Oracle® Solaris Cluster Geographic Edition Remote Replication Guide for Oracle ZFS Storage Appliance

ORACLE®

Oracle Solaris Cluster Geographic Edition Remote Replication Guide for Oracle ZFS Storage Appliance

Part No: E39675

Access to Oracle Support

Contents

Using This Documentation

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

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

Product Documentation Library

Documentation and resources for this product and related products are available at `http://www.oracle.com/pls/topic/lookup?ctx=E56676-01`.

Feedback

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

1

Creating Oracle ZFS Storage Appliance Protection Groups

This chapter contains information about configuring data replication with the remote replication feature of Oracle ZFS Storage Appliance software.

Replication actions can be configured on the appliance software to send updates manually, on a schedule, or continuously. The Geographic Edition software supports the use of continuous mode for data replication for disaster-recovery environments.

During data replication, data from a primary appliance is copied to a backup or secondary appliance. The secondary site can be located at a geographically separated location from the primary site. This distance depends on the distance support that is available from your data replication product.

Before you can replicate data with the appliance software, you must be familiar with the Oracle ZFS Storage Appliance documentation and have the Oracle ZFS Storage Appliance product and the latest Oracle Solaris SRUs installed on your system. For information about configuring a Oracle ZFS Storage appliance, see the Oracle ZFS Storage Appliance product documentation.

This chapter provides the following information:

Replicating Data in an Oracle ZFS Storage Appliance Protection Group (Task Maps)

This section summarizes the tasks for configuring and administering Oracle ZFS Storage Appliance remote replication in a protection group.

TABLE 1   Configuration Tasks for Oracle ZFS Storage Appliance Remote Replication

Task	Description
Plan the Oracle ZFS Storage Appliance replication configuration.	See "Planning and Configuring Remote Replication With Oracle ZFS Storage Appliance Software" on page 13.
Configure remote replication.	See "Configuring Remote Replication With Oracle ZFS Storage Appliance Software" on page 18.
Create a protection group that is configured for Oracle ZFS Storage Appliance replication.	See "How to Create and Configure an Oracle ZFS Storage Appliance Protection Group" on page 24.
Add a remote replication component that is controlled by Oracle ZFS Storage Appliance software.	See "How to Add a Remote Replication Component to an Oracle ZFS Storage Appliance Protection Group" on page 28.
Add application resource groups to the protection group.	See "How to Add an Application Resource Group to an Oracle ZFS Storage Appliance Protection Group" on page 32.
Replicate the protection group configuration to a secondary cluster.	See "How to Replicate the Oracle ZFS Storage Appliance Protection Group Configuration to a Partner Cluster" on page 44.
Activate the protection group.	See "How to Activate a Protection Group" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
Verify the protection group configuration.	Perform a trial a switchover or takeover and test some simple failure scenarios before bringing your system online. See Chapter 3, "Migrating Services That Use Oracle ZFS Storage Appliance Remote Replication".
Check the runtime status of replication.	See "Checking the Runtime Status of Oracle ZFS Storage Appliance Remote Replication" on page 46.

TABLE 2   Administration Tasks for Oracle ZFS Storage Appliance Remote Replication

Task	Description
Detect failure.	See "Detecting Cluster Failure" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
Migrate services by using a switchover.	See "Migrating Replication Services by Switching Over Protection Groups" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
Migrate services by using a takeover.	See "Forcing a Takeover of a Protection Group" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Task	Description
Recover data after forcing a takeover.	See "Recovering Services to a Cluster on a System That Uses Oracle ZFS Storage Appliance Replication" on page 51.

Planning and Configuring Remote Replication With Oracle ZFS Storage Appliance Software

This section contains the following information:

- "Guidelines for Remote Replication With Oracle ZFS Storage Appliance Software" on page 13
- "Overview of the Oracle ZFS Storage Appliance Configuration File" on page 14
- "Geographic Edition Properties to Set for Oracle ZFS Storage Appliance Replication" on page 15
- "Remote Replication Layer Process for Validating the Application Resource Groups and Remote Replication Entities" on page 16

Guidelines for Remote Replication With Oracle ZFS Storage Appliance Software

Observe the following guidelines and restrictions when planning your Oracle ZFS Storage Appliance remote replication configuration:

- **Supported versions of Oracle ZFS Storage Appliance software** – See the Oracle Solaris Cluster 4 Compatibility Guide (`http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf`) for current information about the versions of Oracle ZFS Storage Appliance software that are supported with Geographic Edition software. This document requires My Oracle Support login.

- **Support for zone clusters as partner members** – Zone clusters using Oracle ZFS Storage Appliance replication are supported as members of a Geographic Edition partnership, by themselves or in combination with global clusters.

- **Restriction for synchronous replication** – Continuous replication is asynchronous. Oracle ZFS Storage appliances do not currently support synchronous replication, which does not consider data to be committed to stable storage until it is committed to stable storage on both the primary and secondary storage systems.

- **Quorum devices** – Do not configure a replicated volume as a quorum device. Locate any quorum devices on a shared, unreplicated volume or use a quorum server.

- **Project replication** – Only project level replication is supported.
- **Limit of one action per project** – Each project that is managed by Geographic Edition software can have only one action on the source with its paired package on the target. Multiple actions or packages are not supported for a project that is managed by Geographic Edition software.

For guidelines and requirements by Oracle ZFS Storage Appliance software, see the Oracle ZFS Storage online documentation at `https://`*appliance-hostname*`:215/wiki`, where *appliance-hostname* is the name of your storage appliance.

Overview of the Oracle ZFS Storage Appliance Configuration File

Oracle ZFS Storage Appliance remote replication with Geographic Edition is developed with the script-based plug-in module of Geographic Edition. Your appliance replication configuration must comply with all rules of the script-based plug-in. For each protection group, you must provide a script-based plug-in configuration file on each node. In addition, the Geographic Edition module for appliance replication includes its own configuration file, which is needed only at registration.

Creation of the appliance replication protection group for Geographic Edition is an automated process that takes the appliance configuration file as input and performs the necessary actions. The essential content of this file consists of the following key=value pairs:

PS

Name of the partnership

PG

Name of the protection group

REPCOMP

Name of the appliance project that is replicated from the primary site to the secondary site

REPRS

Name of the replication resource that monitors appliance project replication

REPRG

Name of the replication resource group to contain the replication resource

DESC

Description for the protection group

APPRG

> Application resource groups, one or more, which contain at least an `HAStoragePlus` or `ScalMountPoint` resource. A resource group can belong to only one protection group.

CONFIGFILE

> Configuration file for the script-based plug-in evaluation rules

LOCAL_CONNECT_STRING

> Source appliance connection string, in the form *user*@*hostname* at the local site

REMOTE_CONNECT_STRING

> Target appliance connection string, in the form *user*@*hostname* at the remote site

CLUSTER_DGS

> Oracle Solaris Cluster device groups, separated by commas

For more information, see Chapter 13, "Script-Based Plug-Ins" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Geographic Edition Properties to Set for Oracle ZFS Storage Appliance Replication

This section describes the properties that can be modified for Oracle ZFS Storage Appliance remote replication.

The following table lists the script-based plug-in properties.

Property Type	Properties
Script-based plug-in data replication component properties	■ `local_service_password` ■ `remote_service_password` See "Replicated Component Properties - Overview" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
Script-based plug-in protection group properties	■ `configuration_file` See "configuration_file Property" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

The following table lists the general properties.

