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

Part No: E38181
June 2013, E38181–02

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

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

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

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

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

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

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

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

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

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

Contents

Preface

Oracle Solaris Cluster Geographic Edition Remote Replication Guide for Sun ZFS Storage Appliance provides procedures for administering Sun ZFS Storage Appliance from Oracle (Sun ZFS Storage) remote replication with Oracle Solaris Cluster Geographic Edition (Geographic Edition) software.

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

This document is intended for experienced system administrators with extensive knowledge of Oracle software and hardware. This document is not to be used as a planning or presales guide.

The instructions in this book assume knowledge of the Oracle Solaris operating system and Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

Using UNIX Commands

This document contains information about commands that are used to install, configure, or administer a Geographic Edition configuration. This document might not contain complete information on basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information:

- Online documentation for the Oracle Solaris software system
- Other software documentation that you received with your system
- Oracle Solaris OS man pages

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface	Description	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`

`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`**
		`Password:`

aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* save the file. **Note:** Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows UNIX system prompts and superuser prompts for shells that are included in the Oracle Solaris OS. In command examples, the shell prompt indicates whether the command should be executed by a regular user or a user with privileges.

TABLE P–2 Shell Prompts

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

Related Books

Information about related Geographic Edition topics is available in the documentation that is listed in the following table. All Geographic Edition documentation is available at `http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw`.

Topic	Documentation
Overview	*Oracle Solaris Cluster Geographic Edition Overview*
Installation	*Oracle Solaris Cluster Geographic Edition Installation Guide*
Data Replication	*Oracle Solaris Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility* *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy and Universal Replicator* *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard* *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite* *Oracle Solaris Cluster Geographic Edition Remote Replication Guide for Sun ZFS Storage Appliance*
System administration	*Oracle Solaris Cluster Geographic Edition System Administration Guide*

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table.

Topic	Documentation
Concepts	*Oracle Solaris Cluster Concepts Guide*
Hardware installation and administration	*Oracle Solaris Cluster 3.3 3/13 Hardware Administration Manual* and individual hardware administration guides
Software installation	*Oracle Solaris Cluster Software Installation Guide*
Data service installation and administration	*Oracle Solaris Cluster Data Services Planning and Administration Guide* and individual data service guides
Data service development	*Oracle Solaris Cluster Data Services Developer's Guide*
System administration	*Oracle Solaris Cluster System Administration Guide* *Oracle Solaris Cluster Quick Reference*
Software upgrade	*Oracle Solaris Cluster Upgrade Guide*

Topic	Documentation
Error messages	*Oracle Solaris Cluster Error Messages Guide*
Command and function references	*Oracle Solaris Cluster Reference Manual*
	Oracle Solaris Cluster Data Services Reference Manual

Access to Oracle Support

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

Getting Help

If you have problems installing or using Geographic Edition, contact your service provider and provide the following information.

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Oracle Solaris 10)
- The release number of Geographic Edition (for example, Geographic Edition 3.3)

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

Command	Function
`prtconf -v`	Displays the size of the system memory and reports information about peripheral devices
`psrinfo -v`	Displays information about processors
`showrev -p`	Reports which patches are installed
`prtdiag -v`	Displays system diagnostic information
`/usr/cluster/bin/clnode show-rev -v`	Displays Geographic Edition release and package version information for each node

Also have available the contents of the `/var/adm/messages` file.

◆ ◆ ◆ CHAPTER 1

Configuring and Administering Sun ZFS Storage Appliance Protection Groups

This chapter contains information about configuring and administering data replication with the remote replication feature of Sun 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 Sun ZFS Storage Appliance documentation and have the Sun ZFS Storage Appliance product and the latest Oracle Solaris SRUs installed on your system. For information about configuring a Sun ZFS Storage appliance, see the Sun ZFS Storage Appliance product documentation.

The chapter contains the following sections:

- "Overview of Configuring and Administering Remote Replication in a Sun ZFS Storage Appliance Protection Group" on page 12
- "Planning and Configuring Remote Replication With Sun ZFS Storage Appliance Software" on page 13
- "Creating, Modifying, Validating, and Deleting a Sun ZFS Storage Appliance Protection Group" on page 16
- "Administering Sun ZFS Storage Appliance Data-Replicated Components" on page 30
- "Administering Sun ZFS Storage Appliance Application Resource Groups" on page 36
- "Replicating a Sun ZFS Storage Appliance Protection Group Configuration to a Partner Cluster" on page 41
- "Activating and Deactivating a Sun ZFS Storage Appliance Protection Group" on page 43
- "Resynchronizing a Sun ZFS Storage Appliance Protection Group" on page 47
- "Checking the Runtime Status of Sun ZFS Storage Appliance Remote Replication" on page 48

Overview of Configuring and Administering Remote Replication in a Sun ZFS Storage Appliance Protection Group

