluster Software on All Nodes \(XML\)”](#) in [“Oracle Solaris Cluster Software Installation Guide ”](#) and the `cluster(1CL)` and `clresource(1CL)` man pages.

Example 3-3 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`.

```
# ssh root@phys-newyork-1
phys-newyork-1# geopg 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.

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 69
- [“Displaying a Detailed Oracle Data Guard Runtime Status”](#) on page 70
- [“Enabling Debugging for Runtime Messages”](#) on page 73

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

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

1. Log in to a 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 [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration 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 ODG-configuration-name-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 3-1](#).

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:

```
ODG-protection-group-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:

ODG-configuration-name-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 ODG-configuration-name-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 3-1 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.
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 (<code>dgmgrl</code>).
Faulted	Oracle errors " <i>List-of-ORA-xxxx-errors</i> " were found in the Oracle Data Guard broker (<code>dgmgrl</code>) 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 snapshot standby.
Unknown	Unexpected error - <i>unexpected-error</i>	
Unknown	Oracle Data Guard broker (<code>dgmgrl connect-string</code>) 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 (<code>dgmgrl</code>) did not respond to the <code>show configuration</code> 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 <code>sysdba_username</code> , <code>sysdba_password</code> , <code>local_db_service_name</code> , or <code>remote_db_service_name</code> parameter does not match the information that is maintained by the Geographic Edition software.

How to Check the Overall Runtime Status of Replication

Status	Status Message	Possible Causes
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.
Unknown	A switchover is in progress	Self-explanatory.
Unknown	A failover is in progress	Self-explanatory.
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 sysdba user name for Oracle Data Guard configuration <i>ODG-configuration-name</i> in protection group <i>ODG-protection-group</i>	The field for the sysdba_password was not found in the Cluster Configuration Repository (CCR) or was longer than expected.
Degraded	Local cluster <i>cluster-name</i> is not primary for Oracle Data Guard configuration <i>ODG-configuration-name</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>ODG-configuration-name</i> found does not match <i>ODG-configuration-name</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>ODG-configuration-name</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>ODG-configuration-name</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.
Degraded	The BystandersFollowRoleChange property for Oracle Data Guard Broker configuration <i>Broker_Config_Name_Variable</i> must be set to 'NONE'	The Oracle Data Guard Broker property is not set to 'NONE'.
Degraded	Fast-start failover must be disabled for Oracle Data Guard Broker configuration <i>Broker_Config_Name_Variable</i>	Fast-start failover is enabled.
Online	Online or replicating in <i>replication-mode</i> mode	

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

Enabling Debugging for Runtime Messages

This section provides information to enable debug messages for the Oracle Data Guard control shell scripts as well as for the common agent container, and to verify that debugging is enabled.

▼ How to Enable Debugging for Runtime Messages

Perform this task on each cluster node where you want debugging active for Oracle Data Guard data replication.

1. **Assume the root role on the cluster node.**
2. **Determine whether debugging is active for Oracle Data Guard replication with Geographic Edition.**

```
# grep info /etc/syslog.conf
#
```

Debugging is active if *.info is set in the /etc/syslog.conf file.

3. **If debugging is inactive, add the *.info setting to the /etc/syslog.conf file.**
4. **Verify the addition of the *.info setting to the /etc/syslog.conf file.**

```
# grep info /etc/syslog.conf
*.err;kern.debug;daemon.notice;mail.crit,*.info /var/adm/messages
#
```

5. **Restart the system-log service.**

```
# svcadm restart system-log
```

Next Steps To disable debugging, remove the *.info setting from the /etc/syslog.conf file and restart the system-log service.

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 75](#)
- [“Migrating Services That Use Oracle Data Guard With a Switchover” on page 76](#)
- [“Recovering Oracle Data Guard Data After a Takeover” on page 80](#)
- [“Recovering From an Oracle Data Guard Data Replication Error” on page 85](#)

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`.
- For HA for Oracle configurations, the `dataguard_role` resource property is updated to reflect the new status of the new primary and standby clusters.
- Application services and the shadow Oracle database-server resource groups are brought online on the new primary cluster `cluster-newyork`.