Property Type	Properties
General protection group properties	`RoleChange_ActionCmd` `Timeout` See "Protection Group Properties" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Remote Replication Layer Process for Validating the Application Resource Groups and Remote Replication Entities

During protection group validation, the Oracle ZFS Storage Appliance remote replication layer validates the application resource groups and the replication entities by verifying that an application resource group in the protection group has its `Auto_start_on_new_cluster` property set to `False`.

Setting the `Auto_start_on_new_cluster` property to `False` prevents the Oracle Solaris Cluster resource group manager from automatically starting the application resource groups. In this case, the startup of resource groups is reserved for the Geographic Edition software.

The appliance `geocontrol` module supplies a script that is used by the script-based plug-in module. The script entry points require the same set of arguments. These arguments are validated for semantics and completeness. The following validation checks are performed:

- Are all of the mandatory arguments defined?
- Is the appliance monitoring resource defined?
- Are the hostnames of the local and remote appliances specified?
- Are the login credentials provided to execute `oscgeo7kcli` commands?
- Is the remote replication component name provided?

When the validation is complete, the Geographic Edition software creates and brings online the replication resource group and its resources, if they don't already exist. If a resource group or resource of the same name already exists, the Geographic Edition software might modify its properties. The software cannot create a new resource group or a resource of the same name if one already exists.

Creating and Validating an Oracle ZFS Storage Appliance Protection Group

This section contains the following topics:

- "Strategies for Creating Oracle ZFS Storage Appliance Protection Groups" on page 17
- "Configuring Remote Replication With Oracle ZFS Storage Appliance Software" on page 18
- "How to Create and Configure an Oracle ZFS Storage Appliance Protection Group" on page 24
- "Debugging an Oracle ZFS Storage Appliance Protection Group" on page 27

Note - You can create protection groups that are not configured to use remote replication. To create a protection group that does not use a replication subsystem, omit the -d *data-replication-type* option when you use the `geopg` command. The `geoadm status` command shows a state for these protection groups of `Degraded`.

For more information, see "Creating a Protection Group That Does Not Require Data Replication" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

Strategies for Creating Oracle ZFS Storage Appliance Protection Groups

The following task maps describe the steps to perform:

TABLE 3   Task Map: Creating a Protection Group

Task	Description
1. Create a role and user for each storage appliance. Create projects and enable replication. Configure remote replication for both partner clusters.	See "Configuring Remote Replication With Oracle ZFS Storage Appliance Software" on page 18.
2. Create the protection group from a cluster node.	See "How to Create and Configure an Oracle ZFS Storage Appliance Protection Group" on page 24.
3. Add the remote replication component to the protection group.	See "How to Add a Remote Replication Component to an Oracle ZFS Storage Appliance Protection Group" on page 28.
4. Start the protection group locally.	See "How to Activate a Protection Group" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Task	Description
5. Add the application resource group to the protection group.	See "How to Add an Application Resource Group to an Oracle ZFS Storage Appliance Protection Group" on page 32.
6. From the secondary cluster, retrieve the protection group configuration.	See "How to Replicate the Oracle ZFS Storage Appliance Protection Group Configuration to a Partner Cluster" on page 44.
7. From the secondary cluster, activate the protection group locally.	See "How to Activate a Protection Group" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Configuring Remote Replication With Oracle ZFS Storage Appliance Software

This section describes the steps you must perform before you can configure Oracle ZFS Storage Appliance remote replication with Geographic Edition software. The following procedures are in this section:

- "How to Create a Role and Associated User for the Primary and Secondary Appliances" on page 18
- "How to Create a Project and Enable Replication for the Project" on page 19
- "How to Configure Oracle Solaris Cluster Resources on the Primary Cluster" on page 20
- "How to Configure Oracle Solaris Cluster Resources on the Secondary Cluster" on page 20

▼ How to Create a Role and Associated User for the Primary and Secondary Appliances

If a role and associated user do not yet exist on the source and target appliances, perform this procedure to create them.

1. **Log in to the Oracle ZFS Storage appliance.**

2. **Create a role for remote replication.**

   Configure the role with the following permissions:

   - Object `nas.*.*.*` with permissions `clone`, `destroy`, `rrsource`, `rrtarget`, `createShare`, and `createProject`.

■ Object `workflow.*.*` with permission `read`.

3. **Create a user for replication that is associated with the role you created in Step 2.**

▼ How to Create a Project and Enable Replication for the Project

1. **Log in to the Oracle ZFS Storage appliance on the primary `cluster-paris` site.**

2. **Navigate to Shares > Projects and create the projects that you need for your application.**

3. **In each project, create the file systems and LUNs that you need for your application.**

   Ensure that NFS exceptions and LUN settings are identical on the primary and secondary storage appliances. For more information, see "Copying and Editing Actions" in *Oracle ZFS Storage 7000 System Administration Guide* (`http://docs.oracle.com/cd/E26765_01/html/E26397/`).

4. **For iSCSI LUNs, if you use nondefault targets and target groups, ensure that target groups and initiator groups used by LUNs within the project also exist on the replication target.**

   These groups must use the same name in the replication target as in the source appliance.

   For more information on replicating projects with iSCSI LUNs, see "iSCSI Configurations and Replication" in Oracle ZFS Storage Appliance Administration Guide (`http://docs.oracle.com/cd/E71909_01/html/E71919/goksk.html`)

5. **For each project, navigate to Replication, create an action, and enable the action with continuous mode.**

Troubleshooting If you need to stop Oracle ZFS Storage Appliance replication directly from the Oracle ZFS Storage appliance, you must perform the following tasks in the order shown:

■ Set `continuous=false`.

■ Wait for the update to complete.

■ Set `enabled=false` to stop replication.

Geographic Edition requires that `last_result` of replication be a success. Otherwise, adding a project to a Geographic Edition protection group and protection group validation will fail.

▼ How to Configure Oracle Solaris Cluster Resources on the Primary Cluster

This procedure creates Oracle Solaris Cluster resources on the primary cluster for the application to be protected.

Before You Begin   Ensure that the following tasks are completed on the storage appliance:

- Replication peers are configured by the storage administrator.
- Projects are configured by the storage administrator.
- Replication is enabled for the project.
- For iSCSI LUNs, if you use nondefault target groups, the target groups and initiator groups used by LUNs within the project also exist on the replication target. In addition, these groups use the same names in the replication target as in the source appliance.
- If you use file systems, NFS Exceptions exist for all nodes of both clusters. This ensures that either cluster can access the file systems when that cluster has the primary role.

1. **Create the Oracle Solaris Cluster device groups, file systems, or ZFS storage pools you want to use.**

   Specify the LUNs or file systems in the Oracle ZFS Storage appliance to be replicated.

   For information about creating device groups, file systems, and ZFS storage pools in a cluster configuration, see *Oracle Solaris Cluster System Administration Guide*.

2. **Create an HAStoragePlus resource or a scalable mount-point resource for the device group, file system, or ZFS storage pool you use.**

   This resource manages bringing online the Oracle ZFS Storage Appliance storage on both the primary and secondary clusters.

   For information about creating an HAStoragePlus or scalable mount-point resource, see *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

▼ How to Configure Oracle Solaris Cluster Resources on the Secondary Cluster

This procedure creates Oracle Solaris Cluster resources on the secondary cluster for the application to be protected.