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

TABLE 1–1 Tasks for Sun ZFS Storage Appliance Remote Replication

Task	Description
Plan the Sun ZFS Storage Appliance replication configuration.	See "Planning and Configuring Remote Replication With Sun ZFS Storage Appliance Software" on page 13.
Configure remote replication.	See "Configuring Remote Replication With Sun ZFS Storage Appliance Software" on page 17.
Create a protection group that is configured for Sun ZFS Storage Appliance replication.	See "How to Create and Configure a Sun ZFS Storage Appliance Protection Group" on page 22.
Add a replication component that is controlled by Sun ZFS Storage Appliance software.	See "How to Add a Remote Replication Component to a Sun ZFS Storage Appliance Protection Group" on page 30.
Add application resource groups to the protection group.	See "How to Add an Application Resource Group to a Sun ZFS Storage Appliance Protection Group" on page 37.
Replicate the protection group configuration to a secondary cluster.	See "How to Replicate the Sun ZFS Storage Appliance Protection Group Configuration to a Partner Cluster" on page 41.
Activate the protection group.	See "How to Activate a Sun ZFS Storage Appliance Protection Group" on page 43.
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 2, "Migrating Services That Use Sun ZFS Storage Appliance Remote Replication."
Check the runtime status of replication.	See "Checking the Runtime Status of Sun ZFS Storage Appliance Remote Replication" on page 48.
Detect failure.	See "Detecting Cluster Failure on a System That Uses Sun ZFS Storage Appliance Remote Replication" on page 51.
Migrate services by using a switchover.	See "Migrating Services That Use Sun ZFS Storage Appliance Remote Replication With a Switchover" on page 53.
Migrate services by using a takeover.	See "Forcing a Takeover on a System That Uses Sun ZFS Storage Appliance Remote Replication" on page 55.
Recover data after forcing a takeover.	See "Recovering Services to a Cluster on a System That Uses Sun ZFS Storage Appliance Replication" on page 57.

Planning and Configuring Remote Replication With Sun ZFS Storage Appliance Software

This section contains the following information:

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

Guidelines for Remote Replication With Sun ZFS Storage Appliance Software

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

- **Minimum version of Sun ZFS Storage Appliance software** — Sun ZFS 7000 Storage Appliance 2011.1.5 software is the minimum version supported with Geographic Edition software.
- **Minimum version of Oracle Solaris Cluster Geographic Edition software** – Both clusters in a Geographic Edition partnership that uses Sun ZFS Storage Appliance replication must run a minimum of Oracle Solaris Cluster 3.3 3/13 software.
- **Restriction for synchronous replication** – Continuous replication is asynchronous. Sun 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 Sun ZFS Storage Appliance software, see the Sun ZFS Storage online documentation at `https://`*appliance-hostname*`:215/wiki`, where *appliance-hostname* is the name of your storage appliance.

Overview of the Sun ZFS Storage Appliance Configuration File

Sun 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 10, "Script-Based Plug-Ins," in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Geographic Edition Properties to Set for Sun ZFS Storage Appliance Replication

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

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

Property Type	Properties
Script-based plug-in replicated 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 "General Properties of a Protection Group" 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 Sun 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 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, Modifying, Validating, and Deleting a Sun ZFS Storage Appliance Protection Group

This section contains the following topics:

- "Strategies for Creating Sun ZFS Storage Appliance Protection Groups" on page 16
- "Configuring Remote Replication With Sun ZFS Storage Appliance Software" on page 17
- "How to Install the Sun ZFS Storage Appliance Plug-In for Geographic Edition" on page 21
- "How to Create and Configure a Sun ZFS Storage Appliance Protection Group" on page 22
- "How to Modify a Sun ZFS Storage Appliance Protection Group" on page 25
- "Validating a Sun ZFS Storage Appliance Protection Group" on page 26
- "Debugging a Sun ZFS Storage Appliance Protection Group" on page 28
- "How to Delete a Sun ZFS Storage Appliance Protection Group" on page 28

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 System Administration Guide*.

Strategies for Creating Sun ZFS Storage Appliance Protection Groups

The following task maps describe the steps to perform:

TABLE 1–2 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 Sun ZFS Storage Appliance Software" on page 17.
2. Download and install the Sun ZFS Storage Appliance plug-in for Geographic Edition.	See "How to Install the Sun ZFS Storage Appliance Plug-In for Geographic Edition" on page 21.
3. Create the protection group from a cluster node.	See "How to Create and Configure a Sun ZFS Storage Appliance Protection Group" on page 22.
4. Add the replication component to the protection group.	See "How to Add a Remote Replication Component to a Sun ZFS Storage Appliance Protection Group" on page 30.
5. Start the protection group locally.	See "How to Activate a Sun ZFS Storage Appliance Protection Group" on page 43.
6. Add the application resource group to the protection group.	See "How to Add an Application Resource Group to a Sun ZFS Storage Appliance Protection Group" on page 37.
7. From the secondary cluster, retrieve the protection group configuration.	See "How to Replicate the Sun ZFS Storage Appliance Protection Group Configuration to a Partner Cluster" on page 41.
8. From the secondary cluster, activate the protection group locally.	See "How to Activate a Sun ZFS Storage Appliance Protection Group" on page 43.

