# Installing and Configuring the Disaster Recovery Framework for Oracle® Solaris Cluster 4.4

**Part No: E69315**
January 2019

ORACLE

Installing and Configuring the Disaster Recovery Framework for Oracle Solaris Cluster 4.4

**Part No: E69315**



**Access to Oracle Support**

Oracle customers that have purchased support 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.

**Référence: E69315**



**Accès aux services de support Oracle**

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` ou le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` si vous êtes malentendant.

# Contents

# Tables

# Examples

# Using This Documentation

- **Overview** – Provides guidelines for planning an Oracle Solaris Cluster disaster recovery framework installation and configuration, and includes procedures to install, configure, upgrade, and uninstall the Disaster Recovery framework software
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Advanced experience troubleshooting and replacing hardware

## Product Documentation Library

Documentation and resources for this product and related products are available at http://docs.oracle.com/cd/E69294_01.

## Feedback

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

♦ ♦ ♦ CHAPTER 1

# Planning the Disaster Recovery Framework Installation

This chapter provides planning information and guidelines for installing an Oracle Solaris Cluster disaster recovery framework configuration. This chapter also describes how to plan the data replication between two clusters.

This chapter contains the following sections:



## Installation Process

To successfully install Disaster Recovery framework software, you must complete the following installation phases:

1. Planning your installation.
2. Connecting your hardware.
3. Installing Oracle Solaris Cluster software.
4. Installing data replication products.
5. Installing and configuring the required software.
6. Installing Disaster Recovery framework software.
7. Configuring the Disaster Recovery framework.

This installation process progresses from the initial planning phase to the eventual startup of the Disaster Recovery framework. This guide provides information about phases 1, 6, and 7.

**Note -** You can also install Disaster Recovery framework software at the same time that you install Oracle Solaris Cluster software.

For information about installing Oracle Solaris Cluster software, see the *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

For information about configuring a cluster after startup, see the *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

**Caution -** The `/var/cluster.sav/` directory is created and used during current and future Disaster Recovery framework installation, update, and uninstallation operations. This directory and its contents are for Disaster Recovery framework internal use only, and must not be deleted by the user. User removal of this directory, or any of its contents, would pose a high risk of compromising the cluster during any future installation, update, or uninstallation of Disaster Recovery framework software.

# Planning Cluster Hardware

This section helps you to plan your hardware for the primary cluster, the secondary cluster, and the inter-cluster communication.

The Disaster Recovery framework hardware configuration consists of the following elements:

- At least two separate clusters that are running Oracle Solaris Cluster software with attached data storage. One of these clusters must be designated the primary cluster.

  **Note -** While you can use a single-node cluster at both the primary and backup sites, a single-node cluster offers no internal redundancy. To ensure no single point of failure, you must have a minimum of two nodes in a cluster at the primary site. You can use a single-node cluster at the secondary site as a cost-effective backup solution, if the secondary site is used only for backup purposes and is not for running mission-critical applications.

- Internet connections for inter-cluster management communication between the clusters and for default inter-cluster heartbeats.
- Connections for either host-based or storage-based data replication.
- Connections for custom heartbeats, if any.

The hardware configurations that Disaster Recovery framework software supports are identical to the hardware configurations that the Oracle Solaris Cluster product supports. For use of

Disaster Recovery framework software with storage-based data replication mechanisms, the cluster hardware configurations are those configurations that support the related storage hardware. Partner clusters must be compatibly configured to support data replication between the clusters.

Internet access is required between partner clusters. The communication between partner clusters for inter-cluster management operations is through a logical-hostname IP address. The default inter-cluster heartbeat module also communicates through a logical-hostname IP address.

A cluster in a disaster recovery partnership conforms to the standard configuration rules of a cluster that is running Oracle Solaris Cluster software.

# Planning Required Software

This section helps you to adapt the configuration of your Oracle Solaris Cluster software for the installation of Disaster Recovery framework software. This section also helps you to plan the installation of your data replication software.

The following information is provided in this section:


- "Planning the Disaster Recovery Framework Software" on page 15
- "Planning the Data Replication Software" on page 16
- "Planning Volume Management" on page 17


## Planning the Disaster Recovery Framework Software

The Disaster Recovery framework software must be installed on a cluster that is running the Oracle Solaris OS and Oracle Solaris Cluster software. You can install Disaster Recovery framework software at the same time that you install Oracle Solaris Cluster software or at any time afterwards. The Disaster Recovery framework software configuration is identical to the Oracle Solaris Cluster software configuration.

The clusters in a Disaster Recovery framework configuration can run different versions of Disaster Recovery framework software, as long as they are no more than one consecutive version different. For example, the same Disaster Recovery framework configuration could

have clusters running either version 4.4 or 4.3. But a cluster running version 4.4 cannot run in the same Disaster Recovery framework configuration as a cluster running version 4.0.

# Planning the Data Replication Software

A cluster that is using the Disaster Recovery framework with a data replication product is subject to the standard configuration rules of a cluster that is running the data replication product with Oracle Solaris Cluster software. Partner clusters must have compatible software configurations to support data replication between the clusters.

The Disaster Recovery framework supports the following data replication products:

- MySQL software
- Oracle Data Guard software, in configurations that use Oracle Database software
- Oracle Solaris ZFS snapshots feature
- Oracle ZFS Storage Appliance software
- Disaster recovery framework script-based plug-ins

The following sections describe the types of replication that the above products provide:


- "Host-Based Replication" on page 16
- "Storage-Based Replication" on page 16
- "Replication for Databases" on page 17
- "Built-In Replication" on page 17
- "Custom Replication" on page 17


## Host-Based Replication

The ZFS snapshots feature of Oracle Solaris software is a host-based replication method. This method consists of software installed on a host that controls replication from one server to a secondary server.

## Storage-Based Replication

Oracle ZFS Storage Appliance uses a storage-based method. This method uses replication that is built into the storage hardware. If you use Oracle ZFS Storage Appliance software, you must install the software on each node of the cluster.

### Replication for Databases

Oracle Data Guard functionality is part of the Oracle Database software and so does not require you to install additional software onto your system. The Disaster Recovery framework module for Oracle Data Guard can only be used with Oracle databases.

### Built-In Replication

MySQL database software offers a built-in replication protocol. Configuring the Disaster Recovery framework MySQL replication module enables you to control replication between MySQL instances on each site.

### Custom Replication

The Disaster Recovery framework script-based plug-in enables the user to develop replication modules to integrate additional replication protocols into the Disaster Recovery framework. The plug-in provides the interface to register custom replication control scripts with the Disaster Recovery framework.

## Planning Volume Management

See the Oracle Solaris Cluster 4.4 Compatibility Guide ( `https://community.oracle.com/docs/DOC-997312`) for information about volume management support for your data replication product.

# Planning Resource and Resource Group Names

A partnership requires two clusters to be combined into one environment, and one cluster might be a running production system. Therefore, advance planning of resources and resource groups is essential for a successful installation.

The Disaster Recovery framework requires that resource-group names be identical on each partner cluster to ensure that a resource or resource group can be managed as a single entity across both clusters in the partnership.

# Planning Required IP Addresses and Hostnames

You must have all the required IP addresses and hostnames before you begin the installation process. This section provides information about these requirements.

## IP Address Requirements

You must set up a number of IP addresses for various Disaster Recovery framework components, depending on your cluster configuration. Observe the following guidelines:

- You must have an IP address for the cluster name and for each cluster node.
- See "Public-Network IP Addresses" in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* for a list of components that require IP addresses. Add these IP addresses to any naming services that are used. Also add these IP addresses to the local `/etc/inet/hosts` file on each cluster node after you install Oracle Solaris software.
- You might also need additional IP addresses for data replication products. For more information about requirements for configuring data replication, see the Oracle Solaris Cluster guide for your data replication product.

## Hostname Requirements

Observe the following guidelines:

- **Logical hostname** – A cluster name must be suitable as a hostname because the Disaster Recovery framework creates the logical hostname by using the cluster name. Therefore, the cluster name must be in the naming system.
- **Unique cluster name** – Cluster names must be unique. For example, if you have a cluster wholly within the domain `.france`, you can use hostnames like `paris` and `grenoble`. However, if you have a cross-domain cluster, you must specify the hostnames with enough qualification to identify the host on the network. For example:
  - You can link `paris` and `munich` with hostnames `paris.france` and `munich.germany`, and the cluster names remain `paris` and `munich`.
  - You cannot create a partnership between clusters `paris.france` and `paris.texas` because of a collision on the cluster name `paris`.

# Planning Security

This section contains the following information about securing the Disaster Recovery framework:



## Setting Up and Assigning Rights Profiles With the Disaster Recovery Framework

The Disaster Recovery framework bases its rights profiles on the rights profiles that are used in the Oracle Solaris Cluster software. For general information about setting up and assigning rights profiles with Oracle Solaris Cluster software, refer to Chapter 2, "Oracle Solaris Cluster and User Rights" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

The Disaster Recovery framework adds the following rights entries to the appropriate rights database in the `/etc/security` directory:

- Authentication names to `auth_attr`
- Execution profiles to `prof_attr`
- Execution attributes to `exec_attr`

---

**Note -** The default search order for the `auth_attr` and `prof_attr` databases is `files nis`, which is defined in the `/etc/nsswitch.conf` file. If you have customized the search order in your environment, confirm that `files` is in the search list. Including `files` in the search list enables your system to find the rights database entries that the Disaster Recovery framework defined.