Before You Begin   Ensure that the following tasks are completed on the storage appliance:

- Replication peers are configured by the storage administrator.
- Projects are configured by the storage administrator.

- Replication is enabled for the project.
- For iSCSI LUNs, if you use nondefault target groups, the target groups and initiator groups used by LUNs within the project also exist on the replication target. In addition, these groups must use the same names in the replication target as in the source appliance.
- If you use file systems, NFS Exceptions exist for all nodes of both clusters. This ensures that either cluster can access the file systems when that cluster has the primary role.

1. **On one node of the `cluster-newyork` (partner) site, create the application group.**

   The `Auto_start_on_new_cluster` property must be set to `False`.

   ```
   phys-newyork-1# clresourcegroup create -p Auto_start_on_new_cluster=False \
   application-resource-group
   ```

2. **Determine whether the replicated project contains any LUNs.**

   a. **On the `cluster-paris` (primary) site, access the Oracle ZFS Storage Appliance browser user interface (BUI).**

   b. **Navigate to Shares > Projects and select the project being replicated.**

3. **If the project contains only file systems, perform the following tasks.**

   If the project contains any LUNs, skip to Step 4.

   a. **If replication is not in continuous mode, select Replication for the project and click Update Now or Sync Now.**

   This executes a manual replication to synchronize the two sites.

   b. **On the `cluster-newyork` (partner) site, access the appliance BUI.**

   c. **Navigate to In Projects > Replica and select the project being replicated.**

   d. **Select Replication for the project and click Clone Most Recently Received Project Snapshot.**

   Enter the same project name as on the primary appliance.

4. **If the replicated project contains LUNS, perform the following tasks.**

   a. **Create protection group and add replicated project and resource groups to it.**

   See "How to Create and Configure an Oracle ZFS Storage Appliance Protection Group" on page 24.

> **Note -** Resource groups added to the protection group can be empty on the secondary cluster. The storage and application resources will be created on the secondary cluster in subsequent steps.

b. **From one node of either cluster, start the protection group globally.**

   `# geopg start -e global` *protection-group*

c. **From one node of either cluster, switch over the protection group to the secondary cluster.**

   `# geopg switchover -f -m cluster-newyork` *protection-group*

   The project is made local on the secondary storage.

d. **On secondary cluster, map iSCSI devices from the project on the secondary storage.**

   i. **Map the iSCSI devices to the corresponding DID numbers.**

   ii. **Use the `cldevice list` command to find devices corresponding to the devices being exported from the appliance.**

e. **Create the Oracle Solaris Cluster device groups or import the ZFS storage pools on the LUNs in the project.**

   Specify the LUNs or file systems in the project that is now local on the secondary appliance.

   For information about creating device groups and file systems and adding ZFS storage pools in a cluster configuration, see *Oracle Solaris Cluster System Administration Guide*.

5. **On cluster-newyork, create an HAStoragePlus resource or a scalable mount-point resource for the device group, file system, or ZFS storage pool you use.**

   This resource manages bringing online the Oracle ZFS Storage Appliance storage on the secondary cluster.

   For information about creating an HAStoragePlus or scalable mount-point resource, see *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

6. **Bring up the application on `cluster-newyork` that uses the replicated storage and create corresponding cluster resources.**

7. On cluster-newyork, confirm that the application resource group is correctly configured by bringing it online.

phys-newyork-1# **clresourcegroup online -emM** *application-resource-group*

8. If the replication project contains only file systems, perform the following tasks.

If the project contains any LUNs, skip to Step 9.

a. On a node of the secondary cluster, put the application resource group in the unmanaged state on secondary cluster.

phys-newyork-1# **clresource disable -g** *application-resource-group* **+**
phys-newyork-1# **clresourcegroup offline** *application-resource-group*
phys-newyork-1# **clresourcegroup unmanage** *application-resource-group*

b. If you created a file system and it is mounted, unmount the file system.

phys-newyork-1# **umount /mounts/***file-system*

c. If the Oracle Solaris Cluster device group is online, take it offline.

phys-newyork-1# **cldevicegroup offline** *raw-disk-group*

d. Destroy the clone on the Oracle ZFS Storage appliance.

i. Access the appliance BUI on the `cluster-newyork` site.

ii. Navigate to Shares > Projects and select the project that is cloned.

iii. Select Remove or Destroy entry for the cloned project.

9. If the replicated project contains any LUNs, from one node of either cluster, switch over the protection group to the primary cluster.

This step takes offline the configuration on the secondary cluster and brings it online on the primary cluster.

geopg switchover -f -m cluster-paris *protection-group*

Next Steps ■ If the replicated project contains any LUNs, initial configuration on the primary and secondary clusters is now complete.

■ If the replicated project contains only file systems, go to "How to Create and Configure an Oracle ZFS Storage Appliance Protection Group" on page 24.

▼ How to Create and Configure an Oracle ZFS Storage Appliance Protection Group

Perform this procedure from a node of the primary cluster.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, and in the Protection Groups section click `Create`. Use data replication type `Sbp`. For more information about Oracle Solaris Cluster Manager, see Chapter 13, "Using the Oracle Solaris Cluster GUI" in *Oracle Solaris Cluster System Administration Guide*.

Before You Begin Ensure that the following conditions are met:

- The Geographic Edition software is installed on the primary and secondary storage appliances.
- You have reviewed the information in "Planning and Configuring Remote Replication With Oracle ZFS Storage Appliance Software" on page 13.
- You have created a remote replication role and user on each appliance. See "How to Create a Role and Associated User for the Primary and Secondary Appliances" on page 18.
- You have created the projects you need. See "How to Create a Project and Enable Replication for the Project" on page 19.
- The local cluster is a member of a partnership.
- The protection group you are creating does not already exist on either partner cluster.

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

   For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

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

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

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

2. **Copy the default `zfssa_geo_config` file to another location.**

The `/var/tmp/` directory is used as an example location in this step and the next step.

```
# cp /opt/ORCLscgrepzfssa/etc/zfssa_geo_config /var/tmp/
```

3. **On all nodes of both clusters, create or update an `/etc/opt/SUNWscgrepsbp/configuration` file to contain the script-based plug-in evaluation rules.**
Update the file so that it contains one line that contains the rule information for the remote replication component.

project-name|any|*nodelist*

project-name

> Name of the project.

nodelist

> The name of one or more cluster nodes where the plug-in is to validate the configuration.

For example, assuming that the nodes of cluster `cluster-newyork` are `phys-newyork-1` and `phys-newyork-2`, on each node of cluster `cluster-newyork`, you would issue the following commands:

```
phys-newyork-N# mkdir /etc/opt/SUNWscgrepsbp
phys-newyork-N# echo "trancos|any|phys-newyork-1,phys-newyork-2" > /etc/opt/
SUNWscgrepsbp/configuration
```

Assuming that the nodes of cluster `paris` are `phys-paris-3` and `phys-paris-4`, on each node of cluster `paris`, you would issue the following commands:

```
phys-paris-N# mkdir /etc/opt/SUNWscgrepsbp
phys-paris-N# echo "trancos|any|phys-paris-3,phys-paris-4" > /etc/opt/SUNWscgrepsbp/
configuration
```

For more information about configuration files, see "configuration_file Property" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

4. **Specify the configuration values in the temporary `/var/tmp/zfssa_geo_config` file.**
The following list uses sample values:

```
PS=zfssa-partnership
PG=zfssa-pg
REPCOMP=trancos
REPRS=zfssa-replication-resource
REPRG=zfssa-replication-resource-group
DESC="ZFS Storage Appliance replication protection group"
APPRG=usa-rg
CONFIGFILE=/etc/opt/SUNWscgrepsbp/configuration
```

```
LOCAL_CONNECT_STRING=user@local-appliance.example.com
REMOTE_CONNECT_STRING=user@remote-appliance.example.com
CLUSTER_DGS=
```

> **Note -** For the `LOCAL_CONNECT_STRING` and `REMOTE_CONNECT_STRING` variables, provide the user that you created in Step 3 of "How to Create a Role and Associated User for the Primary and Secondary Appliances" on page 18.