Configuring Remote Replication With Sun ZFS Storage Appliance Software

This section describes the steps you must perform before you can configure Sun 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 18
- "How to Configure Oracle Solaris Cluster Resources on the Primary Cluster" on page 18
- "How to Configure Oracle Solaris Cluster Resources on the Secondary Cluster" on page 19

▼ 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 Sun ZFS Storage appliance.**

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

Configure the role with the following permissions:

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

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

▼ 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 Sun ZFS Storage appliance to be replicated.

If you create a ZFS storage pool, observe the following requirements and restrictions:

- Ensure that the zpool version on the cluster where you create the zpool is supported by the Oracle Solaris OS version of the partner cluster nodes. This is necessary so that the zpool can be imported by the partner cluster nodes, when that cluster becomes primary. You can do this by setting the zpool version to the default zpool version of the cluster that is running the earlier version of Oracle Solaris software.
- Mirrored and unmirrored ZFS storage pools are supported.
- ZFS storage pool spares are not supported with storage-based replication in a Geographic Edition configuration. The information about the spare that is stored in the storage pool results in the storage pool being incompatible with the remote system after it has been replicated.
- ZFS can be used with either Synchronous or Asynchronous mode. If you use Asynchronous mode, ensure that SRDF is configured to preserve write ordering, even after a rolling failure.

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

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

3. **Select Replication for the project and click Update Now.**

   This executes a manual replication to synchronize the two sites.

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

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

6. **Select Replication for the project and click the Reverse the Direction of Replication icon.**

   Replication is reversed.

7. **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 Sun ZFS Storage appliance to be replicated.

   If you create a ZFS storage pool, observe the following requirements and restrictions:

   - Ensure that the zpool version on the cluster where you create the zpool is supported by the Oracle Solaris OS version of the partner cluster nodes. This is necessary so that the zpool can be imported by the partner cluster nodes, when that cluster becomes primary. You can do this by creating the zpool on the cluster that runs the lowest version of Oracle Solaris software.
   - Mirrored and unmirrored ZFS storage pools are supported.
   - ZFS storage pool spares are not supported with storage-based replication in a Geographic Edition configuration. The information about the spare that is stored in the storage pool results in the storage pool being incompatible with the remote system after it has been replicated.

- ZFS can be used with either Synchronous or Asynchronous mode. If you use Asynchronous mode, ensure that SRDF is configured to preserve write ordering, even after a rolling failure.

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

8 **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 Sun 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*.

9 **Confirm that the application resource group is correctly configured by bringing it online and taking it offline again.**

```
phys-newyork-1# clresourcegroup online -emM app-resource-group
phys-newyork-1# clresourcegroup offline app-resource-group
```

10 **If you created a file systems and it is mounted, unmount the file system.**

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

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

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

12 **Reverse the replication on the primary site.**

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

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

c. **Select Replication for the project and click the Reverse the Direction of Replication icon**

▼ How to Install the Sun ZFS Storage Appliance Plug-In for Geographic Edition

Perform this procedure on all nodes of both clusters in the partnership.

1 **In a web browser, go to the Oracle ZFS Storage Appliance Plugin Download site, `http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html`.**

2 **Click the Accept License Agreement button.**

3 **Click the Download link for the latest Oracle Solaris Cluster Geographic Edition Plugin for Solaris 10.**

The zip file containing the `ORCLscgezfssacli` package is downloaded. Unzip the file to extract the package.

4 **As the root role, on all nodes of the global cluster, navigate to the directory containing the extracted `ORCLscgezfssacli` package and install it.**

The installation should be done within the global zone.

```
# pkgadd -d . ORCLscgezfssacli
```

▼ How to Create and Configure a Sun ZFS Storage Appliance Protection Group

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 Sun 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 18.
- You have installed the Sun ZFS Storage Appliance plug-in package, `ORCLscgezfssacli`. See "How to Install the Sun ZFS Storage Appliance Plug-In for Geographic Edition" on page 21.
- The local cluster is a member of a partnership.
- The protection group you are creating does not already exist on either partner cluster.

Perform this procedure from a node of the primary 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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun 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 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-ps
PG=zfssa-pg
REPCOMP=trancos
```

```
REPRS=zfssa-rep-rs
REPRG=zfssa-rep-rg
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 Sun 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-ps* is the partnership name and *zfssa-pg* is the protection group name:

```
phys-newyork-1# geopg get --partnership zfssa-ps zfssa-pg
```

7 **Verify the protection group configuration.**

```
phys-newyork-1# geoadm status
phys-newyork-1# clresource status zfssa-rep-rs
```

zfssa-rep-rs
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 Sun ZFS Storage Appliance Remote Replication From the Primary Cluster to the Secondary Cluster" on page 53.