---

## Disaster Recovery Framework Rights Profiles

The Disaster Recovery framework CLI and the Oracle Solaris Cluster Manager browser interface use rights profiles to control end-user access to operations. The general conventions

for these rights are described in Table 1, “Disaster Recovery Framework Rights Profiles,” on page 20.

**TABLE 1** Disaster Recovery Framework Rights Profiles

| Management Rights Profile | Authorizations | Rights Granted |
|---|---|---|
| Geo Management | `solaris.cluster.geo.read` | Read information about the Disaster Recovery framework entities |
| | `solaris.cluster.geo.admin` | Perform administrative tasks with the Disaster Recovery framework |
| | `solaris.cluster.geo.modify` | Modify the configuration of the Disaster Recovery framework |
| Basic Solaris User | Oracle Solaris authorizations | Perform the same operations that the Basic Solaris User role identity can perform |
| | `solaris.cluster.geo.read` | Read information about the Disaster Recovery framework entities |

# Configuring Firewalls

The Disaster Recovery framework partner clusters communicate using transport services and ICMP echo requests and replies (pings). Their packets must therefore pass data center firewalls, including any firewalls configured on cluster nodes in partner clusters. The table below contains a list of required and optional services and protocols used by Disaster Recovery framework partnerships, and the associated ports that you must open in your firewalls for these services to function. The ports listed are defaults, so if you customize the port numbers serving the specified transfer protocols, the customized ports must be opened instead.

Ports other than those listed in Table 2, “Ports and Protocols Used by Disaster Recovery Framework Partnerships - Required Services,” on page 20 and Table 3, “Ports and Protocols Used by Disaster Recovery Framework Partnerships - Optional Services,” on page 21 might be required by storage replication services. See product documentation for details.

**TABLE 2** Ports and Protocols Used by Disaster Recovery Framework Partnerships - Required Services

| Port Number | Protocols | Use in Disaster Recovery framework partnership |
|---|---|---|
| 22 | UDP and TCP | Secure shell (`ssh`). Used during the initial certificate transfer that establishes trust between partner clusters. |
| 2084 | UDP (default), TCP | Intercluster heartbeat |

| Port Number | Protocols | Use in Disaster Recovery framework partnership |
|---|---|---|
| 11162 | TCP | The Java Management Extensions (JMX) port (`jmxmp-connector-port`). A messaging protocol used for the exchange of configuration and status information between the two sites in a partnership. |
| - | ICMP Echo Request/ Reply | Backup heartbeat between partner clusters |

**TABLE 3** Ports and Protocols Used by Disaster Recovery Framework Partnerships - Optional Services

| Port Number | Protocols | Use in Disaster Recovery framework partnership |
|---|---|---|
| 161 | TCP and UDP | Simple Network Management Protocol (SNMP) communications |
| 162 | TCP and UDP | SNMP traps |

# Securing Inter-Cluster Communication

This section provides the information about the following methods to secure communication between partner clusters:

- “Security Certificates” on page 21
- “IP Security (IPsec)” on page 21

## Security Certificates

You must configure the Disaster Recovery framework for secure communication between partner clusters. The configuration must be reciprocal, so cluster `cluster-paris` must be configured to trust its partner cluster `cluster-newyork`, and cluster `cluster-newyork` must be configured to trust its partner cluster `cluster-paris`.

For information and procedures to set up security certificates for partner clusters, see “Configuring Trust Between Partner Clusters” on page 43.

## IP Security (IPsec)

You can use IP Security Architecture (IPsec) to configure secure communication between partner clusters. IPsec enables you to set policies that permit or require either secure datagram authentication, or actual data encryption, or both, between machines communicating by using IP.

Consider using IPsec for secure TCP/UDP heartbeat communications.

IPsec uses two configuration files:

- **IPsec policy file**, `/etc/inet/ipsecinit.conf`. Contains directional rules to support an authenticated, encrypted heartbeat. The contents of this file are different on the two clusters of a partnership.
- **IPsec keys file**, `/etc/init/secret/ipseckeys`. Contains keys files for specific authentication and encryption algorithms. The contents of this file are identical on both clusters of a partnership.

Observe the following guideline when using IPsec for secure inter-cluster communication:

- Oracle Solaris Cluster software and Disaster Recovery framework software support IPsec by using only manual keys. Keys must be stored manually on the cluster nodes for each combination of server and client IP address. The keys must also be stored manually on each client.
- In the Disaster Recovery framework infrastructure, the hostname of a logical host is identical to the cluster name. The logical hostname is a special HA resource. You must set up a number of IP addresses for various Disaster Recovery framework components, depending on your cluster configuration.
- On each partner cluster, you must configure encryption and authorization for exchanging inbound and outbound packets from a physical node to the logical-hostname addresses. The values for the Oracle Solaris IP Security Architecture (IPsec) configuration parameters on these addresses must be consistent between partner clusters.

Refer to *Securing the Network in Oracle Solaris 11.4* for more information about IPsec.

# Planning the Disaster Recovery Framework Environment

This section provides guidelines for planning and preparing the following components for Disaster Recovery framework software installation:

# License Certificates

Ensure that you have available all necessary license certificates before you begin software installation. The Disaster Recovery framework software does not require a license certificate. However, each node that is installed with Disaster Recovery framework software must be covered under your Oracle Solaris Cluster software license agreement.

For licensing requirements for data replication software and application software, see the installation documentation for those products.

# Logical Hostnames

The Disaster Recovery framework uses the logical hostname of a cluster for inter-cluster management communication and heartbeat communication. The IP address for a cluster name must be available for the Disaster Recovery framework to wrap a logical hostname around the IP address when the software is started by using the `geoadm start` command.

You can use the `cluster` command to find the name of the cluster when you need to verify that the cluster name is suitable for use as a hostname. To find the name of the cluster, run the following command:

```
# cluster list
```

For more information, see the `cluster`(8CL) man page.

# Zone Clusters

In some Disaster Recovery framework configurations, a zone cluster can be configured as a cluster partner. Observe the following guidelines for the use of zone clusters in a cluster partnership.

- **Public-network IP addresses** - A zone cluster that is configured in a Disaster Recovery framework configuration must meet the following public-network requirements:
  - Each zone-cluster node must have a public-network IP address that corresponds to the zone-cluster node's hostname.
  - The zone-cluster node's public-network IP address must be accessible by all nodes in the Disaster Recovery framework configuration's partner cluster.

    - Each zone-cluster node must have a failover IP address that maps to the hostname that corresponds to the zone-cluster name.
- **Mixed cluster types** – The partnership can use other zone clusters or a combination of zone clusters and global clusters.
- **Framework packages** – The Disaster Recovery framework packages are required in the global zones in all cases, even if the Disaster Recovery framework is only going to be enabled in the zone clusters. The Disaster Recovery framework group package is `ha-cluster/geo/geo-framework`.
- **Starting the infrastructure** – You can start the disaster recovery framework from within a zone cluster node, but not from within any other type of non-global zone.

# Cluster Partnerships

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

You create only one partnership between two specific clusters by using the `geops(8)` command. After you have created a partnership, you can use this command to modify the properties of this partnership.

Observe the following guidelines:

- **Unique cluster names** – When creating partnerships, ensure that the name of all the clusters in the partnership are unique. For example, if you have a cluster wholly within the domain `.france`, you can use hostnames like `paris` and `grenoble`. However, if you have a cross-domain cluster, you must specify the hostnames with enough qualification to identify the host on the network. For example:
    - You can link `paris` and `munich` with hostnames `paris.france` and `munich.germany`, and the cluster names remain `paris` and `munich`.
    - You cannot create a partnership between clusters `paris.france` and `paris.texas` because of a collision on the cluster name `paris`.
- **Application resource group names** – The names of the application resource groups that are managed by the Disaster Recovery framework must be the same on both partner clusters. You can configure the names of these resource groups manually.
- **Single partnership between cluster pairs** – You can define only one partnership between two specific clusters. A single cluster can participate in other partnerships with different clusters.

- **Device groups** – You cannot add device groups to a protection group that does not use data replication.

# Cluster Protection Groups

Protection groups enable a set of clusters to tolerate and recover from disaster by managing the resource groups for services. A protection group contains application resource groups and properties for managing data replication for those application resource groups.

Observe the following general guidelines when you configure a protection group:

- **Partnerships** – You must create a partnership before you can create a protection group for that partnership. Protection groups can exist only in a partnership.
- **Unique naming** – Protection groups must have unique name across partnerships. If this is not the case the protection group will go to mismatch state.
- **Duplicate application resource group configuration** – You can duplicate the application resource group configuration on partner clusters. The configuration for a protection group is identical on partner clusters, so partner clusters must have the application resource groups of the protection group defined in their configuration. The Disaster Recovery framework propagates protection group configurations between partners.
- **Data replication** – You can specify a data replication type in the protection group to indicate the mechanism that is used for data replication between partner clusters. When a service is protected from disaster by data replication, the protection group also contains replication resource groups. Protection groups link an application in a resource group with the application data that should be replicated. This linkage and replication enable the application to fail over seamlessly from one cluster to another cluster.
- **Replicating the Oracle Solaris boot environment** – Do not replicate an Oracle Solaris boot environment between two systems. Doing so is not appropriate for disaster recovery environments, as it might introduce instability in the target boot environment.
- **Application resource groups** – Certain `rg_affinities` settings, such as the "++" affinity will disallow the dependee resource group from going to the unmanaged state if the dependent resource group is not in unamanaged state. Some resource dependencies (such as `resource_dependencies_offline_restart`) will result in the dependee resource going offline if the dependent resource is brought offline. A protection group switchover disables application resources in the resource groups it contains, then offlines these resource groups and lastly, it unmanages these resource groups. If any of these operations fail, the switchover will fail. If you must keep dependent resource groups outside of the protection group managing dependee resource group(s), the protection group's `External_Dependency_Allowed` property must be set to `true`. In addition, consider the following options:

- Put dependent resource groups in other protection group(s), then put these protection groups in the same multigroup as the dependee protection group with dependency from protection group that contains dependent resource group(s) on this dependee protection group. To initiate a switchover use the `geomg` command with the multigroup and *not* `geopg` with the individual protection groups.
- If you keep the dependent resource groups outside of protection groups, you will always always need to bring them to the unmanaged state before initiating the switchover of the protection groups containing the dependee resource groups.