For more information about the `zfssa_geo_config` file, see "Overview of the Oracle ZFS Storage Appliance Configuration File" on page 14.

5. **Execute the `zfssa_geo_register` script on the primary cluster.**

   For example:

   ```
   phys-newyork-1# /opt/ORCLscgrepzfssa/util/zfssa_geo_register -f /var/tmp/
   zfssa_geo_config
   ```

6. **Replicate the protection group to the partner cluster.**

   The final messages of the registration script outline the required `geopg get` command. You must log in to one node of the partner cluster and execute that exact command.

   For example, where *zfssa-partnership* is the partnership name and *zfssa-protection-group* is the protection group name:

   ```
   phys-paris-1# geopg get --partnership zfssa-partnership zfssa-protection-group
   ```

7. **Verify the protection group configuration.**

   ```
   phys-newyork-1# geoadm status
   phys-newyork-1# clresource status zfssa-replication-resource
   ```

   zfssa-replication-resource

   Specifies the name of the replication resource.

8. **Verify that you can switch over from one cluster to the other.**

   See "How to Switch Over Replication From the Primary Cluster to the Secondary Cluster" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Troubleshooting   If you experience failures while performing this procedure, enable debugging. See "Debugging an Oracle ZFS Storage Appliance Protection Group" on page 27.

Debugging an Oracle ZFS Storage Appliance Protection Group

If you encounter problems when creating a protection group or replicating a protection group with the `geopg get` command, you can set the `DEBUG` property of the `/opt/ORCLscgrepzfssa/etc/config` file to run trace logs. These logs will display on the terminal.

After an Oracle ZFS Storage Appliance remote replication component is added to the protection group, you instead enable debugging by directly setting the `Debug_level` property of the Oracle ZFS Storage Appliance resource with the `clresource set` command. Debug messages will display on the terminal.

 # clresource set -p Debug_level=_N_ _zfssa-replication-resource_

The following values are valid for the `DEBUG` and `Debug_level` properties:

`0`   No trace. This is the default.

`1`   Function trace.

`2`   Trace everything.

Additionally, logs of `oscgeo7kcli` calls and their results are recorded in `/var/cluster/geo/zfssa/`_replication-component_`_logfile` files on each cluster node.

Adding a Remote Replication Component to an Oracle ZFS Storage Appliance Protection Group

A protection group is the container for the application resource groups and remote replication components, which contain data for services that are protected from disaster. The Geographic Edition software protects the data by replicating it from the primary cluster to the secondary cluster. By adding a remote replication component to a protection group, the software monitors the replication status of an appliance project. The software also controls the role and state of the project during protection group operations such as start, stop, switchover, and takeover.

This section provides the following information for administering remote replication components in an Oracle ZFS Storage Appliance protection group:

- "How to Add a Remote Replication Component to an Oracle ZFS Storage Appliance Protection Group" on page 28

- "Remote Replication Subsystem Process for Verifying the Data Replication Component" on page 31

▼ How to Add a Remote Replication Component to an Oracle ZFS Storage Appliance Protection Group

Perform this procedure to add a remote replication component to an existing Oracle ZFS Storage Appliance protection group.

Note - When the protection group is initially created, any remote replication components that are specified in the `zfssa_geo_config` configuration file are added to the protection group. Thus, you only need to run this procedure to add more remote replication components to existing Oracle ZFS Storage Appliance protection groups.

Before You Begin Before you add a remote replication component to a protection group, ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The protection group is offline on the local cluster and the partner cluster, if the partner cluster can be reached.
- The underlying project exists on the appliance that is connected to the local cluster.
- The replication action must exist.

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

   For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

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

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

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

2. On all nodes of both clusters, create or update a `/var/tmp/zfssa_geo_eval_rules` configuration file to contain the script-based plug-in evaluation rules.

Update the file so that it contains one line that contains the rule information for the remote replication component.

project-name|any|*nodelist*

project-name

Name of the project.

nodelist

The name of one or more cluster nodes where the plug-in is to validate the configuration.

For more information about configuration files, see "configuration_file Property" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

3. Update the appliance configuration files with information for the new remote replication component.

a. Make a copy of the default `/opt/ORCLscgrepzfssa/etc/zfssa_geo_config` file to a different location, such as `/var/tmp/`.

```
# cp /opt/ORCLscgrepzfssa/etc/zfssa_geo_config /var/tmp/zfssa_geo_config
```

b. Edit the copy of the `zfssa_geo_config` file with updates for the new remote replication component.

Define the following key=value pairs as shown:

```
PS=partnership
PG=protection-group
REPCOMP=project
REPRS=resource
REPRG=resource-group
CONFIGFILE=eval-rules-configuration-file
LOCAL_CONNECT_STRING=user@source-appliance
REMOTE_CONNECT_STRING=user@target-appliance
```

PS=*partnership*

Specifies the name of the existing partnership.

PG=*protection-group*

Specifies the name of the protection group that you are adding the remote replication component to.

REPRS=*resource*

> Specifies the name of a resource *other than* the existing resource.

REPRG=*resource-group*

> Specified an existing resource group that contains the remote replication resources for this protection group.

CONFIGFILE=*eval-rules-configuration-file*

> Specifies the edited copy of the `zfssa_geo_eval_rules` file that you created in Step 2.

LOCAL_CONNECT_STRING=*user@local-appliance*

> Specifies the source user and hostname.

REMOTE_CONNECT_STRING=*user@remote-appliance*

> Specifies the target user and hostname.

Note - For the `LOCAL_CONNECT_STRING` and `REMOTE_CONNECT_STRING` variables, provide the user that you created in Step 3 of "How to Create a Role and Associated User for the Primary and Secondary Appliances" on page 18.

---

4. **Add a remote replication component to the protection group.**

    Use the `zfssa_geo_register` script with a new configuration file.

    ```
    phys-newyork-1# /opt/ORCLscgrepzfssa/util/zfssa_geo_register -f /var/tmp/
    zfssa_geo_config
    ```

    The command adds a remote replication component to a protection group on the local cluster. If the partner cluster contains a protection group with the same name, the command also propagates the new configuration to the partner cluster.

    Note - Because the add operation for the remote replication component is performed during the scripted registration, an example is not provided here.

Troubleshooting If you have difficulties adding the remote replication component to the protection group, see "Debugging an Oracle ZFS Storage Appliance Protection Group" on page 27.

Remote Replication Subsystem Process for Verifying the Data Replication Component

During protection group validation, the Oracle ZFS Storage Appliance remote replication layer validates the application resource groups and the replication entities by verifying that an application resource group in the protection group has its `Auto_start_on_new_cluster` property set to `False`.