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 offline the application resource groups in the protection group 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 online all the shadow Oracle database-server resource groups and any other application resource groups in the protection group

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. For HA for Oracle configurations, the `dataguard_role` resource property is also updated with the status of the new primary and standby clusters. The shadow Oracle database-server resource group and any other application resource groups in the protection group are online on the new primary cluster. Data replication from the new primary cluster to the new standby cluster begins.

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 50](#).

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 a Protection Group” in “Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

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 D, “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 shows one of the following states:

- A SUCCESS state
- Busy performing a health check (ORA-16610)
- The remote database is unreachable (ORA-16625)
- The `show configuration` command times out

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 offline the application resource groups in the protection group 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` command. If this command fails and the system remains a standby, the `failover to standby-database-name immediate` command is run.
- 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 online all shadow Oracle database-server resource groups and application resource groups in the protection group. For HA for Oracle configurations, the `dataguard_role` resource property is also updated to reflect the new primary and standby clusters.

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 shadow Oracle database-server resource group and any other application resource group in the protection 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 50.

If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures described in [“How to Resynchronize a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

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.

▼ How to Perform a Failback Switchover or Failback Takeover

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

Note - You can also accomplish some steps in this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, highlight the protection group name, and click the button for the action you want to perform. For more information about Oracle Solaris Cluster Manager see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in [“Oracle Solaris Cluster System Administration Guide ”](#).

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

Before You Begin Ensure that the clusters 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 protection group could be 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. **For HA for Oracle, on the original primary cluster, verify that the `SUNW.oracle_server` resource is in a healthy state.**

If HA for Oracle is not running on the original primary cluster, omit this step.

- **If the resource is not in a faulted state, unmonitor the database resource.**

```
# clresource unmonitor oracle_server-rs
```

- **If the resource is in a faulted state or repeatedly restarts, perform the following steps:**

- a. **Disable the HA for Oracle resource or resource group.**

- **If the `dataguard_role` property is set to STANDBY, disable the HA for Oracle resource.**

A STANDBY value is set if the takeover was performed when the old primary was running at the time of the takeover.

```
phys-paris-1# clresource disable oracle_server-rs
```

- **If the `dataguard_role` property is set to PRIMARY, disable the HA for Oracle resource group.**

A PRIMARY value is set if the takeover was performed when the old primary was down during the takeover.

```
phys-paris-1# clresourcegroup quiesce -k oracle_server-rg
phys-paris-1# clresource disable oracle_server-rs
phys-paris-1# clresourcegroup offline oracle_server-rg
phys-paris-1# clresourcegroup online oracle_server-rg
```

Note - When the cluster restarts, an attempt is made to start a database that needs to be reinstated. Therefore, you must disable the resource as soon as possible. You might need to quiesce the HA for Oracle resource group if the RGM has already attempted to bring it online.

If the Oracle Database resource is in the `stop_failed` state, clear the `stop_failed` flag by using the following command.

```
phys-paris-1# clresource clear oracle_server-rs
```

b. Determine whether the database is shut down on the cluster nodes.

- c. If the database is not shut down, become the Oracle user on that node and stop the database by using one of the following methods:**

First method:

```
phys-paris-1$ srvctl stop database -d database_name
```

Second method:

```
phys-paris-1$ ORACLE_SID=db_SID export ORACLE_SID
phys-paris-1$ sqlplus /nolog
SQL> connect sys/sysdba password as sysdba
SQL> shutdown immediate
SQL> exit
```

d. Start mount the database.

```
phys-paris-1$ sqlplus /nolog
SQL> connect sys/sysdba password as sysdba
SQL> startup mount
...
SQL> exit
```

- 3. Reinstate the old Oracle Data Guard primary database to become the standby for the current primary database.**

If you issue the `dgmgrl` command from the old primary cluster, include the new primary's database service name in the connection string.

```
phys-newyork-1$ dgmgrl
DGMGRL> connect sys/password[@new_primary_service_name]
DGMGRL> reinstate database old_primary_database_name
...
DGMGRL> exit
```

Note - If the database cannot be reinstated, you might need to re-create it or otherwise recover the database by using an appropriate method. For instructions, refer to [“Using Flashback Database After a Failover”](#) in *Oracle Data Guard Concepts and Administration*.