Each data replication product has its own additional requirements when configuring a protection group. For more information, see the appropriate Oracle Solaris Cluster guide for the data replication software that you will use:

- *Oracle Solaris Cluster Data Replication Guide for MySQL*
- *Oracle Solaris Cluster Data Replication Guide for Oracle Data Guard*
- *Oracle Solaris Cluster Data Replication Guide for ZFS Snapshots*
- *Oracle Solaris Cluster Remote Replication Guide for Oracle ZFS Storage Appliance*

## Sites for Protection Groups

A site is a group of clusters for which you want to manage sets of protection groups, or multigroup, in a single operation. When you perform a switchover or takeover of a multigroup, a site is specified as the target.

Observe the following general guidelines when you configure a site:

- **First site controller** – The cluster from which you create a new site is automatically configured as a site controller.
- **Multiple controllers** – To avoid a single point of failure, configure at least two controller clusters in a site.
- **Zone clusters** – A zone cluster can be a site member.
- **Remote management of an Oracle Data Guard database** – A cluster can be configured with a protection group to manage a remote Oracle Data Guard Database instance that is not running on an Oracle Solaris Cluster configuration.

## Cluster Multigroups

A multigroup is a set of protection groups that you can manage in a single operation.

Observe the following general guidelines when you configure a multigroup:

- **Single site** – All protection groups in a multigroup must be configured on clusters that are part of the same site.
- **Unique name** – Multigroup names must be unique throughout the site. In addition, if multiple sites share a common cluster, those sites cannot contain multigroups of the same name.
- **Site-to-cluster configurations** – A multigroup can consist of protection groups where one of the partner clusters is not configured in a site. In such a configuration, multigroup operations can only be performed from a cluster that is in a site. To manage protection groups from the partner cluster that is not in a site, you must manage the protection groups individually using the `geopg` command.
- **Protection group dependencies** – One or more protection groups can be configured to have a strong dependency on a third protection group in the multigroup. When the protection groups in a multigroup are taken offline for a switchover or takeover, the depended-on protection group is taken offline after the protection groups that depend on it are taken offline. And when a multigroup is brought online, the depended-on protection group is brought online before the protection groups with a dependency on it are brought online.
- **Nested protection group dependencies** – A protection group that other protection groups depend on can itself have a dependency on another protection group.
- **Single dependency** – A protection group cannot have a direct dependency on more than one protection group.

♦ ♦ ♦ CHAPTER 2

# Installing and Configuring the Disaster Recovery Framework Software

This chapter describes the steps for enabling and configuring the Disaster Recovery framework. This chapter contains the following sections:



## Installation Overview

You can install Disaster Recovery framework software on a running cluster without disruption. Because the Disaster Recovery framework installation process does not require you to restart Oracle Solaris Cluster software, the cluster remains in production with services running.

---

**Note -** Ensure that you have installed all the required software updates for your cluster configuration on each node of every cluster before you start installing the software. See "Updating to a New Oracle Solaris Cluster Version" in *Updating Your Oracle Solaris Cluster 4.4 Environment* for installation instructions.

---

For zone clusters that are already created, when you install the Disaster Recovery framework software, the software is propagated to the zone-cluster nodes by default. If you create a zone cluster after Disaster Recovery framework software is installed in the global cluster, you must install the Disaster Recovery framework software in the new zone-cluster nodes.

This section contains the following lists of tasks to perform to create a Disaster Recovery framework configuration:



# Prerequisite Configuration Tasks

Before you begin administering the Disaster Recovery framework, you must identify the Oracle Solaris Cluster installations you need to host protection groups. Then, you need to adjust the Oracle Solaris Cluster configuration and environment to support the formation of partnerships and protection groups with the Disaster Recovery framework. The following table describes these prerequisite tasks.

**TABLE 4** Prerequisite Configuration Tasks

| Task | Description |
|---|---|
| 1. Set the cluster name to the name you want to use with the Disaster Recovery framework. | Use the `cluster` command. For more information about this requirement, see "How to Enable the Disaster Recovery Framework" on page 40. |
| 2. Set up the IP address and host maps for the cluster that is enabled to run Geographic Edition software. | See Chapter 2, "Installing Software on Global-Cluster Nodes" in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*. |
| 3. Install and configure your data replication product. **Note -** This step is required before you can create protection groups with the `geopg create` command. | See the Oracle Solaris Cluster data replication guide for the product you are using. A list of available guides is provided in "Cluster Protection Groups" on page 25. |
| 4. Port and configure application configuration and corresponding resource groups on clusters that are candidates for partnership. | See the guidelines and prerequisites in "Creating a Partnership" on page 45. |

# Installation and Configuration Tasks

After you have completed the prerequisite configuration tasks, you can install and configure the Disaster Recovery framework as described in the following table.

**TABLE 5** Disaster Recovery Framework Installation and Configuration Tasks

| Task | Description and Documentation |
|---|---|
| 1. Install Disaster Recovery framework software. | See "Installing Disaster Recovery Framework Software" on page 31. |

| Task | Description and Documentation |
|---|---|
| 2. Set up security between the candidate partner clusters. | ■ Exchange certificates, as described in "Security Certificates" on page 21.<br>■ (Optional) Configure a secure logical hostname that uses IP Security Architecture (IPsec), as described in "IP Security (IPsec)" on page 21. |
| 3. If using a zone cluster as a partner, prepare the zone cluster for membership. | See "Preparing a Zone Cluster for Partner Membership" on page 37. |
| 4. Enable the Disaster Recovery framework software. | Issue the `geoadm start` command. For more information, see "Enabling the Disaster Recovery Framework Infrastructure" on page 40. |
| 5. Create partnerships. | See "Configuring a Partnership" on page 43. |
| 6. Configure data replication. | See the Oracle Solaris Cluster data replication guide for the product you use. A list of available guides is provided in "Cluster Protection Groups" on page 25. |
| 7. Create and validate protection groups. | See the Oracle Solaris Cluster guide for the data replication product you use.<br><br>To create a protection group that does not require data replication, see "Creating a Protection Group That Does Not Require Data Replication" on page 51. |
| 8. Add data replication device groups and application resource groups to the protection group. | See the Oracle Solaris Cluster guide for the data replication product you use. |
| 9. Bring online (activate) the protection groups. | See "How to Activate a Protection Group" on page 57. |
| 10 (Optional) Create sites and multigroups. | Set up sets of clusters and protection groups on which to perform switchover or takeover in a single operation. For more information, see "Configuring Sites and Multigroups" on page 60. |
| 11. Test the configured partnership and protection groups to validate the setup. | Perform a trial switchover or takeover and test some simple failure scenarios. See Chapter 11, "Migrating Services" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4* and procedures to migrate services in the Oracle Solaris Cluster guide for your data replication product.<br>**Note -** You cannot perform personality swaps if you are running EMC Symmetrix Remote Data Facility/Asynchronous data replication. |

# Installing Disaster Recovery Framework Software

You must install Disaster Recovery framework software on every node of each cluster in your geographically separated cluster by using the `pkg add` command.

# How to Install Disaster Recovery Framework Software

This procedure explains how to install Disaster Recovery framework software. Perform the procedure in the global zone for each node of a global cluster or zone cluster that you are configuring in a partnership.

**Before You Begin**

Before you begin to install software, make the following preparations:

- Ensure that the Oracle Solaris OS is installed to support Disaster Recovery framework software.

  If Oracle Solaris software is already installed on the node, you must ensure that the Oracle Solaris installation meets the requirements for Disaster Recovery framework software and any other software that you intend to install on the cluster.

  ---

  **Note -** If you want to use the Oracle Solaris Cluster Manager browser interface to administer Disaster Recovery framework components, ensure that all cluster nodes have the same root password.

  ---

- Read Chapter 1, “Planning the Disaster Recovery Framework Installation”.
- Read the following guides, which contain information that can help you plan your configuration and prepare your installation strategy:
  - *Oracle Solaris Cluster 4.4 Release Notes* – Restrictions, bug workarounds, and other late-breaking information.
  - *Disaster Recovery Framework Concepts for Oracle Solaris Cluster 4.4*.
  - Documentation for all third-party software products.

1. **Become the `root` role in the global zone of the node where you intend to run the Disaster Recovery framework.**

   ---

   **Note -** The Disaster Recovery framework software must be installed in the global zone for all nodes of each cluster in the partnership, whether the partner cluster is a global cluster or a zone cluster. For a zone cluster that will be configured in a partnership, Disaster Recovery framework software must be installed in both the zone cluster nodes and on the underlying global cluster nodes.