When you bring a protection group online on the primary cluster, bring the application resources groups participating in that protection group online only on the same primary cluster. Setting the `Auto_start_on_new_cluster` property to `False` prevents the Oracle Solaris Cluster Resource Group Manager (RGM) from automatically starting the application resource groups. In this case, the startup of resource groups is reserved for the Geographic Edition software.

Application resource groups must be online only on the primary cluster when the protection group is activated.

The appliance `geocontrol` module supplies a script that is used by the script-based plug-in module. The script entry points require the same set of arguments. These arguments are validated for semantics and completeness. The following validation checks are performed:

- Are all of the mandatory arguments defined?
- Is the remote replication component configuration file for the script-based plug-in evaluation rules defined?
- Is the specified replication resource configured with a correct start command, if the resource exists already?

When the Oracle ZFS Storage Appliance remote replication component is added to a protection group, the data replication layer makes the following validations:

- Only one replication action is defined at the primary appliance, if the protection group role is `PRIMARY`.
- Only one replication package exists on the secondary appliance, if the protection group role is `SECONDARY`.
- The replication mode is set to `continuous`.
- The ZFS storage appliance can be reached. When an Oracle ZFS Storage Appliance remote replication component is added to a protection group, an Oracle Solaris Cluster data replication resource as defined by the property `REPRS` in the configuration file is created by this command. This resource monitors the data replication state.

> ⚠️ **Caution -** Do not change, remove, or take offline these resources or resource groups. Use only Geographic Edition commands to administer replication resource groups and resources that are internal entities managed by Geographic Edition software. Altering the configuration or state of these entities directly with Oracle Solaris Cluster commands might result in unrecoverable failure.

When the validation is complete, the Geographic Edition software adds the application resource group to the protection group.

> **Note -** Every entry point of the underlying script-based plug-in has a validation method. In the case of Oracle ZFS Storage Appliance remote replication, all the validation methods are the same.

Adding an Application Resource Group to an Oracle ZFS Storage Appliance Protection Group

To make an application highly available, the application must be managed as a resource in an application resource group.

The initial registration of the protection group is performed with the `zfssa_geo_register` script. This section explains how to manage the application resource groups on their own.

All the entities you configure for the application resource group on the primary cluster, such as application data resources, application configuration files, and the resource groups, must be replicated manually on the secondary cluster. The resource group names must be identical on both clusters. Also, the data that the application resource uses must be replicated on the secondary cluster.

▼ How to Add an Application Resource Group to an Oracle ZFS Storage Appliance Protection Group

Perform this procedure to add an existing resource group to the list of application resource groups for a protection group.

> **Note -** You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, click the protection group name, and in the Resource Groups section click `Add Resource Groups`. For more information about Oracle Solaris Cluster Manager, see Chapter 13, "Using the Oracle Solaris Cluster GUI" in *Oracle Solaris Cluster System Administration Guide*.

> **Note -** When the protection is initially created, any resource groups that are specified in the `zfssa_geo_config` configuration file are automatically created as well. Thus, you do not need to perform this procedure to add the resource groups specified in the `zfssa_geo_config` file at the time the protection group was created.

The protection group can be activated or deactivated and the resource group can be either online or unmanaged.

- If the resource group is unmanaged and the protection group is activated after the configuration of the protection group has changed, the local state of the protection group becomes `Error`.
- If the resource group to add is online and the protection group is deactivated, the request is rejected. You must activate the protection group before adding an online resource group.

Before You Begin Ensure that the following conditions are met:

- The protection group is defined.
- The resource group to be added already exists on both clusters and is in an appropriate state.

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

   For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

   > **Note -** If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.
   >
   > `# chmod A+user:`*username*`:rwx:allow /var/cluster/geo`
   >
   > The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and Oracle ZFS Storage Appliance software.

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.

Geographic edition requires that the application resource group be in unmanaged state on the secondary cluster. If the protection group is stopped on the primary cluster, the application resource group must be unmanaged on the primary cluster also. If the protection group is active on the primary cluster, the application resource group must be in unmanaged or online state on the primary cluster.

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 Geographic Edition software, see Appendix B, "Legal Names and Values of Geographic Edition Entities," in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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.

6. **If necessary, bring online the application resource group.**

   `# clresourcegroup online` *application-resource-group*

Example 1 Adding an Application Resource Group to an Oracle ZFS Storage Appliance Protection Group

The following example adds two application resource groups, `apprg1` and `apprg2`, to the `zfssa-pg` protection group.

`# geopg add-resource-group apprg1,apprg2 zfssa-pg`

Troubleshooting If the add operation is unsuccessful on the local cluster, the configuration of the protection group is not modified. Otherwise, the configuration status is set to `OK` on the local cluster.

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

♦♦♦ **C H A P T E R   2**

2

Administering Oracle ZFS Storage Appliance Protection Groups

This chapter contains information about administering data replication with the remote replication feature of Oracle ZFS Storage Appliance software.

The chapter contains the following sections:

- "Administering Oracle ZFS Storage Appliance Remote Replication Components" on page 37
- "Administering Oracle ZFS Storage Appliance Application Resource Groups" on page 41
- "Replicating an Oracle ZFS Storage Appliance Protection Group Configuration to a Partner Cluster" on page 44
- "Checking the Runtime Status of Oracle ZFS Storage Appliance Remote Replication" on page 46

Administering Oracle ZFS Storage Appliance Remote Replication Components

This section provides the following information for administering remote replication components in an Oracle ZFS Storage Appliance protection group:

- "How to Modify an Oracle ZFS Storage Appliance Remote Replication Component" on page 38
- "How to Remove a Remote Replication Component From an Oracle ZFS Storage Appliance Protection Group" on page 40

▼ How to Modify an Oracle ZFS Storage Appliance Remote Replication Component

> **Note -** You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, click the protection group name, in the Data Replication Components section click the remote replication component name, and click `Edit`. For more information about Oracle Solaris Cluster Manager, see Chapter 13, "Using the Oracle Solaris Cluster GUI" in *Oracle Solaris Cluster System Administration Guide*.

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

   For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

   > **Note -** If you use a role with Geo Management RBAC rights, ensure that the `/var/cluster/geo` ACLs are correct on each node of both partner clusters. If necessary, assume the `root` role on the cluster node and set the correct ACLs.
   >
   > **# chmod A+user:***username***:rwx:allow /var/cluster/geo**
   >
   > The `/var/cluster/geo` directory must have the correct access control lists (ACL) applied for compatibility between the Geo Management RBAC rights profile and Oracle ZFS Storage Appliance software.

2. **The script-based plug-in evaluation rules for the replication component is correct in the file that is already associated with the protection group.**

   The protection group configuration output of `geopg show` will display which file is associated with the protection group.

   For more information about configuration files, see "configuration_file Property" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

3. **Update the ZFS appliance configuration file with the changed information for an existing replication component.**

   For the parameters that you do not want to modify, specify the same values that exist in the replication component.

> **Note -** You must specify existing values for unchanged parameters and new values for the parameters to be modified. Hence, it is recommended to save the configuration file after you create the replication component, so that you can edit required values when you want to modify the replication component.

a. **If you have a saved copy of the configuration file used to create the replication component, use that file.**

    Otherwise, make a copy of the `/opt/ORCLscgrepzfssa/etc/zfssa_geo_config` replication configuration file to a different location, such as `/var/tmp/geo/zfssa/zfssa_geo_config`.