4. **To perform a failback takeover instead of a failback switchover, flashback your primary database to the point at which the original takeover occurred.**
5. **For HA for Oracle, update and re-enable the HA for Oracle resource on the original primary cluster.**

If HA for Oracle is not running on the original primary cluster, omit this step.

```
phys-paris-1# clresource set -p dataguard_role=STANDBY oracle_server-rs
phys-paris-1# clresource enable oracle_server-rs
    Restore monitoring if monitoring of the resource was previously disabled
phys-paris-1# clresource monitor oracle_server-rs
```

6. **If the original primary cluster was down at the point of failure, update the original primary cluster to be the secondary.**

- a. **From a node of the original primary cluster, stop the protection group.**

If the original primary cluster was down at the time of takeover, the protection group should already be stopped.

```
phys-paris-1# geopg stop -e local protection-group
```

```
-e local
```

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

```
protection-group
```

Specifies the name of the protection group.

- b. **Verify that the protection group is stopped.**

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

c. Update the protection group.

```
phys-paris-1# geopg update protection-group
```

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

For more information about synchronizing protection groups, see [“Resynchronizing a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

7. From one node in each cluster, locally validate the configuration for each protection group.

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

```
phys-paris-1# geopg validate protection-group
phys-newyork-1# geopg validate protection-group
```

For more information, see [“How to Validate an Oracle Data Guard Protection Group”](#) on page 50.

8. From one node in either cluster, globally activate the protection group on both clusters.

```
phys-node-n# geopg start -e global protection-group
```

9. From one node in either cluster, switch over the protection group to the original primary.

```
phys-node-n# geopg switchover -f -m cluster-paris protection-group
```

For more information, see [“Migrating Replication Services by Switching Over Protection Groups”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

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.

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 `OnLine` because the probe is still running correctly.

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

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

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
\  
ODG-configuration-name protection-group
```

3. Revalidate the protection group configuration.

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

```
phys-paris-1# geogg validate protection-group
```

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

4. Review the status of the protection group configuration.