   ---

2. **Set up the repository for the Oracle Solaris Cluster software packages.**

   - **If the cluster nodes have direct access or web proxy access to the Internet, perform the following steps.**

a. **Go to `https://pkg-register.oracle.com`.**

b. **Choose `Oracle Solaris Cluster software`.**

c. **Accept the license.**

d. **Request a new certificate by choosing `Oracle Solaris Cluster software` and submitting a request.**

The certification page is displayed with download buttons for the key and the certificate.

e. **Download the key and certificate files and install them as described in the returned certification page.**

f. **Configure the ha-cluster publisher with the downloaded SSL keys and set the location of the Oracle Solaris Cluster 4.4 repository.**

In the following example the repository name is https://pkg.oracle.com/solaris/cluster/.

```
# pkg set-publisher -p \
-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.key.pem \
-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.certificate.pem \
```

`-k /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.key.pem`

Specifies the full path to the downloaded SSL key file.

`-c /var/pkg/ssl/Oracle_Solaris_Cluster_4.4.certificate.pem`

Specifies the full path to the downloaded certificate file.

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

- **If you are using an ISO image of the software, perform the following steps.**

a. **Download the Oracle Solaris Cluster 4.4 ISO image from Oracle Software Delivery Cloud at `https://edelivery.oracle.com/`.**

**Note -** A valid Oracle license is required to access Oracle Software Delivery Cloud.

Oracle Solaris Cluster software, which includes Disaster Recovery framework software, is part of the Oracle Solaris Product Pack. Follow online instructions to complete selection of the media pack and download the software.

**b. Make the Oracle Solaris Cluster 4.4 ISO image available.**

```
# lofiadm -a path-to-iso-image
/dev/lofi/N
# mount -F hsfs /dev/lofi/N /mnt
```

-a *path-to-iso-image*

Specifies the full path and file name of the ISO image.

**c. Set the location of the Oracle Solaris Cluster 4.4 package repository.**

```
# pkg set-publisher -p file:///mnt/repo
```

**3. Ensure that the `solaris` and `ha-cluster` publishers are valid.**

```
# pkg publisher
PUBLISHER                             TYPE     STATUS   P  LOCATION
solaris                               origin   online   F  ha-cluster-repository
solaris                               origin   online   F  solaris-repository
ha-cluster                            origin   online   F  ha-cluster-repository
```

For information about setting the `solaris` publisher, see "Adding, Modifying, or Removing Package Publishers" in *Updating Systems and Adding Software in Oracle Solaris 11.4*.

---

**Tip -** Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created. The `-v` option also shows any release notes that apply to this particular install or update operation.

---

If you do not get any error messages when you use the `-nv` options, run the command again without the `-n` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI to get more information to help you diagnose and fix the problem. For troubleshooting information, see Appendix A, "Troubleshooting Package Installation and Update," in *Updating Systems and Adding Software in Oracle Solaris 11.4*.

**4. Install the Disaster Recovery framework 4.4 software.**

```
# /usr/bin/pkg install ha-cluster-geo-full
```

5. **Verify that the package installed successfully.**

   Output is similar to the following example, which checks the installation state of the `ha-cluster-geo-full` group package.

```
% pkg info ha-cluster/group-package/ha-cluster-geo-full
Name: ha-cluster/group-package/ha-cluster-geo-full
Summary: Oracle Solaris Cluster Disaster Recovery framework full group package
Description: Oracle Solaris Cluster Disaster Recovery framework full group package
Category: Meta Packages/Group Packages
State: Installed
Publisher: ha-cluster
Version: 4.4 (Oracle Solaris Cluster Disaster Recovery Framework 4.4.0.10.0)
Branch: 0.10.0
Packaging Date: Fri Sep 23 10:44:06 2017
Last Install Time: Mon Oct 24 21:25:53 2017
Size: 77.00 B
FMRI: pkg://ha-cluster/ha-cluster/group-package/ha-cluster-geo-full@version:dateTtimeZ
```

6. **Repeat this procedure on each node of each partner cluster.**

**Next Steps** Install any required software updates. Go to Chapter 3, “Updating Disaster Recovery Framework Software”.

Configure the Disaster Recovery framework on the clusters. Go to Chapter 2, “Installing and Configuring the Disaster Recovery Framework Software”.

# Securing Disaster Recovery Framework Software

This section provides procedures to configure IPsec to secure communication between partner clusters.

For additional information about configuring secure communication between partner clusters, see “Planning Security” on page 19.

## ▼ How to Configure IPsec for Secure Cluster Communication

The following example procedure configures a cluster, `cluster-paris`, for IPsec secure communication with another cluster, `cluster-newyork`. The procedure assumes that the local logical hostname on `cluster-paris` is `lh-paris-1` and that the remote logical hostname is `lh-`

`newyork-1`. Inbound messages are sent to `lh-paris-1` and outbound messages are sent to `lh-newyork-1`.

Perform the following procedure on each node of `cluster-paris`.

1. **Log in to the first node of the primary cluster, `phys-paris-1`, as the `root` role.**

   For a reminder of which node is `phys-paris-1`, see "Example Disaster Recovery Framework Cluster Configuration" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

2. **Set up an entry for the local address and remote address in the IPsec policy file.**

   The policy file is located at `/etc/inet/ipsecinit.conf`. Permissions on this file should be `644`. For more information about this file, see the `ipsecconf(8)` man page.

   For information about the names and values that are supported by the Disaster Recovery framework, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

   a. **Configure the communication policy.**

      The default port for the `tcp_udp` plug-in is 2084. You can specify this value in the `etc/cacao/instances/default/modules/com.sun.cluster.geocontrol.xml` file.

      The following entry in the `/etc/inet/ipsecinit.conf` file configures a policy with no preference for authorization or encryption algorithms.

      ```
      # {raddr lh-newyork-1 rport 2084} ipsec {auth_algs any encr_algs any \
      sa shared} {laddr lh-paris-1 lport 2084} ipsec {auth_algs any encr_algs \
      any sa shared}
      ```

      When you configure the communication policy on the secondary cluster, `cluster-newyork`, you must reverse the policies.

      ```
      # {laddr lh-newyork-1 lport 2084} ipsec {auth_algs any encr_algs \
      any sa shared} {raddr lh-paris-1 rport 2084} ipsec {auth_algs any encr_algs \
      any sa shared}
      ```

   b. **Add the policy by rebooting the node or by running the following command.**

      ```
      # ipsecconf -a /etc/inet/ipsecinit.conf
      ```

3. **Set up encryption and authentication keys for inbound and outbound communication.**

   The communication file is located at `/etc/init/secret/ipseckeys`. Permissions on the file should be `600`.

Add keys:

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

Key entries have the following general format:

```
# inbound to cluster-paris
add esp spi paris-encr-spi dst lh-paris-1 encr_alg paris-encr-algorithm \
encrkey paris-encrkey-value
add ah spi newyork-auth-spi dst lh-paris-1 auth_alg paris-auth-algorithm \
authkey paris-authkey-value

# outbound to cluster-newyork
add esp spi newyork-encr-spi dst lh-newyork-1 encr_alg newyork-encr-algorithm \
encrkey newyork-encrkey-value
add ah spi newyork-auth-spi dst lh-newyork-1 auth_alg newyork-auth-algorithm \
authkey newyork-authkey-value
```

For more information about the communication files, see the `ipsecconf(8)` man page.

**Next Steps** If you are configuring a zone cluster as a member of a partnership, go to "Preparing a Zone Cluster for Partner Membership" on page 37.

Otherwise, go to "Enabling the Disaster Recovery Framework Infrastructure" on page 40.

# Preparing a Zone Cluster for Partner Membership

To enable a zone cluster to function as a member of a Disaster Recovery framework partnership, the common agent container must be manually configured within the zone cluster.

## ▼ How to Prepare a Zone Cluster for Partner Membership

This procedure configures common agent container security in a zone cluster to prepare the zone cluster for use in a cluster partnership.

**Before You Begin** Ensure that the following conditions are met:

- The zone cluster is created. See "Creating and Configuring a Zone Cluster" in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

- You have read the requirements for using a zone cluster in a cluster partnership. See "Zone Clusters" on page 23.
- Disaster recovery framework software is installed in the global cluster that supports the zone cluster you are configuring.

1. **Assume the `root` role on a node of the global cluster that supports the zone cluster you are configuring.**

2. **Set up the network address for the zone cluster.**

```
phys-schost# clzonecluster configure zone-cluster-name
clzc:zone-cluster-name> add net
clzc:zone-cluster-name:net> set address=zone-cluster-name
clzc:zone-cluster-name:net> end

clzc:zone-cluster-name> verify
clzc:zone-cluster-name> commit
clzc:zone-cluster-name> exit
```

3. **Copy the security files for the common agent container to all zone cluster nodes.**

   This step ensures that security files for the common agent container are identical on all cluster nodes and that the copied files retain the correct file permissions.

   Perform all steps in the zone cluster.

   a. **Log in to each node of the zone cluster.**

```
phys-schost# zlogin zone-cluster-name
zcname#
```

   b. **On each node, stop the common agent container.**

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

   c. **On one node, create the security keys.**

```
zcname# cacaoadm create-keys --force
```

   d. **Create a tar file of the `/etc/cacao/instances/default/security` directory.**

```
zcname# cd /etc/cacao/instances/default
zcname# tar cf /tmp/SECURITY.tar ./security
```

   e. **Copy the `/tmp/SECURITY.tar` file to each of the other cluster nodes.**

**f. On each node to which you copied the `/tmp/SECURITY.tar` file, extract the security files.**

Any security files that already exist in the `/etc/cacao/instances/default/security` directory are overwritten.

```
zcname# cd /etc/cacao/instances/default
zcname# tar xf /tmp/SECURITY.tar
```

**g. Delete the `/tmp/SECURITY.tar` file from each node in the cluster.**

You must delete each copy of the tar file to avoid security risks.

```
zcname# rm /tmp/SECURITY.tar
```

**h. On each node, set the common agent container network-bin address.**

```
zcname# cacaoadm set-param network-bind-address=0.0.0.0
```

**i. On each node, enable and start the common agent container.**

```
zcname# /usr/sbin/cacaoadm enable
zcname# /usr/sbin/cacaoadm start
```

**4. Verify that the Disaster Recovery framework modules are loaded on the zone cluster node.**

```
phys-schost# cacaoadm status com.sun.cluster.geocontrol
phys-schost# cacaoadm status com.sun.cluster.geoutilities
phys-schost# cacaoadm status com.sun.cluster.notifier
```

- If a module is loaded, command output would be similar to the following. You can safely ignore the message `Module is not in good health`.

```
Operational State:ENABLED
Administrative State:LOCKED
Availability Status:[]
Module is not in good health.
```

- If a module is not loaded, command output would be similar to the following.

```
Module com.sun.cluster.geocontrol has not been loaded.
Cause of the problem:[DEPENDENCY]
```

See the Troubleshooting section at the end of this procedure.