Troubleshooting If you experience failures while performing this procedure, enable debugging. See "Debugging a Sun ZFS Storage Appliance Protection Group" on page 28.

▼ How to Modify a Sun ZFS Storage Appliance Protection Group

Before You Begin Before modifying the configuration of your protection group, ensure that the protection group you want to modify exists locally.

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

For more information about RBAC, see "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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 Sun ZFS Storage Appliance software.

---

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

```
# geopg set-prop -p property [-p...] pg-name
```

`-p` *property*
Specifies a property of the protection group.

For more information about the properties you can set, see Appendix A, "Standard Geographic Edition Properties," in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

pg-name
Specifies the name of the protection group.

This command modifies the properties of a protection group on all nodes of the local cluster. If the partner cluster contains a protection group of the same name, this command also propagates the new configuration information to the partner cluster.

For information about the properties you can set, see "Property Descriptions for Script-Based Plug-Ins" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

For information about the names and values that are supported by the Geographic Edition software, see Appendix B, "Legal Names and Values of Geographic Edition Entities," in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

Example 1–1 Modifying the Configuration of a Protection Group

The following example modifies the `Timeout` property of the `zfssa-pg` protection group.

```
# geopg set-prop -p Timeout=300 zfssa-pg
```

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

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

Validating a Sun ZFS Storage Appliance Protection Group

During protection group validation, the Sun ZFS Storage remote replication layer of the Geographic Edition software validates that the following conditions are met:

- The specified device group is a valid Oracle Solaris Cluster device group. The replication layer uses the `cldevicegroup list` command if the `Cluster_dgs` property is specified. The replication layer also verifies that the device group is of a valid type.
- The properties are valid for each Sun ZFS Storage component that has been added to the protection group.

▼ How to Validate a Sun ZFS Storage Appliance Protection Group

Before You Begin Ensure that the protection group you want to validate exists locally and that the common agent container is online on all nodes of both clusters in the partnership.

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

For more information about RBAC, see "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun ZFS Storage Appliance software.

---

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

Note – 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 pg-name
```

pg-name
Specifies a unique name that identifies a single protection group

Example 1–2 Validating the Configuration of a Protection Group

The following example validates the protection group `zfssa-pg`.

```
# geopg validate zfssa-pg
```

Troubleshooting 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.

- 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.

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

```
# clresource set -p Debug_level=N zfssa-rep-rs
```

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.

▼ How to Delete a Sun ZFS Storage Appliance Protection Group

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

- The protection group you want to delete exists locally.
- The protection group is offline on the local cluster.

Note – To keep the application resource groups online while deleting the protection group, you must remove the application resource groups from the protection group .

Perform this procedure from a node in the cluster where you want to delete 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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun ZFS Storage Appliance software.

---

2 **Delete the protection group from the local cluster.**

```
# geopg delete pg-name
```

pg-name
Specifies the name of the protection group

This command deletes the configuration of the protection group from the local cluster. The command also removes the replication resource group for each Sun ZFS Storage Appliance component in the protection group. This command does not alter the replication state of the component.

3 **To delete the protection group on the secondary cluster, repeat Step 1 and Step 2 on `cluster-newyork`.**

Example 1–3 Deleting a Sun ZFS Storage Appliance Protection Group

The following example deletes the protection group `zfssa-pg` from both partner clusters. The protection group is offline on both partner clusters. In this example, `cluster-paris` is the primary cluster and `cluster-newyork` is the partner cluster.

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

Example 1–4 Deleting a Sun ZFS Storage Appliance Protection Group While Keeping Application Resource Groups Online

The following example keeps online two application resource groups, `apprg1` and `apprg2`, while deleting their protection group, `zfssa-pg` from both partner clusters. First the application resource groups are removed from the protection group, then the protection group is deleted from the primary cluster phys-paris and the partner cluster phys-newyork.

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

Troubleshooting If the deletion is unsuccessful, the configuration status is set to Error. Fix the cause of the error, and rerun the geopg delete command.

Administering Sun ZFS Storage Appliance Data-Replicated Components

A protection group is the container for the application resource groups and 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 data-replicated 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 data-replicated components in a Sun ZFS Storage Appliance protection group:

- "How to Add a Remote Replication Component to a Sun ZFS Storage Appliance Protection Group" on page 30
- "Remote Replication Subsystem Process for Verifying the Replicated Component" on page 33
- "How to Modify a Sun ZFS Storage Appliance Data-Replicated Component" on page 34
- "How to Remove a Data-Replicated Component From a Sun ZFS Storage Appliance Protection Group" on page 35

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

Perform this procedure to add a replication component to an existing Sun 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 replication components to existing Sun ZFS Storage Appliance protection groups.

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

- You have installed the Sun ZFS Storage Appliance plug-in package, `ORCLscgezfssacli`. See "How to Install the Sun ZFS Storage Appliance Plug-In for Geographic Edition" on page 21.
- 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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun 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 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 data-replicated 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 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 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 data-replicated 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 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 replication component is performed during the scripted registration, an example is not provided here.

Troubleshooting If you have difficulties adding the component to the protection group, see "Debugging a Sun ZFS Storage Appliance Protection Group" on page 28.

Remote Replication Subsystem Process for Verifying the Replicated Component

During protection group validation, the Sun 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 replicated 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 Sun ZFS Storage Appliance 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 Sun ZFS Storage Appliance 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 Sun ZFS Storage Appliance remote replication, all the validation methods are the same.

▼ How to Modify a Sun ZFS Storage Appliance Data-Replicated Component

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

For more information about RBAC, see "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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 Sun ZFS Storage Appliance software.

---

2 **Modify the replication component.**

The following command modifies the properties of a remote replication component in a protection group on the local cluster. Then, the command propagates the new configuration to the partner cluster if the partner cluster contains a protection group with the same name.