```
phys-paris-1# geogg list protection-group
```

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

◆◆◆ APPENDIX A

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 Replication Component Properties

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

Data replication component 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 database on the local cluster.

Category:

Required

Default:

None

Tunable:

At creation

Data replication component 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 component property: `local_oracle_svr_rg_name` (string)

Name of the local Oracle database server resource group that manages the local database in the Oracle Data Guard Broker configuration. A shadow Oracle database-server resource group shadows the real resource group. If you want, add the shadow to the protection group application resource group list.

Note - The previous name of this property, `local_rac_proxy_svr_rg_name`, is still valid.

Category:

Required

Default:

None

Tunable:

At creation

Data replication component 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 database on the remote cluster.

Category:

Required

Default:

None

Tunable:

At creation

Data replication component 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 component property: `remote_oracle_svr_rg_name` (string)

Name of the remote Oracle database-server resource group on the partner cluster that manages the remote database in the Oracle Data Guard Broker configuration. A shadow Oracle database-server resource group shadows the real resource group. If you want, add the shadow to the protection group application resource group list.

Note - The previous name of this property, `remote_rac_proxy_svr_rg_name`, is still valid.

Category:

Required

Default:

None

Tunable:

At creation

Data replication component 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 component 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`, `physical`, and `snapshot`.

Category:

Required

Default:

None

Tunable:

Any time

Data replication component 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 if an Oracle wallet is not used

Default:

None

Tunable:

Any time

Data replication component 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 if an Oracle wallet is not used

Default:

None

Tunable:

Any time

Index

A

- administering
 - data replication with Oracle Data Guard, 15, 53
 - Oracle Data Guard Broker configurations, 63
- application resource groups
 - administering, 61
 - creating, 46
 - removing, 61

C

- caution notice
 - HA for Oracle switchover, 16
 - suspending replication after takeover, 79
 - switchover outside Geographic Edition control, 16
 - sysdba_password property, 55
- configuration summary, 16
- configuring
 - Oracle Data Guard Broker configurations, 35
 - Oracle Data Guard configuration, 21
 - Oracle Data Guard software, 20
 - protection groups, 39
- creating
 - application resource group, 46
 - protection groups, 39
 - replication Oracle Data Guard Broker configurations, 41

D

- data recovery, 80
 - failback switchover, 81
- debugging
 - enabling, 73
- deleting

- application resource group, 61
 - replication Oracle Data Guard Broker configuration, 65
- detecting failure, 75

E

- enabling
 - debugging, 73

F

- failback switchover, 81
- failure
 - detecting, 75
 - primary cluster, 75
 - standby cluster, 76

L

- local_database_name, 43, 87
- local_db_service_name, 43, 87, 88
- local_oracle_svr_rg_name, 44, 88

M

- messages
 - enabling debugging, 73
- migrating services, 75
 - data recovery after, 80
 - with a switchover, 76
 - with a takeover, 78
- modifying

- replication Oracle Data Guard Broker configurations, 64

- O**
- Oracle Clusterware, 21, 34
- Oracle Data Guard
 - administering data replication with, 15, 53
 - configuring software, 20
 - detecting failure, 75
 - initial replication configuration, 19
 - migrating services that use, 75
 - properties of
 - local_database_name, 43, 87
 - local_db_service_name, 43, 87
 - local_oracle_svr_rg_name, 44, 88
 - remote_database_name, 44, 88
 - remote_db_service_name, 44, 88
 - remote_oracle_svr_rg_name, 44, 89
 - replication_mode, 44, 89
 - standby_type, 44, 89
 - sysdba_password, 44, 90
 - sysdba_username, 44, 90
 - replication component properties for, 87
 - replication resource groups, 19
 - runtime status, 69
 - overall, 69
 - shadow resource groups, 17
- Oracle Data Guard Broker configurations
 - adding to protection group, 41
 - administering, 63
 - configuring, 35
 - modifying, 64
 - removing, 65
- Oracle Data Guard configuration
 - configuring, 21
 - setting up primary database, 21
- Oracle Grid Infrastructure, 21, 34
- Oracle wallet, 44, 55, 90
- Oracle Wallet Manager, 55

- P**
- primary cluster
 - data recovery, 80
 - failure detection, 75
 - switchover, 76
 - takeover, 78
- properties
 - Oracle Data Guard, 87
 - local_database_name, 43, 87
 - local_db_service_name, 43, 87
 - local_oracle_svr_rg_name, 44, 88
 - remote_database_name, 44, 88
 - remote_db_service_name, 44, 88
 - remote_oracle_svr_rg_name, 44, 89
 - replication_mode, 44, 89
 - standby_type, 44, 89
 - sysdba_password, 44, 90
 - sysdba_username, 44, 90
- protection groups
 - adding application resource group to, 46
 - adding Oracle Data Guard Broker configurations to, 41
 - adding shadow Oracle database-server resource group to, 46
 - configuring, 39
 - creating, 39
 - creation strategies, 53
 - modifying Oracle Data Guard Broker configurations for, 64
 - removing application resource group, 61
 - removing Oracle Data Guard Broker configuration from, 65
 - removing shadow Oracle database-server resource group, 61
 - replicating configuration of, 66
 - validating, 50

- R**
- recovery *See* data recovery
 - from replication error, 85
- remote_database_name, 44, 88
- remote_db_service_name, 44
- remote_oracle_svr_rg_name, 44, 89
- replication
 - adding replication component, 41
 - initial configuration of, 19

- migrating services, 75
- modifying Oracle Data Guard Broker configurations, 64
- Oracle Data Guard, 15, 53
- protection group configuration, 66
- recovering from errors, 85
- removing Oracle Data Guard Broker configuration, 65
- resource groups, 19
- runtime status details, 70, 70
- runtime status overview, 69
- replication resource groups and status, 71
- replication_mode, 44, 89
- resource groups
 - application, 61
 - replication, 19
 - shadow, 17
- runtime status
 - replication, 69
 - state and status messages, 71

S

- shadow resource groups, 17
- standby cluster
 - failure detection, 76
 - switchover, 76
 - takeover, 78
- standby_type, 44, 89
- switchover, 76
 - actions performed during, 77
 - caution when outside Geographic Edition control, 16
- sysdba_password, 44, 90
- sysdba_password property
 - caution notice, 55
- sysdba_username, 44, 90

T

- takeover, 78
 - actions performed during, 79
 - caution notice, 79
 - data recovery after, 80
 - failback switchover, 81

V

- validating protection groups, 50