**5. Exit the zone cluster node.**

```
zcname# exit
phys-schost#
```

**Troubleshooting** If a Disaster Recovery framework module is not loaded, check that the zone cluster configuration is correct.

After you have verified that the configuration is complete and correct, and you have fixed any errors, do one of the following:

- On each zone cluster node, restart the common agent container.

  ```
  zcnode# /usr/sbin/cacaoadm restart
  ```

- From a global-cluster node, reboot the zone cluster.

  ```
  phys-schost# clzonecluster reboot zone-cluster-name
  ```

After processing is complete on all zone cluster nodes, check that the Disaster Recovery framework modules are now loaded. If any modules are still not loaded, contact your Oracle service representative for assistance.

**Next Steps** Go to "Enabling the Disaster Recovery Framework Infrastructure" on page 40.

# Enabling the Disaster Recovery Framework Infrastructure

When Disaster Recovery framework software is enabled, the cluster is ready to enter a partnership with another enabled cluster.

For more information about setting up and installing the Disaster Recovery framework, see Chapter 3, "Administering the Disaster Recovery Framework" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

## ▼ How to Enable the Disaster Recovery Framework

This procedure enables the Disaster Recovery framework infrastructure on the local cluster only. Repeat this procedure on all the clusters of your geographically separated cluster.

**Before You Begin** Ensure that the following conditions are met:

- The cluster is running the Oracle Solaris OS and Oracle Solaris Cluster software.

- If you want to use the Oracle Solaris Cluster Manager browser interface to administer Disaster Recovery framework components, ensure that all cluster nodes have the same root password.
- The Oracle Solaris Cluster management-agent container for Oracle Solaris Cluster Manager is running.
- The Disaster Recovery framework software is installed.
- The cluster has been configured for secure cluster communication by using security certificates, that is, nodes within the same cluster must share the same security certificates. This is configured during Oracle Solaris Cluster installation.

1. **Assume the `root` role on a cluster node.**

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

```
# cluster list
```

For global clusters, if the cluster name is not the name that you want to use, change the cluster name with the following command:

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

`-c` *new-cluster-name*

Specifies the new cluster name.

*cluster-name*

The cluster whose name you are changing.

For more information, see the `cluster`(8CL) man page.

---

**Note -** After you have enabled the Disaster Recovery framework infrastructure, you must not change the cluster name while the infrastructure is enabled.

---

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

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

4. **On a node of the cluster, start the Disaster Recovery framework infrastructure.**

```
# geoadm start
```

The `geoadm start` command enables the Disaster Recovery framework infrastructure on the local cluster only. For more information, see the `geoadm`(8) man page.

5. **Verify that you have enabled the infrastructure and that the Disaster Recovery framework resource groups are online.**

```
# geoadm show
# clresourcegroup status geo-clusterstate geo-infrastructure
# clresource status -g geo-clusterstate,geo-infrastructure
```

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

The output for the `clresourcegroup status` and `clresource status` commands display that the `geo-failovercontrol`, `geo-hbmonitor`, and `geo-clustername` resources and the `geo-infrastructure` resource group is online on one node of the cluster. The `geo-clusterstate` resource group is online on both nodes.

For more information, see the `clresourcegroup`(8CL) and `clresource`(8CL) man pages.

**Example 1** Enabling the Disaster Recovery Framework Infrastructure on a Cluster

This example enables the Disaster Recovery framework on the `cluster-paris` cluster.

1. Start the Disaster Recovery framework infrastructure on `cluster-paris`.

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

2. Ensure that the Disaster Recovery framework infrastructure was successfully enabled.

```
phys-paris-1# geoadm show

--- CLUSTER LEVEL INFORMATION ---
Oracle Solaris Cluster Disaster Recovery framework is active on cluster-paris from
 node phys-paris-1
Command execution successful
phys-paris-1#
```

3. Verify the status of the Disaster Recovery framework resource groups and resources.

```
phys-paris-1# clresourcegroup status geo-clusterstate geo-infrastructure

=== Cluster Resource Groups ===

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

geo-infrastructure            phys-paris-1          No                  Online
```

```
                                phys-paris-2          No                  Offline

phys-paris-1# clresource status -g geo-clusterstate,geo-infrastructure

=== Cluster Resources ===

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

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

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

**Next Steps** Configure trust between partner clusters. Go to "How to Configure Trust Between Two Clusters" on page 43.

# Configuring a Partnership

This section provides the following Information:

- "Configuring Trust Between Partner Clusters" on page 43
- "Creating a Partnership" on page 45
- "Joining an Existing Partnership" on page 48

## Configuring Trust Between Partner Clusters

This section provides procedures to configure secure communication, or trust, between the two clusters you want to be in a partnership.

### ▼ How to Configure Trust Between Two Clusters

Before you create a partnership between two clusters, you must configure the Disaster Recovery framework for secure communication between the two clusters. The configuration must be

reciprocal. For example, you must configure the cluster `cluster-paris` to trust the cluster `cluster-newyork`, and you must also configure the cluster `cluster-newyork` to trust the cluster `cluster-paris`.

---

**Note -** You can also perform this task by using the Oracle Solaris Cluster Manager browser interface. Click Partnerships, then click Add Partner Trust. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

---

**Before You Begin** Ensure that the following conditions are met:

- The cluster on which you want to create the partnership is running.
- The `geoadm start` command has already been run on this cluster and the partner cluster. For more information about using the `geoadm start` command, see "Enabling the Disaster Recovery Framework Infrastructure" on page 40.
- The cluster name of the partner cluster is known.
- The host information of the partner cluster is defined in the local hosts file. The local cluster needs to know how to reach the partner cluster by name.

1. **Assume the `root` role on a cluster node.**

2. **Import the public keys from the remote cluster to the local cluster.**
Run the following command on one node of the local cluster to import the keys from the remote cluster to one node of the cluster.

   *local-cluster*`#` **`geops add-trust -c`** *remote-cluster*

   `-c` *remote-cluster*

   Specifies the logical hostname of the cluster with which to form a partnership. The logical hostname is used by the Disaster Recovery framework and maps to the name of the remote partner cluster. For example, a remote partner cluster name might resemble the following:

   `cluster-paris`

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

   *remote-cluster*`.certificate[0-9]*`

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

   For more information about the `geops` command, refer to the `geops(8)` man page.

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

4. **Verify trust from one node of each cluster.**

**Note -** You can also accomplish this step by using the Oracle Solaris Cluster Manager browser interface. Click Partnerships, then click Verify Partner Trust. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

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

**Next Steps** Configure the partnership. Go to "Creating a Partnership" on page 45.

**See Also** To remove trust, see "Removing Trust Between Partner Clusters" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

# Creating a Partnership

This section provides procedures to create a Disaster Recovery framework partnership between two clusters:

## ▼ How to Create a Partnership

**Note -** You can also accomplish this procedure by using the Oracle Solaris Cluster Manager browser interface. Click Partnerships, and then click Create. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

**Before You Begin** Ensure that the following conditions are met:

- The cluster on which you want to create the partnership is up and running.
- The `geoadm start` command must have already been run on the this cluster and the partner cluster. For more information about using the `geoadm start` command, see "Enabling the Disaster Recovery Framework Infrastructure" on page 40.
- The cluster name of the partner cluster is known.
- The host information of the partner cluster must defined in the local host file. The local cluster needs to know how to reach the partner cluster by name.
- Security has been configured on the two clusters by installing the appropriate certificates. See "Configuring Trust Between Partner Clusters" on page 43 for more information.

1. **Log in to a cluster node.**

   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see "Disaster Recovery Framework Rights Profiles" on page 19.

2. **Create the partnership.**

   *local-partner-cluster*`# geops create -c` *remote-partner-cluster*`[.`*domain-name*`] [-h` *heartbeat*`] \`
   `[-p` *property-setting* `[-p…]]` *partnership*

   `-c` *remote-partner-cluster*[.*domain-name*]

   Specifies the name of the remote cluster that will participate in the partnership. If clusters in the partnership are in different domains, you must also specify the domain name of the remote cluster.

   This name matches the logical hostname used by the Disaster Recovery framework infrastructure on the remote cluster.

   `-h` *heartbeat*

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

   If you omit this option, the default Disaster Recovery framework heartbeat is used.

   Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Oracle specialist for assistance if your system requires the use of custom heartbeats. For more information about configuring custom heartbeats, see Chapter 6, "Administering Heartbeats" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

   **Note -** A custom heartbeat prevents the default heartbeat from being used during partnership creation. If you want to use the default heartbeat for your partnership, you must delete the custom heartbeat before you run the `geops create` command.

   `-p` *property-setting*

   Specifies the value of partnership properties with a string of *property=value* pair statements.

   Specify a description of the partnership with the `Description` property.

   You can configure heartbeat-loss notification with the `Notification_emailaddrs` and `Notification_actioncmd` properties. For more information about configuring heartbeat-

loss notification, see "Configuring Heartbeat-Loss Notification" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

For more information about the properties you can set, see Appendix A, "Standard Disaster Recovery Framework Properties," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

*partnership*

Specifies the name of the partnership.

For information about the names and values that are supported by the Disaster Recovery framework, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

For more information about the `geops` command, refer to the `geops`(8) man page.

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

**Note -** You can also accomplish this step by using the Oracle Solaris Cluster Manager browser interface. Click Partnerships to view partnership information. For additional details, click the partnership name. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

The partnership states will be Degraded and the heartbeat state will be Offline. These states will change after the partnership is joined from the partner cluster.

*local-partner-cluster#* **`geoadm status`**

**Example 2** Creating a Partnership

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