```
# geopg modify-replication-component -p  property [-p...] zfssa-rep-cmp  zfssa-pg
```

`-p` *property*
Specifies the properties of the data-replicated component.

zfssa-rep-cmp
Specifies the name of the data-replicated component.

zfssa-pg
Specifies the name of the protection group that will contain the new data-replicated component.

Example 1–5 Modifying the Properties of a Sun ZFS Storage Appliance Remote Replication Component

The following example modifies the `Timeout` property of remote replication component `trancos` that is part of the appliance protection group, `zfssa-pg`.

```
# geopg modify-replication-component -p Timeout=215 trancos zfssa-pg
```

▼ How to Remove a Data-Replicated Component From a Sun ZFS Storage Appliance Protection Group

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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun ZFS Storage Appliance software.

---

2 **Remove the replicated component.**

```
# geopg remove-replication-component zfssa-rep-cmp zfssa-pg
```

zfssa-rep-cmp
Specifies the name of the data-replicated component.

zfssa-pg
Specifies the name of the protection group.

Example 1–6 Removing a Replicated Component From a Sun ZFS Storage Appliance Protection Group

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

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

Administering Sun 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 the following tasks:

- "How to Add an Application Resource Group to a Sun ZFS Storage Appliance Protection Group" on page 37
- "How to Delete an Application Resource Group From a Sun ZFS Storage Appliance Protection Group" on page 39

▼ How to Add an Application Resource Group to a Sun 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 – 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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun 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 **If the application resource group must have dependencies on resource groups and resources that are not managed by this protection group, ensure that the `External_dependencies_allowed` property of the protection group is set to `TRUE`.**

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

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

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

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

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

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

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

```
# geopg add-resource-group app-resource-group pg-name [-p external_dependencies_allowed=TRUE]
```

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

pg-name
Specifies the name of the protection group.

`-p external_dependencies_allowed=TRUE`
Permits the application resource group to have dependencies on resource groups and resources that are outside 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 app-resource-group
```

Example 1–7 Adding an Application Resource Group to a Sun 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.

▼ How to Delete an Application Resource Group From a Sun 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.

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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun ZFS Storage Appliance software.

---

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

```
# geopg remove-resource-group app-resource-group pg-name
```

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

pg-name
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 1–8 Deleting a Sun 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 a Sun 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 Sun ZFS Storage Appliance Protection Group Configuration to a Partner Cluster

Before You Begin Before you replicate the configuration of a Sun 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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun 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 ps-name pg-name
```

-s *ps-name*
: Specifies the name of the partnership from which the protection group configuration information is retrieved.

pg-name
: 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 1–9 Replicating a Sun 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 replicated component and application group that are configured almost identically to the replicated 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 a Sun ZFS Storage Appliance Protection Group" on page 28.

Activating and Deactivating a Sun ZFS Storage Appliance Protection Group

When you activate a protection group, the protection group assumes the role that you assigned to it during configuration. When you deactivate a protection group, its application resource groups are also unmanaged. You can activate or deactivate a protection group in the following ways:

- Globally – Activates or deactivates a protection group on both clusters where the protection group is configured
- On the primary cluster only – Secondary cluster remains inactive
- On the secondary cluster only – Primary cluster remains inactive

When a protection group is activated on the primary cluster, the application resource groups that are configured for the protection group are also started. The Geographic Edition software uses the following Oracle Solaris Cluster commands on the primary cluster to bring the resource groups online:

```
# clresourcegroup online -eM rglist
```

▼ How to Activate a Sun ZFS Storage Appliance 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 RBAC rights profile.**

For more information about RBAC, see "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun ZFS Storage Appliance software.

---

2 **Activate the protection group.**

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

-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.

pg-name
Specifies the name of the protection group.

The `geopg start` command uses the `clresourcegroup online -eM` *resourcegrouplist* command to bring resource groups and resources online. For more information about using this command, see the `clresourcegroup`(1CL) man page.

Example 1–10 Globally Activating a Sun ZFS Storage Appliance Protection Group

The following example globally activates a protection group.

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

Example 1–11 Locally Activating a Sun ZFS Storage Appliance Protection Group

The following 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 zfssa-pg
```

▼ How to Deactivate a Sun ZFS Storage Appliance Protection Group

This procedure deactivates the protection group on all nodes of the primary and secondary clusters, depending on the scope of the command. When you deactivate a protection group, its application resource groups are also unmanaged.

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

For more information about RBAC, see "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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 Sun ZFS Storage Appliance software.

---

2 **Deactivate the protection group.**

When you deactivate a protection group on the primary cluster, its application resource groups are also taken offline.

```
# geopg stop -e scope [-D] pg-name
```

`-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 where the protection group is deployed.

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

`-D`

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

If you omit this option, the replication subsystem and the protection group are both stopped. If the role of the protection group on the local cluster is primary, omitting the `-D` option also results in taking the application resource groups offline and putting them in an unmanaged state.

pg-name

Specifies the name of the protection group.

Example 1–12 Deactivating a Sun ZFS Storage Appliance Protection Group on All Clusters

The following example deactivates a protection group on all clusters.

```
# geopg stop -e global zfssa-pg
```

Example 1–13 Deactivating a Sun ZFS Storage Appliance Protection Group on a Local Cluster

The following example deactivates a protection group on the local cluster.

```
# geopg stop -e local zfssa-pg
```

Example 1–14 Stopping Sun ZFS Storage Appliance Remote Replication While Leaving the Protection Group Online

The following example stops replication on the local cluster only.

```
# geopg stop -e local -D zfssa-pg
```

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

```
# geopg stop -e local zfssa-pg
```

Example 1–15 Deactivating a Sun ZFS Storage Appliance Protection Group While Keeping Application Resource Groups Online

The following example keeps online two application resource groups, apprg1 and apprg2, while deactivating their protection group, zfssa-pg, on both clusters.

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

   ```
   # geopg remove-resource-group apprg1,apprg2 zfssa-pg
   ```
2. Deactivate the protection group.

   ```
   # geopg stop -e global zfssa-pg
   ```

Troubleshooting If the geopg stop command fails, run the geoadm status command to obtain the status of each component. For example, the configuration status might be set to Error depending on the cause of the failure. The protection group might remain activated even though some resource groups might be unmanaged. The protection group might be deactivated with replication running.

If the configuration status is set to Error, revalidate the protection group. See "Validating a Sun ZFS Storage Appliance Protection Group" on page 26.

Resynchronizing a Sun ZFS Storage Appliance Protection Group

You can resynchronize the configuration information of the local protection group with the configuration information that is retrieved from the partner cluster. You need to resynchronize a protection group when its `Synchronization` status in the output of the `geoadm status` command is `Error`.

For example, you might need to resynchronize protection groups after booting the cluster. For more information, see "Booting a Cluster" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Resynchronizing a protection group updates only entities that are related to Geographic Edition software.

▼ How to Resynchronize a Sun ZFS Storage Appliance Protection Group

Before You Begin The protection group must be deactivated on the cluster where you are running the `geopg update` command. For information about deactivating a protection group, see "Activating and Deactivating a Sun ZFS Storage Appliance Protection Group" on page 43.

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

For more information about RBAC, see "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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 Sun ZFS Storage Appliance software.

---

2 **Resynchronize the protection group.**

```
# geopg update pg-name
```

pg-name
Specifies the name of the protection group.

This command synchronizes the local Geographic Edition protection group configuration information on the local cluster with the protection group configuration information retrieved from the partner cluster.

Example 1–16 Resynchronizing a Sun ZFS Storage Appliance Protection Group

The following example resynchronizes a protection group.

```
# geopg update zfssa-pg
```

Checking the Runtime Status of Sun 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 a Sun ZFS Storage Appliance Protection Group" on page 22.

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 a Sun ZFS Storage Appliance Runtime Status" on page 48
- "How to Check the Runtime Status of Sun ZFS Storage Appliance Replication" on page 49
- "Sun ZFS Storage Appliance Replication Resource Group Runtime Status and Status Messages" on page 49

Overview of Displaying a Sun ZFS Storage Appliance Runtime Status

The status of each Sun ZFS Storage Appliance remote replication resource indicates the status of replication on a particular 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 a Sun ZFS Storage Appliance component to a protection group, Geographic Edition software creates a resource for each replication component. This resource monitors the status of replication for its replication component.