    ```
    # cp /opt/ORCLscgrepzfssa/etc/zfssa_geo_config /var/tmp/geo/zfssa/
    ```

b. **Edit this replication configuration file with updates for the existing replication component.**

    Fill in the existing values for unchanged parameters and new values for the parameters that require to be modified.

    For example, if you want to modify the user for connecting to the local and remote Oracle ZFS Storage Appliances for an existing replication component `trancos` in a protection group `zfssa-pg` configured in partnership `zfssa-partnership`. If the replication configuration file to use for `trancos` is `/var/tmp/geo/zfssa/zfssa_geo_config`, then set these new values in that file, keeping all other parameters unchanged.

    ```
    PS=zfssa-partnership
    PG=zfssa-pg
    REPCOMP=trancos
    REPRS=zfssa-replication-resource
    REPRG=zfssa-replication-resource-group
    DESC="ZFS Storage Appliance replication protection group"
    APPRG=usa-rg
    CONFIGFILE=/etc/opt/SUNWscgrepsbp/configuration
    LOCAL_CONNECT_STRING=newuser@local-appliance.example.com
    REMOTE_CONNECT_STRING=newuser@remote-appliance.example.com
    CLUSTER_DGS=
    MODIFY_PASSPHRASE=
    ```

4. **On the node of the local partner having the changed replication configuration file, run the `zfssa_geo_register` script to modify the replication component.**

    Use the `zfs_snap_geo_register` script with the modified configuration file.

    ```
    # /opt/ORCLscgrepzfssa/util/zfssa_geo_register -f /var/tmp/geo/zfssa/zfssa_geo_config
    ```

The command modifies the replication component in a protection group on the local cluster. If the partner cluster contains a protection group with the same name, the command also propagates the new configuration to the partner cluster.

You can perform this procedure from either cluster. However, if you are using a saved configuration file that was used to create replication component from the primary cluster, perform this procedure on the primary cluster since the values in the file will be appropriate from the primary cluster.

5. **Check that the new parameters reflect in the configuration of the replication component on each partner cluster.**

   For example, you can type the following command on one node of `paris` cluster and one node of `newyork` cluster:

   ```
   # geopg show zfssa-pg
   .....
   SBP Script-based plug-in configurations:
   trancos
   .....
   repcom1_create_config_args :
   ......paris:newuser@local-appliance.example.com,newyork:newuser@remote-appliance.
   example.com......
   ```

▼ How to Remove a Remote Replication Component From an Oracle ZFS Storage Appliance Protection Group

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, click the protection group name, in the Data Replication Components section highlight the remote replication component name, and click Remove. For more information about Oracle Solaris Cluster Manager, see Chapter 13, "Using the Oracle Solaris Cluster GUI" in *Oracle Solaris Cluster System Administration Guide*.

Before You Begin Ensure that the following conditions are met:

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

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

    For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

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

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

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

    ---

2. **Remove the remote replication component.**

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

    zfssa-replication-component

    > Specifies the name of the remote replication component.

    zfssa-protection-group

    > Specifies the name of the protection group.

Example 2   Removing a Remote Replication Component From an Oracle ZFS Storage Appliance Protection Group

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

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

Administering Oracle ZFS Storage Appliance Application Resource Groups

To make an application highly available, the application must be managed as a resource in an application resource group.

The initial registration of the protection group is performed with the `zfssa_geo_register` script. This section explains how to manage the application resource groups on their own.

All the entities you configure for the application resource group on the primary cluster, such as application data resources, application configuration files, and the resource groups, must be replicated manually on the secondary cluster. The resource group names must be identical on both clusters. Also, the data that the application resource uses must be replicated on the secondary cluster.

This section contains information about deleting an application resource group from an Oracle ZFS Storage Appliance protection group.

▼ How to Delete an Application Resource Group From an Oracle ZFS Storage Appliance Protection Group

You can remove an application resource group from a protection group without altering the state or contents of an application resource group.

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, click the partnership name, click the protection group name, highlight the resource group name, and click `Remove`. For more information about Oracle Solaris Cluster Manager, see Chapter 13, "Using the Oracle Solaris Cluster GUI" in *Oracle Solaris Cluster System Administration Guide*.

Before You Begin   Ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The resource group to be removed is part of the application resource groups of the protection group.

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

For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

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

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

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

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

```
# geopg remove-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 removes an application resource group from a protection group on the local cluster. If the partner cluster contains a protection group of the same name, the command also removes the application resource group from the protection group on the partner cluster.

Example 3 Deleting an Oracle ZFS Storage Appliance Application Resource Group From a Protection Group

The following example removes two application resource groups, apprg1 and apprg2, from the zfssa-pg protection group.

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

Troubleshooting If the remove operation is unsuccessful on the local cluster, the configuration of the protection group is not modified. Otherwise, the configuration status is set to OK on the local cluster.

If the configuration status is OK on the local cluster but the remove operation is unsuccessful on the partner cluster, the configuration status is set to Error on the partner cluster.

Replicating an Oracle ZFS Storage Appliance Protection Group Configuration to a Partner Cluster

After you have configured remote replication, resource groups, and resources on your primary and secondary clusters and you have created a protection group for those entities on the primary cluster, you can replicate the configuration of the protection group to the secondary cluster.

▼ How to Replicate the Oracle ZFS Storage Appliance Protection Group Configuration to a Partner Cluster

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

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

- The protection group is defined on the remote cluster, not on the local cluster.
- The application resource groups in the protection group on the remote cluster exist on the local cluster.
- Application resource groups are online only on the primary cluster and in the unmanaged state on the secondary cluster.

Perform this procedure from `phys-newyork-1`, which is a node on the secondary cluster. For a reminder of which node is `phys-newyork-1`, see "Example Geographic Edition Cluster Configuration" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

   For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

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

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

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

---

2. **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 `True`.

 # clresourcegroup set -p Auto_start_on_new_cluster=True *resource-group*

3. **Replicate the protection group configuration to the partner cluster.**

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

-s *partnership*

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

protection-group

> Specifies the name of the protection group.
>
> If no protection group is specified, then all protection groups that exist in the specified partnership on the remote partner are created on the local cluster.

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

Note - The `geopg get` command replicates Geographic Edition related entities. To replicate Oracle Solaris Cluster resource groups, resource types, and resources, use the `cluster export -t rg,rt,rs` command to generate an XML cluster configuration file, modify the XML file for the expected configuration on the secondary cluster. Then run the `clresource create` command with the -a option to apply the configuration updates. For more information, see "How to Configure Oracle Solaris Cluster Software on All Nodes (XML)" in *Oracle Solaris Cluster Software Installation Guide* and the `cluster(1CL)`and `clresource(1CL)` man pages.

Example 4   Replicating an Oracle ZFS Storage Appliance Protection Group to a Partner Cluster

The following example replicates the configuration of `zfssa-pg` from `cluster-paris` to `cluster-newyork`.