```
cluster-paris.usa1# geops create -c cluster-newyork.usa \
-p Description=Transatlantic \
-p Notification_emailaddrs=sysadmin@example.com \
paris-newyork-ps
```

**Next Steps** To finalize the new partnership, the remote partner cluster must join the partnership. Go to "Joining an Existing Partnership" on page 48.

**See Also** To remove a partnership between two clusters, see "How to Remove Trust Between Two Clusters" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

# Joining an Existing Partnership

When you define and configure a partnership, the partnership specifies a second cluster to be a member of that partnership. Then, you must configure this second cluster to join the partnership.

## ▼ How to Join a Partnership

**Note -** You can also accomplish this procedure by using the Oracle Solaris Cluster Manager browser interface. Click Partnerships, then click Join Partnership. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

**Before You Begin** Ensure that the following conditions are met:

- The local cluster is enabled to run the Disaster Recovery framework.
- The partnership you want the cluster to join is defined and configured on another cluster (`cluster-paris`) and the local cluster (`cluster-newyork`) is specified as a member of this partnership. See "Creating a Partnership" on page 45.
- Security has been configured on the clusters by installing the appropriate certificates.

  See "Security Certificates" in *Installing and Configuring the Disaster Recovery Framework for Oracle Solaris Cluster 4.4* for more information.

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

   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see "Disaster Recovery Framework Rights Profiles" on page 19.

2. **Confirm that the remote cluster that originally created the partnership, `cluster-paris`, can be reached at its logical hostname.**

   *local-partner-cluster#* **`ping lh-paris-1`**

   For information about the logical hostname of the cluster, see "How to Enable the Disaster Recovery Framework" on page 40.

3. **Join the partnership.**

   *local-partner-cluster#* **`geops join-partnership`** `[`**`-h`** *heartbeat*`]` *remote-partner-cluster partnership*

-h *heartbeat*

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

If you omit this option, the default Disaster Recovery framework heartbeat is used.

Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Oracle specialist for assistance if your system requires the use of custom heartbeats. For more information about configuring custom heartbeats, see Chapter 6, "Administering Heartbeats" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

*remote-partner-cluster*

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

*partnership*

Specifies the name of the partnership.

For information about the names and values that are supported by Disaster Recovery framework software, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

For more information about the `geops` command, refer to the `geops`(8) man page.

4. **Verify that the cluster was added to the partnership and that the partnership properties were defined correctly.**

```
local-partner-cluster# geops list
local-partner-cluster# geoadm status
```

**Example 3** Joining a Partnership

This example joins the `cluster-newyork` cluster to the `paris-newyork-ps`partnership.

```
phys-newyork-1# geops join-partnership cluster-paris paris-newyork-ps
phys-newyork-1# geops list
phys-newyork-1# geoadm status
```

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

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

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

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

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

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

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

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

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

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

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

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

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

6. Verify that the partnership has been created successfully.

```
# geops list
# geoadm status
```

**Next Steps** Configure protection groups. See “Configuring Protection Groups” on page 51 and the Oracle Solaris Cluster guide for the data replication product you will use.

# Configuring Protection Groups

This section contains the following information:



Also see the appropriate Oracle Solaris Cluster guide for procedures to create a protection group for your data replication product.

## Creating a Protection Group That Uses Data Replication

**Note -** If you do not need to use data replication, see "Creating a Protection Group That Does Not Require Data Replication" on page 51.

The procedures to configure a protection group that uses data replication vary, depending on the data replication product you use. See the appropriate Oracle Solaris Cluster guide for your data replication product for guidelines and procedures to configure a protection group:

- *Oracle Solaris Cluster Data Replication Guide for MySQL*
- *Oracle Solaris Cluster Data Replication Guide for Oracle Data Guard*
- *Oracle Solaris Cluster Data Replication Guide for ZFS Snapshots*
- *Oracle Solaris Cluster Remote Replication Guide for Oracle ZFS Storage Appliance*

After you create the protection group and add application resource groups and data-replicated components, validate the protection group. Go to "Validating a Protection Group" on page 56.

## Creating a Protection Group That Does Not Require Data Replication

Some protection groups do not require data replication. If you are using the Disaster Recovery framework to manage only resource groups, you can create protection groups that do not replicate data.

This section provides the following procedures:

- "How to Create a Protection Group That Is Configured Not to Use Data Replication" on page 52
- "How to Add an Application Resource Group to a Protection Group That Does Not Use Data Replication" on page 54

**Note -** You cannot add device groups to a protection group that does not use data replication.

To create a protection group that uses data replication, see the appropriate Oracle Solaris Cluster guide for your data replication product:

- *Oracle Solaris Cluster Data Replication Guide for MySQL*
- *Oracle Solaris Cluster Data Replication Guide for Oracle Data Guard*
- *Oracle Solaris Cluster Data Replication Guide for ZFS Snapshots*
- *Oracle Solaris Cluster Remote Replication Guide for Oracle ZFS Storage Appliance*

## ▼ How to Create a Protection Group That Is Configured Not to Use Data Replication

**Note -** You can also accomplish this procedure by using the Oracle Solaris Cluster Manager browser interface. Click Partnerships, click the partnership name to go to its page, and click Create in the Protection Groups section. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

**Before You Begin** Before you create a protection group without data replication, ensure that the following conditions are met:

- The local cluster is a member of a partnership.
- The protection group that you are creating does not already exist.

**Note -** Protection group names are unique in the global Disaster Recovery framework namespace. You cannot use the same protection group name in more than one partnership on the same system.

1. **Log in to a cluster node.**
   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see "Disaster Recovery Framework Rights Profiles" on page 19.

2. **Create a new protection group by using the `geopg create` command.**

This command creates a protection group on the local cluster.

```
# geopg create -s partnership -o local-role \
[-p property [-p …]] protection-group
```

`-s` *partnership*

Specifies the name of the partnership.

`-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.

`External_Dependency_Allowed`

Specifies whether to allow any dependencies between resource groups and resources that belong to this protection group and resource groups and resources that do not belong to this protection group.

`RoleChange_ActionArgs`

Specifies a string that follows system-defined arguments at the end of the command line when the role-change callback command runs.

`RoleChange_ActionCmd`

Specifies the path to an executable command. This path should be valid on all nodes of all partner clusters that can host the protection group. The script is invoked during a switchover or takeover on the new primary cluster when the protection group is started on the new primary cluster. The script is invoked on the new primary cluster after the data replication role changes from secondary to primary and before the application resource groups are brought online. If the data replication role change does not succeed, then the script is not called.

`Timeout`

Specifies the timeout period for the protection group in seconds. You can change the timeout period from the default value depending on the complexity of your data replication configuration.

For more information about the properties you can set, see Appendix A, "Standard Disaster Recovery Framework Properties," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

*protection-group*

Specifies the name of the protection group.

3. Log in to a node on the partner cluster to replicate the protection group to the partner cluster.