You can monitor the status of replication of this 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 Sun ZFS Storage Appliance Replication

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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun ZFS Storage Appliance component.

```
# clresourcegroup status zfssa-rep-rg
# clresource status zfssa-rep-rs
```

Refer to the `Status` and `Status Message` fields for the replication component you want to check. See "Sun ZFS Storage Appliance Replication Resource Group Runtime Status and Status Messages" on page 49 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
```

Sun 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 Sun ZFS Storage Appliance replication resource group is `Online`.

TABLE 1–3 Status and Status Messages of an Online Sun ZFS Storage Appliance Replication Resource Group

Status	Status Message
Online	Sending update

TABLE 1–3 Status and Status Messages of an Online Sun ZFS Storage Appliance Replication Resource Group *(Continued)*

Status	Status Message
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.
Degraded	A remote procedure call failed on the remote peer. The target system may be running incompatible software.
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 Sun ZFS Storage Appliance documentation.

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

◆ ◆ ◆ CHAPTER 2

Migrating Services That Use Sun 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:

- "Detecting Cluster Failure on a System That Uses Sun ZFS Storage Appliance Remote Replication" on page 51
- "Migrating Services That Use Sun ZFS Storage Appliance Remote Replication With a Switchover" on page 53
- "Actions Performed by the Geographic Edition Software During a Switchover" on page 54
- "Forcing a Takeover on a System That Uses Sun ZFS Storage Appliance Remote Replication" on page 55
- "Recovering Services to a Cluster on a System That Uses Sun ZFS Storage Appliance Replication" on page 57
- "Recovering From a Sun ZFS Storage Appliance Remote Replication Error" on page 66

Detecting Cluster Failure on a System That Uses Sun ZFS Storage Appliance Remote Replication

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

- "Detecting Primary Cluster Failure" on page 51
- "Detecting Secondary Cluster Failure" on page 52

Detecting Primary Cluster Failure

When the primary cluster for a protection group fails, the secondary cluster in the partnership detects the failure. The cluster that fails might be a member of more than one partnership, resulting in multiple failure detections.

The following actions take place when a primary cluster failure occurs. During a failure, the appropriate protection groups are in the `Unknown` state on the cluster that failed.

- Heartbeat failure is detected by a partner cluster.
- The heartbeat is activated in emergency mode to verify that the heartbeat loss is not transient and that the primary cluster has failed. The heartbeat remains in the `Online` state during this default time-out interval, while the heartbeat mechanism continues to retry the primary cluster.

  This query interval is set by using the `Query_interval` heartbeat property. If the heartbeat still fails after the interval you configured, a heartbeat-lost event is generated and logged in the system log. When you use the default interval, the emergency-mode retry behavior might delay heartbeat-loss notification for about nine minutes. Messages are displayed in the graphical user interface (GUI) and in the output of the `geoadm status` command.

  For more information about logging, see "Viewing the Geographic Edition Log Messages" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.
- If the partnership is configured for heartbeat-loss notification, then one or both of the following actions occurs:
  - An email is sent to the address specified in the `Notification_emailaddrs` property.
  - The script defined in `Notification_actioncmd` is executed.

  For more information about configuring heartbeat-loss notification, see "Configuring Heartbeat-Loss Notification" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Detecting Secondary Cluster Failure

When a secondary cluster for a protection group fails, a cluster in the same partnership detects the failure. The cluster that failed might be a member of more than one partnership, resulting in multiple failure detections.

During failure detection, the following actions take place:

- Heartbeat failure is detected by a partner cluster.
- The heartbeat is activated in emergency mode to verify that the secondary cluster is dead.
- When a failure is confirmed by the Geographic Edition software, the cluster notifies the administrator. The system detects all protection groups for which the cluster that failed was acting as secondary. The state of the appropriate protection groups is marked `Unknown`.

Migrating Services That Use Sun ZFS Storage Appliance Remote Replication With a Switchover

Perform a switchover of a Sun ZFS Storage protection group when you want to migrate services to the partner cluster in an orderly fashion.

A switchover consists of the following actions:

- Application services are unmanaged on the former primary cluster.
- The replication role is reversed and now continues to run from the new primary.
- Application services are brought online on the new primary cluster.

▼ How to Switch Over Sun ZFS Storage Appliance Remote Replication From the Primary Cluster to the Secondary Cluster

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

- Data replication is active between the primary cluster and the secondary cluster. That is, the replication is enabled from the source to the target Sun ZFS Storage appliances.
- The Geographic Edition replication resource for this appliance replication shows the `Online` state.
- The Geographic Edition software is running on both clusters.
- The secondary cluster is a member of a partnership.
- Both cluster partners can be reached.
- The overall state of the protection group is set to `OK`.
- All components on the primary and secondary sites are up and well.

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

For more information about RBAC, see "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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 Sun ZFS Storage Appliance software.

2 **Initiate the switchover.**

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