```
# rlogin phys-newyork-1 -l root
phys-newyork-1# geopg get -s paris-newyork-ps zfssa-pg
```

Troubleshooting   If the validation is successful, the configuration status is set to OK, and the protection group is created on the local cluster. This protection group contains a remote replication component and application group that are configured almost identically to the remote replication component and application group on the remote cluster.

If the validation fails, the protection group is not created on the local cluster. Fix the cause of the error, and replicate it again.

If you have difficulties adding the component to the protection group, see "Debugging an Oracle ZFS Storage Appliance Protection Group" on page 27.

Checking the Runtime Status of Oracle ZFS Storage Appliance Remote Replication

The Geographic Edition software internally creates and maintains one replication resource group for each protection group. The name of the replication resource group is specified by the user in the configuration as described in "How to Create and Configure an Oracle ZFS Storage Appliance Protection Group" on page 24.

You can obtain an overall view of the status of replication as well as a more detailed runtime status of the appliance replication resource groups. The following sections describe the procedures for checking each status:

- "Overview of Displaying an Oracle ZFS Storage Appliance Runtime Status" on page 47
- "How to Check the Runtime Status of Oracle ZFS Storage Appliance Replication" on page 47
- "Oracle ZFS Storage Appliance Replication Resource Group Runtime Status and Status Messages" on page 48

Overview of Displaying an Oracle ZFS Storage Appliance Runtime Status

The status of each Oracle ZFS Storage Appliance remote replication resource indicates the status of replication on a particular remote replication component. The status of all the resources under a protection group are aggregated in the replication status. This replication status is the second component of the protection group state. For more information about the states of protection groups, refer to "Monitoring the Runtime Status of the Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

If you add an Oracle ZFS Storage Appliance component to a protection group, Geographic Edition software creates a resource for each remote replication component. This resource monitors the status of replication for its remote replication component.

You can monitor the status of replication of this remote replication component by checking the `Status` and `Status Message` of this resource. Use the `clresourcegroup status` command to display resource status and the status message.

▼ How to Check the Runtime Status of Oracle ZFS Storage Appliance Replication

Note - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships and click the partnership name. For more information about Oracle Solaris Cluster Manager, see Chapter 13, "Using the Oracle Solaris Cluster GUI" in *Oracle Solaris Cluster System Administration Guide*.

1. **Access a node of the cluster where the protection group has been defined.**

   You must be assigned the Basic Solaris User RBAC rights profile to complete this procedure. For more information about RBAC, see "Securing Geographic Edition Software" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

2. **Check the runtime status of replication.**

   ```
   # geoadm status
   ```

   Refer to the `Protection Group` section of the output for replication information. The information that is displayed by this command includes the following:

   - Whether the local cluster is enabled for partnership participation

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

3. **Check the runtime status of replication for each Oracle ZFS Storage Appliance component.**

   # `clresourcegroup status` *zfssa-replication-resource-group*
   # `clresource status` *zfssa-replication-resource*

   Refer to the `Status` and `Status Message` fields for the remote replication component you want to check. See "Oracle ZFS Storage Appliance Replication Resource Group Runtime Status and Status Messages" on page 48 for a list of possible status values and status messages.

4. **List the status of components that are managed by Geographic Edition software.**

   # `clresource status -t ORCL.repzfssa`

Oracle ZFS Storage Appliance Replication Resource Group Runtime Status and Status Messages

The following table lists the `Status` and `Status Message` values that are returned by the `clresource status` command when the `State` of the Oracle ZFS Storage Appliance replication resource group is `Online`.

TABLE 4   Status and Status Messages of an Online Oracle ZFS Storage Appliance Replication Resource Group

Status	Status Message
Online	Sending update
Online	Receiving update
Online	Idle
Degraded	The most recent replication update failed because the target system has reached the maximum number of concurrent replication updates.
Degraded	The appliance failed to contact the remote peer. There might be a network connectivity issue or the management software on the target might have failed.

Status	Status Message
Degraded	A remote procedure call failed on the remote peer. The target system may be running incompatible software.
Degraded	ZFSSA cluster heads are not in clustered state. Perform a failback operation.
Faulted	The most recent replication update was cancelled by an administrator.
Faulted	The most recent replication update failed because replication is disabled globally or disabled for this package on the target appliance.
Faulted	The most recent replication update failed because there is insufficient space on this system to create a new project-level snapshot.
Faulted	The most recent replication update failed because the target is running incompatible software.
Faulted	The most recent replication update failed because the target package contains data from a previous replication update that could not be used for an incremental update.
Faulted	The most recent replication update failed because no replication package exists on the target for this replication action.
Faulted	The appliance could not verify the identity of the remote peer.
Unknown	The most recent replication update failed. No additional information is available. Check replication status on the target system. See the replication documentation for more details.
Unknown	Replication is disabled for the project.
Unknown	Continuous mode is set to `false` for the project.
Unknown	Failed to obtain replication key for the project.

For more information about these values, refer to the Oracle ZFS Storage Appliance documentation.

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

3

Migrating Services That Use Oracle ZFS Storage Appliance Remote Replication

This chapter provides information about migrating services for maintenance or as a result of cluster failure. This chapter contains the following sections:

Recovering Services to a Cluster on a System That Uses Oracle ZFS Storage Appliance Replication

This section contains the following information:

Overview of Recovering Services

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

The Geographic Edition software supports the following kinds of failback:

- **Failback-switchover.** During a failback-switchover, applications are brought online again on the original primary cluster after the data of the original primary cluster was resynchronized with the data on the secondary cluster.
- **Failback-takeover.** During a failback-takeover, applications are brought online again on the original primary cluster and use the current data on the original primary cluster. Any updates that occurred on the secondary cluster while it was acting as primary are discarded.

If you want to leave the new primary as the primary cluster and the original primary cluster as the secondary after the original primary restarts, you can resynchronize and revalidate the protection group configuration without performing a switchover or takeover.

▼ How to Perform a Failback-Switchover on a System That Uses Oracle ZFS Storage Appliance Replication

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

Note - The failback procedures apply only to clusters in a partnership. You need to perform the following procedure only once per partnership.

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

- If the original primary cluster had been down, the cluster has been booted and that the Geographic Edition infrastructure is enabled on the cluster. For more information about booting a cluster, see "Booting a Cluster" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
- The protection group on the current primary cluster has the `primary` role.
- The protection group on the original primary cluster has either the `primary` role or `secondary` role, depending on whether the original primary cluster can be reached during the takeover from the current primary cluster.

This procedure uses the example names `cluster-paris` for the original primary cluster and `cluster-newyork` for the current primary cluster.

1. **Synchronize replication from the `newyork` appliance to the `paris` appliance.**

This task is necessary to finish recovery if the cluster had experienced a complete site failure or a takeover. Data stores at `cluster-newyork` will have changed and will need to be replicated back to `cluster-paris` when it is put back in service.

Perform these steps for each project that is replicated.

   a. Access the Oracle ZFS Storage Appliance browser user interface (BUI) on the `cluster-newyork` site.

   b. Navigate to Shares > Projects and select the project being replicated.

   c. Select Replication for the project and click Update now.

     This executes a manual replication to synchronize the two sites.

2. Ensure that the protection group is stopped at the `cluster-paris` site.

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

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

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

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

   protection-group
     Specifies the name of the protection group

   c. Verify that the protection group is stopped.

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

3. Remove obsolete projects from the appliance at the `cluster-paris` site.

   a. Access the BUI on the `cluster-paris` site.

   b. Navigate to Shares > Projects.

   c. If any projects in the protection group are listed, manually delete them.

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

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

a. On cluster-paris, resynchronize the partnership.

phys-paris-1# **geops update** *partnership*

partnership

Specifies the name of the partnership

Note - Perform this step only once per partnership, even if you are performing a failback-switchover for multiple protection groups in the partnership.

For more information about synchronizing partnerships, see "Resynchronizing a Partnership" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

b. On cluster-paris, resynchronize each protection group.

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

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

For more information about synchronizing protection groups, see "Resynchronizing a Protection Group" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

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

phys-paris-1# **geopg validate** *protection-group*

protection-group

Specifies a unique name that identifies a single protection group

For more information, see "Validating a Protection Group" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

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

Because the protection group on cluster-paris has a role of secondary, the geopg start command does not restart the application on cluster-paris.

phys-paris-1# **geopg start -e global** *protection-group*

-e global

> Specifies the scope of the command. By specifying a global scope, the command operates on both clusters.

protection-group

> Specifies the name of the protection group.

Note - Do not use the -n option when performing a failback-switchover. The data must be synchronized from the current primary, cluster-newyork, to the current secondary, cluster-paris.

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

For more information about the geopg start command, see "How to Activate a Protection Group" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

7. **Confirm that the data is completely synchronized.**

The data is completely synchronized when the state of the protection group on cluster-newyork is OK. The protection group has a local state of OK when the appliance data store on cluster-newyork is being updated to the cluster-paris cluster.

To confirm that the state of the protection group on cluster-newyork is OK, use the following command:

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

Refer to the Protection Group section of the output.

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