```
# geopg get --partnership <partnership-name> <protection group name>
```

For information about the names and values that are supported by Disaster Recovery framework software, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

For more information about the `geopg` command, refer to the `geopg`(8) man page.

**Example 5** Creating and Configuring a Protection Group That Is Configured to Not Use Data Replication

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

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

**Next Steps** Go to "How to Add an Application Resource Group to a Protection Group That Does Not Use Data Replication" on page 54.

**See Also** To delete a protection group, see "Deleting Protection Groups and Data Replication Components" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

## ▼ How to Add an Application Resource Group to a Protection Group That Does Not Use Data Replication

1. **Log in to a cluster node.**
You must be assigned the Geo Management rights profile to complete this procedure. For more information, see "Disaster Recovery Framework Rights Profiles" on page 19.

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_Dependency_Allowed` property of the protection group is set to `TRUE`.**

   ```
   # geopg show protection-group | grep -i external_dependency_allowed
   ```

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

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

4. **Start the protection group or change the state of the application resource group to a state that is required for the addition to be allowed.**

   - The Disaster Recovery framework requires that the application resource group be in the `UNMANAGED` state on the secondary cluster.
   - If the protection group is stopped on the primary cluster, the application resource group must also be unmanaged on the primary cluster.
   - If the protection group is active on the primary cluster, the application resource group must be in the `UNMANAGED` or `ONLINE` state on the primary cluster.

   For instructions, see one of the following procedures:

   - "How to Disable a Resource and Move Its Resource Group Into the UNMANAGED State" in *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*
   - "How to Bring Resource Groups Online" in *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.

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

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

   | | |
   |---|---|
   | *application-resource-group* | Specifies the name of an 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. The command adds an application resource group to a protection group on the local cluster. Then, if the partner cluster contains a protection group of the same name, the command propagates the new configuration information to the partner cluster. |

   For information about the names and values that are supported by the Disaster Recovery framework, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

# Validating a Protection Group

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

## ▼ How to Validate a Protection Group

This procedure 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.

---

**Note -** You can also accomplish this procedure by using the Oracle Solaris Cluster Manager browser interface. Click Partnerships, click the partnership name to go to its page, highlight the protection group name, and click Validate. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

---

**Before You Begin** Ensure that the following conditions are met:

- The protection group you want to validate exists locally.
- The common agent container is online on all nodes of both clusters in the partnership.

1. **Log in to one of the cluster nodes.**

   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see "Disaster Recovery Framework Rights Profiles" on page 19.

2. **Validate the configuration of the protection group.**

   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.

   ```
   # geopg validate protection-group
   ```

*protection-group*

Specifies a unique name that identifies a single protection group

- 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 an error message and the configuration remains in the `Error` state. Fix the error in the configuration, then rerun the `geopg validate` command.

**Next Steps** Go to "Activating a Protection Group" on page 57.

# Activating a Protection Group

When configuration of a protection group is complete, activate the protection group to put its configuration into service.

## ▼ How to Activate a Protection Group

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

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

   For more information, see "Disaster Recovery Framework Rights Profiles" on page 19.

   ---

   **Note -** If you use a role with Geo Management 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 rights profile and data replication software.

   ---

2. **Activate the protection group.**

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

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

`-e` *scope*

Specifies the scope of the command.

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

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

`-n`

Prevents the start of replication at protection group startup.

If you omit this option, the replication subsystem starts at the same time as the protection group. In addition, the following data replication products have additional behaviors when the `-n` option is omitted:

- **MySQL** – The `geopg start` command performs the following actions, if the role of the protection group is secondary on the local cluster:
  - Starts the MySQL slave threads
  - Prevents modification by non-`root` roles if this option is configured
  - Prepares the `my.cnf` file to start the database with modifications prevented for non-`root` roles if this option is configured
- **Oracle Data Guard** – The `geopg start` command performs the following operations on each Oracle Data Guard broker configuration in the protection group:
  - Verifies that the resource group that is named in the `local_oracle_svr_rg_name` property contains a resource of type `SUNW.scalable_rac_server_proxy` for a scalable resource group or a resource of type `SUNW.oracle_server` for a failover resource group.
  - Verifies that the Oracle Data Guard `dgmgrl` command can connect using the values that are given for `sysdba_username`, `sysdba_password`, and `local_db_service_name`. Or if the `sysdba_username` and `sysdba_password` properties are null, verifies that the `dgmgrl` command can connect using the Oracle wallet connection format, `dgmgrl /@`*local_db_service_name*.
  - Verifies that the role configured for the replication resource is the same as the role of the protection group on the local cluster.
  - Verifies that the Oracle Data Guard broker configuration details match those that are held by the Disaster Recovery framework. The details to check include which cluster is primary, the configuration name, the database mode (for both the primary

and standby clusters), the replication mode, the standby type, that `FAST_START FAILOVER` is disabled, and that `BystandersFollowRoleChange` is equal to `NONE`.

*protection-group*

Specifies the name of the protection group.

The `geopg start` command uses the `clresourcegroup online -eM` *resource-group-list* command to bring resource groups and resources online. See the `clresourcegroup`(8CL) man page for more information.

If the role of the protection group is primary on the local cluster, the `geopg start` command performs the following operations:

- Runs a script that is defined by the `RoleChange_ActionCmd` property.
- Brings the application resource groups in the protection group online on the local cluster. For Oracle Data Guard, this includes the shadow Oracle database server resource groups.

The `geopg start` command also performs additional operations for MySQL data replication:

- Prepares the `my.cnf` file to start the database without the slave threads
- Brings online the application resource groups in the protection group on the local cluster

**Example 6** Globally Activating a Protection Group

This example globally activates a protection group.

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

**Example 7** Locally Activating a Protection Group

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

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

**Troubleshooting**

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

If the `Configuration` status is set to `Error`, revalidate the protection group by using the procedures that are described in "Validating a Protection Group" on page 56.

**Next Steps**

If you want to administer a set of protection groups as a single entity, go to "Configuring Sites and Multigroups" on page 60.

**See Also** To deactivate a protection group, see "Activating and Deactivating a Protection Group" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

# Configuring Sites and Multigroups

This section contains the following procedures:



## ▼ How to Create a Site

Perform this procedure to configure a new site.

---

**Note -** You can also use the Oracle Solaris Cluster Manager browser interface to perform this task. Click Sites, then click Create. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

---

**Before You Begin** Determine which clusters the site will contain and whether each cluster will be a site controller or a site member. The cluster from which you create the new site is automatically configured as a site controller. To avoid a possible single point of failure, configure at least two clusters as site controllers.

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

   For more information, see "Securing Disaster Recovery Framework Software" on page 35.

   ---

   **Note -** If you use a role with Geo Management 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 rights profile and data replication software.

   ---

2. **Ensure that all nodes in the cluster are online.**

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

--- Node Status ---

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

If any node is offline, wait until the node is back online before you create the new site. The creation of a new site, or the acceptance by invited members to join the site, will fail if any node in the issuing cluster is not online.

3. **Configure the new site.**

The issuing cluster is automatically configured as a site controller, so it is not necessary to specify that cluster name to the `geosite create` command. You can specify the `-c` option and the `-m` option in the same `geosite create` command.

*first-site-controller-cluster-node*`# geosite create [-c` *cluster*`[,…]] [-m` *cluster*`[,…]]` *site*

`-c` *cluster*

The name of a cluster to configure as a site controller. You can specify multiple cluster names, separated by commas (,).

`-m` *cluster*

The name of a cluster to configure as a site member. You can specify multiple cluster names, separated by commas (,).

*site*

The name to give the site that you are creating.

The command issues an invitation to each cluster that is specified to the `geosite create` command. No site-based operations are accepted from a cluster until the cluster accepts the invitation to join the site.

4. **For each cluster that was invited, accept the invitation to join the new site.**

   a. **Ensure that the common agent container is enabled on all nodes of this cluster and all nodes of the cluster that this cluster is joining.**

```
# /usr/lib/cacao/bin/cacaoadm status
```

**b. If the common agent container is not running on any of the cluster nodes, start it.**

```
# /usr/lib/cacao/bin/cacaoadm start
```

**c. From one node, join the site.**

*invited-cluster-node*# **geosite join** *first-site-controller-cluster site*

*first-site-controller-cluster*

The name of the cluster that issued the invitation to join the site.

**5. Verify the site configuration.**

```
# geosite status site
```

**Example 8** Creating a New Site

The following example creates a new site named `europe`. The issuing cluster, `london`, is automatically configured as a site controller. The cluster `madrid` is configured as a second site controller, and the clusters `berlin` and `paris` are configured as site members. The invited clusters accept the invitation from the `london` cluster to join the `europe` site.

```
phys-london-1# geosite create -c madrid -m berlin,paris europe

phys-madrid-1# geosite join london europe
phys-berlin-1# geosite join london europe
phys-paris-1# geosite join london europe
```

**Next Steps** Go to "How to Create a Multigroup" on page 62.

**See Also** To delete a site, see "Deleting a Site" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

# ▼ How to Create a Multigroup

Perform this procedure to configure a multigroup to manage designated sets of protection groups.

**Note -** You can also use the Oracle Solaris Cluster Manager browser interface to perform this task. Click Sites, click the site name to go to its page, and click Add. For Oracle Solaris Cluster Manager log-in instructions, see "How to Access Oracle Solaris Cluster Manager" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

**Before You Begin**

- Ensure that the protection groups you want to contain in the multigroup are configured and working properly. See the Oracle Solaris Cluster guide for your data replication product for procedures to configure a protection group.
- Ensure that a partner cluster for each protection group to configure in the multigroup is configured in the same site. See "How to Create a Site" on page 60.

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

   For more information, see "Disaster Recovery Framework Rights Profiles" on page 19.

   **Note -** If you use a role with Geo Management 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 rights profile and data replication software.

2. **Create the multigroup.**

   *site-controller-cluster-node#* **`geomg create -s`** *site multigroup*

   `-s` *site*

   The name of the site.

   *multigroup*

   The name to assign the new multigroup. The name must be unique throughout the specified site. If the name is not unique, the command fails with an error.

3. **From a node of a site controller cluster, add protection groups to the multigroup.**

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

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

*cluster*:*protection-group*

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

*cluster1*:*protection-group1*/*cluster2*:*protection-group2*

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

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

The comma (,) separates multiple protection group names in the same command.

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

Specifies that multiple protection groups, *cluster1*:*protection-group2* and *cluster2*:*protection-group1*, all have a dependency on the *cluster3*:*protection-group1* protection group. Parentheses can only be used to enclose multiple protection groups with a dependency on another, single protection group. Only one protection group can be specified as the depended-on protection group.

4. **Verify the multigroup configuration.**