```
# geoadm status
```

9. **On either cluster, perform a switchover from cluster-newyork to cluster-paris for each protection group.**

```
# geopg switchover [-f] -m cluster-paris protection-group
```

For more information, see "How to Switch Over Replication From the Primary Cluster to the Secondary Cluster" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

10. **Ensure that the switchover was performed successfully.**

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

```
# geoadm status
```

Check the runtime status of the application resource group and replication for each protection group.

```
# clresourcegroup status -v protection-group
```

Refer to the `Status` and `Status Message` fields that are presented for the remote replication component you want to check.

For more information about the runtime status of replication, see "Checking the Runtime Status of Oracle ZFS Storage Appliance Remote Replication" on page 46.

For more information on how to perform the switchover, see **Broken Link (Target ID: GNMHD)**.

▼ How to Perform a Failback-Takeover on a System That Uses Oracle ZFS Storage Appliance Replication

Use this procedure to restart an application on the original primary cluster and use the current data on the original primary cluster. Any updates that occurred on the secondary cluster while it was acting as primary are discarded.

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

Note - To resume using the data on the original primary you must not have replicated data from the new primary to the original primary cluster, `cluster-paris`, at any point after the takeover operation on the current primary cluster. To prevent replication between the current primary and the original primary, you must have used the `-n` option whenever you used the `geopg start` command.

Before You Begin   Ensure that the clusters have the following roles:

■ If the original primary cluster had been down, the cluster is booted and the Geographic Edition infrastructure is enabled on the cluster. For more information about booting a

cluster, see "Booting a Cluster" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

- The protection group on the current primary cluster has the `primary` role.
- The protection group on the original primary cluster has either the `primary` role or `secondary` role, depending on whether the original primary can be reached during the takeover from the current primary.

This procedure uses the example names `cluster-paris` for the original primary cluster and `cluster-newyork` for the current primary cluster.

1. Log in to the Oracle ZFS Storage appliance on the `cluster-paris` site.

2. On the `paris` appliance, remove the replication action for the project.

3. On the `paris` appliance, re-add the replication action and enable it with continuous mode.

The package is created on the original primary appliance, `paris`. The corresponding package is created on the original secondary appliance, `newyork`.

4. On the original secondary cluster, `cluster-newyork`, stop the protection group locally.

`phys-newyork-1# `**`geopg stop -e local`** *protection-group*

`-e local`

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

protection-group

Specifies the name of the protection group.

Note - Wait for the replica package to appear on `cluster-newyork` before you continue to the next step.

---

5. Make the protection group primary on `cluster-paris` and secondary on `cluster-newyork`.

- **If `cluster-paris` has the secondary role, run the following command from `cluster-paris`:**

  `phys-paris-1# `**`geopg takeover`** *protection-group*

- If `cluster-paris` has the primary role, run the following command from `cluster-newyork`:

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

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

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

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

   For more information, see "Validating a Protection Group" in *Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide*.

7. **From either cluster, start the protection group globally.**

   ```
   phys-paris-1# geopg start -e global protection-group
   ```

   The protection group on `cluster-paris` now has the `primary` role, and the protection group on `cluster-newyork` has the role of `secondary`. The application services are now online on `cluster-paris`.

   For more information, see "How to Activate a Protection Group" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

8. **Remove obsolete projects from the `newyork` appliance.**

   a. **Ensure that the protection group is activated and that all components are in the OK state on both clusters.**

      ```
      # geoadm status
      ```

   b. **Access the BUI on the `newyork` appliance.**

   c. **Navigate to Shares > Projects.**

   d. **If any projects in the protection group are listed, manually delete them.**

Recovering From an Oracle ZFS Storage Appliance Remote Replication Error

When an error occurs at the replication level, the error is reflected in the status of the resource in the replication resource group of the relevant remote replication component. This changed status appears in the Remote Replication status field in the output of the `geoadm status` command for that protection group.

This section contains the following procedures:

-
-

▼ How to Detect Remote Replication Errors

1. Check the status of the replication resources by using the `clresource status` command.

phys-paris-1# **clresource status -v** *zfssa-replication-resource*

zfssa-replication-resource

Specifies the name of the Oracle ZFS Storage Appliance resource.

For information about how different `Resource status` values map to actual replication pair states, see Table 4, "Status and Status Messages of an Online Oracle ZFS Storage Appliance Replication Resource Group," on page 48.

Running the `clresource status` command might return output similar to the following example:

```
…
-- Resources --

Resource Name        Node Name           State     Status Message
-------------        ---------           -----     --------------
Resource: zfssa-replication-resource        phys-paris-1      Online    Faulted  - The
most recent replication update was canceled by an administrator.
Resource: zfssa-replication-resource        phys-paris-2      Offline   Offline
…
```

2. **Display the aggregate resource status for all components in the protection group by using the `geoadm status` command.**

   For example, the output of the `clresource status` command in the preceding example indicates that the Oracle ZFS Storage Appliance replication state of the protection group is in the `Faulted` state on `cluster-paris`.

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

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

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

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

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

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

▼ How to Recover From an Oracle ZFS Storage Appliance Remote Replication Error

To recover from an error state, you might perform some or all of the steps in the following procedure.

1. **Use the procedures in the Oracle ZFS Storage Appliance documentation to determine the causes of the `Faulted` state.**

2. Recover from the `Faulted` state by using the Oracle ZFS Storage Appliance procedures.

If the recovery procedures change the state of the component, this state is automatically detected by the resource and is reported as a new protection group state.

3. Revalidate the protection group configuration.

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

protection-group

  Specifies the name of the Oracle ZFS Storage Appliance protection group.

- If the geopg validate command determines that the configuration is valid, the state of the protection group changes to reflect that fact.
- If the configuration is not valid, the geopg validate command returns a failure message.

4. Review the status of the protection group configuration.

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

5. Review the runtime status of the protection group.

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

Index

Z