```
# geomg status multigroup
```

**See Also** To delete a multigroup, see "Deleting a Multigroup" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

To administer a multigroup, see Chapter 9, "Administering Multigroups" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

♦ ♦ ♦ CHAPTER 3

# Updating Disaster Recovery Framework Software

This chapter describes how to update Disaster Recovery framework software to a new release, or install Software Repository Updates (SRU), in the global cluster or in an Oracle Solaris brand zone cluster.

**Note -** If you update Disaster Recovery framework software to a version that is more than one consecutive version different than the Disaster Recovery framework version running on the nodes of its partner cluster, you must also update the partner cluster nodes to a supported Disaster Recovery framework version. Do not start the Disaster Recovery framework on nodes of an updated cluster unless the version of Disaster Recovery framework software on each node of the partner cluster is no more than one consecutive version different.

If you are updating the Oracle Solaris Cluster software, the Disaster Recovery framework software is automatically updated at the same time.

## Updating a Disaster Recovery Framework Configuration

This section provides the following information to update a cluster to a new Disaster Recovery framework software version:

# Update Requirements and Software Support Guidelines

This section provides requirements and software-support guidelines for all clusters that have a partnership with the global cluster or Oracle Solaris brand zone cluster that you are updating.



## Requirements When Updating to a New Release

Observe the following requirements if you are updating your cluster to the Disaster Recovery framework 4.4 version.

- **Supported hardware** – The cluster hardware must be a supported configuration for Disaster Recovery framework 4.4 software. See the Oracle Solaris Cluster 4 Compatibility Guide for the Disaster Recovery framework configurations that are currently supported.
- **Minimum Oracle Solaris OS version** – Oracle Solaris 11.4 software is the minimum required to support Oracle Solaris Cluster 4.4..
- **Minimum Oracle Solaris Cluster version** – The cluster must run on or be updated to Oracle Solaris Cluster 4.3 software, before updating to Oracle Solaris Cluster 4.4 software. Updating the core Oracle Solaris Cluster software also updates the Disaster Recovery framework software.

---

**Note -** All clusters in a partnership must run either Oracle Solaris Cluster 4.3 or Oracle Solaris Cluster 4.4 software.

---

- **Supported Disaster Recovery framework versions in cluster partnerships** – All clusters that are in a partnership with the cluster that you are updating to Disaster Recovery framework 4.4 software must run either version 4.3 or 4.4 of the Disaster Recovery framework software. If any node on the partner cluster does not already run one of these versions of Disaster Recovery framework software, you must also update that node to a supported version before you restart the Disaster Recovery framework infrastructure on the updated cluster.

### Requirements When Updating to an SRU

Observe the following requirements if you are installing a software update of your Disaster Recovery framework 4.4 configuration.

- You must run the same software updates for Oracle Solaris Cluster software and the common agent container software on all nodes of the same cluster.
- Within a cluster, the software updates for each node on which you have installed Disaster Recovery framework software must meet the Oracle Solaris Cluster software update requirements.
- All nodes in the same cluster must have the same version of Disaster Recovery framework software and the same software updates. However, primary and secondary clusters can run different versions of Disaster Recovery framework software, provided that each version of the Disaster Recovery framework software is correctly updated and the versions are no more than one release different.
- To ensure that the updates have been installed properly, install the software updates on your secondary cluster before you install the software updates on the primary cluster.

## ▼ How to Update the Disaster Recovery Framework Software

Perform this procedure on each cluster node where you want the Disaster Recovery framework to run. Update the secondary cluster before you update the primary cluster, to permit testing. You can perform this procedure on more than one node at the same time.

**Caution -** The cluster in a partnership with the cluster you are updating must also be installed with Oracle Solaris Cluster 4.3 or 4.4 software before you can restart the Disaster Recovery framework 4.4 infrastructure on the updated cluster.

**Before You Begin** Perform this procedure on the cluster you are updating. Perform all steps from the global zone only.

1. **Prepare the cluster for an update.**

   a. **Ensure that the configuration meets the requirements for the update. See “Update Requirements and Software Support Guidelines” on page 66.**

   b. **Have available the installation media or the IPS publisher configured, documentation, and software updates for all software products that you are**

**updating, including Oracle Solaris OS, Oracle Solaris Cluster software, and Disaster Recovery framework software.**

c. **Ensure that you have installed all the required software updates for your cluster configuration on each node of the cluster before you start updating the software.**

d. **If you are installing a software update, ensure that you have read the `README` file for each software update you will install.**

e. **Availability Suite software is no longer supported in Oracle Solaris Cluster 4.4. You must remove Oracle Solaris Availability Suite protection groups and Oracle Solaris Availability Suite software before updating to Oracle Solaris Cluster 4.4 and Oracle Solaris 11.4. See "How to Uninstall Availability Suite Software" in *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Solaris Availability Suite*.**

2. **Ensure that the cluster is functioning properly.**

   a. **From any node, view the current status of the cluster.**

   ```
   % cluster status
   ```

   See the `cluster`(8CL) man page for more information.

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

3. **Run the update. See *Updating Your Oracle Solaris Cluster 4.4 Environment*.**

   The Disaster Recovery framework software is updated when the core Oracle Solaris Cluster 4.4 software is updated. Choose the appropriate update method for your environment.

4. **If you performed a Rolling Update of Oracle Solaris Cluster, after the update is complete, perform the following steps from one node of the cluster(s), global clusters and or Oracle Solaris brand zone clusters, where the disaster recovery framework is running:**

   ```
   # clrs delete -F geo-servicetag
   # clrs create -g geo-clusterstate -t SUNW.SCGeoInitSvc -y \
   R_description="Oracle Solaris Cluster disaster recovery framework initilization
    resource" geo-init-svc
   ```

**Next Steps** Go to "How to Verify Update of the Disaster Recovery Framework Software" on page 69.

# How to Verify Update of the Disaster Recovery Framework Software

Perform this procedure to verify that the cluster is successfully updated to Disaster Recovery framework 4.4 software.

**Before You Begin** Ensure that all update procedures are completed for all cluster nodes that you are updating.

1. **Login to a node in the global cluster or Oracle Solaris brand zone cluster where the Disaster Recovery framework is running and assume the `root` role.**

   If the Disaster Recovery framework is running in an Oracle Solaris brand zone cluster, log in to an Oracle Solaris brand zone cluster node.

2. **View the installed levels of Disaster Recovery framework software.**

   ```
   # pkg list ha-cluster-geo-incorporation
   ```

   **Note -** The version number that the `geo* -v` commands return does not always coincide with the marketing release version numbers.

3. **Repeat the preceding steps for each cluster node you updated.**

4. **Ensure that the cluster is running properly.**

   ```
   # geoadm status
   ```

5. **(Optional) Perform a switchover to ensure that Disaster Recovery framework software is installed properly.**

   ```
   # geopg switchover -m remote-cluster protection-group
   ```

   You must test your geographically separated cluster properly, so that no problems prevent a switchover. Updating only the secondary cluster first and switching over to it enables you to verify that switchover still works. If the switchover fails, the primary site is untouched and you can switch back. If switchover works on the secondary site, then after a certain 'soak time' you can update the primary site as well.

   **Note -** A switchover shuts down the applications in the protection group on the original primary cluster and migrates the applications to the secondary cluster. You should carefully plan the required tasks and resources before you perform a switchover.

♦ ♦ ♦ CHAPTER 4

# Uninstalling Disaster Recovery Framework 4.4 Software

This chapter describes how to uninstall the Disaster Recovery framework software.

When you uninstall Disaster Recovery framework 4.4 software, the node or cluster is no longer a part of the geographically separated cluster.

**Note -** You must uninstall Disaster Recovery framework software before you uninstall Oracle Solaris Cluster software.

## Uninstalling Disaster Recovery Framework Software

### ▼ How to Uninstall Disaster Recovery Framework Software

Use this procedure to uninstall Disaster Recovery framework software that was installed with the `pkg add` command. Remove Disaster Recovery framework software from all nodes in the cluster, unless you are removing the software from node that you are also removing from the cluster. You can continue to run applications during the uninstallation of Disaster Recovery framework software.

1. **Assume the `root` role on the node where you intend to uninstall Disaster Recovery framework software.**

2. **Unmanage data replication resource groups or remove the protection groups on the local cluster.**

Use one of the following methods, depending on whether you might want to reinstall Disaster Recovery framework software at a future time.

- **If you want to remove Disaster Recovery framework software from the cluster but retain the protection group configuration for possible future use, unmanage data replication resource groups for each protection group on the local cluster.**

  Unmanaging these resource groups prevents them from attempting to interact with Disaster Recovery framework software after the software is uninstalled.

  ```
  # clresourcegroup offline data-replication-resource-group
  # clresource disable -g data-replication-resource-group +
  # clresourcegroup unmanage data-replication-resource-group
  ```

  | | |
  |---|---|
  | `-g` | Specifies the data replication resource group to disable. |
  | `+` | Performs the operation on all resources. |

- **If you do not intend to reinstall Disaster Recovery framework software on the cluster in the future, deactivate and remove each protection group from the local cluster.**

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

  | | |
  |---|---|
  | `-e local` | Performs the operation only on the local cluster. |

3. **Stop the Disaster Recovery framework infrastructure on the local cluster.**

   ```
   # geoadm stop
   ```

   For more information about disabling the Disaster Recovery framework on a cluster, see "Disabling the Disaster Recovery Framework" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

4. **Remove the `ha-cluster-full` group package from each node in the local cluster.**

   You must remove the core Oracle Solaris Cluster group package before you can remove the Disaster Recovery framework software. However, this does not remove the installed Oracle Solaris Cluster software.

   ```
   # pkg uninstall ha-cluster-full
   ```

5. **Uninstall all Disaster Recovery framework software packages from each node in the local cluster.**

For a list of the Disaster Recovery framework 4.4 packages, see “How to Install Disaster Recovery Framework Software” on page 32.

```
# pkg uninstall ha-cluster/geo* ha-cluster/group-package/ha-cluster-geo*
```

6. **Verify that all Disaster Recovery framework packages are removed.**

```
# pkg info | grep geo
```

# Index

## E



## F



## G



## H



## I



## J



## L

## M



## N



## O



## P

## R



## S



## T



## U



## V

## Z