```
phys-paris-1# geopg switchover [-f] -m new-primary-cluster pg-name
```

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

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

pg-name
Specifies the name of the protection group.

Example 2–1 Forcing a Switchover From the Primary Cluster to the Secondary Cluster

The following example performs a switchover to the secondary cluster.

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

Actions Performed by the Geographic Edition Software During a Switchover

When you run the `geopg switchover` command, the software confirms that the secondary cluster does indeed allow the replication role reversal. The command creates copies of the replicated source projects and exports the shares for use on the target appliance to ensure that no projects and mount point conflicts can occur. The command then destroys such created clones on the target appliance and confirms that the actual reverse replication operation can be performed on the target appliance. The software performs the following actions on the original primary cluster:

- Takes offline the application resource groups in the protection group and places them in the `Unmanaged` state.

- Performs a switchover to the secondary cluster for each Sun ZFS Storage Appliance replication configuration in the protection group.

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

- Runs the script that is defined in the `RoleChange_ActionCmd` property
- Brings online all application resource groups in the protection group

If the command completes successfully, the secondary cluster becomes the new primary cluster for the protection group. The original primary cluster becomes the new secondary cluster. The application resource groups in the protection group are brought online on the new primary cluster and replication from the appliance that is connected from the new primary cluster to the new secondary cluster begins.

The `geopg switchover` command returns an error if any of the previous operations fails. Run the `geoadm status` command to view the status of each component. For example, the Configuration status of the protection group might be set to `Error`, depending on the cause of the failure. The protection group might be activated or deactivated.

If the Configuration status of the protection group is set to `Error`, revalidate the protection group by using the procedures that are described in "How to Validate a Sun ZFS Storage Appliance Protection Group" on page 26.

If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures that are described in "How to Resynchronize a Sun ZFS Storage Appliance Protection Group" on page 47.

Forcing a Takeover on a System That Uses Sun ZFS Storage Appliance Remote Replication

Perform a takeover when applications need to be brought online on the secondary cluster, regardless of whether the data is completely consistent between the primary Sun ZFS Storage appliance and the secondary appliance. The information in this section assumes that the protection group has been started.

The following steps occur after you initiate a takeover:

- If the former primary cluster can be reached and the protection group is not locked for notification handling or some other reason, the protection group is deactivated.
- Projects that are replicated in the appliance replication configurations, which are present in the protection group that is being taken over from the former primary cluster `cluster-paris`, are taken over by the new primary cluster.

Note – This data might not be consistent with the original replicas. Data replication from the new primary cluster to the former primary cluster is stopped.

- Application services are brought online on the new primary cluster.
- The protection group is activated without replication enabled.

For more details about takeover and the effects of the `geopg takeover` command, see Appendix C, "Disaster Recovery Administration Example," in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

For details about the possible conditions of the primary and secondary cluster before and after a takeover, see Appendix D, "Takeover Postconditions," in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

How to Force Immediate Takeover of Sun ZFS Storage Appliance Services by a Secondary Cluster

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

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

Perform this procedure from a node in the secondary 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 "Geographic Edition Software and RBAC" in *Oracle Solaris Cluster Geographic Edition System Administration 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 Sun ZFS Storage Appliance software.

2 **Initiate the takeover.**

```
phys-newyork-1# geopg takeover [-f] pg-name
```

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

pg-name
Specifies the name of the protection group.

Example 2–2 Forcing a Takeover by a Secondary Cluster

The following example forces the takeover of the protection group `zfssa-pg` by the secondary cluster `cluster-newyork`.

The node `phys-newyork-1` is the first node of 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*.

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

Next Steps For information about the state of the primary and secondary clusters after a takeover, see Appendix D, “Takeover Postconditions,” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

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

This section contains the following information:

- “Overview of Recovering Services” on page 57
- “How to Resynchronize and Revalidate the Protection Group Configuration” on page 58
- “How to Perform a Failback-Switchover on a System That Uses Sun ZFS Storage Appliance Replication” on page 60
- “How to Perform a Failback-Takeover on a System That Uses Sun ZFS Storage Appliance Replication” on page 64

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 Resynchronize and Revalidate the Protection Group Configuration

Use this procedure to resynchronize and revalidate data on the original primary cluster with the data on the current primary cluster.

Before You Begin Before you resynchronize and revalidate the protection group configuration, a takeover has occurred on the current primary cluster. Ensure that the clusters now have the following roles:

- If the original primary cluster had been down, the cluster has been 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 protection group could be reached during the takeover.

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

1 Resynchronize the original primary cluster with the current primary cluster.

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 ps-name
```

ps-name
Specifies the name of the partnership.

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

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

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

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

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

pg-name
Specifies the name of the protection group.

For more information about synchronizing protection groups, see "Resynchronizing a Sun ZFS Storage Appliance Protection Group" on page 47.

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

```
phys-paris-1# geopg validate pg-name
```

pg-name
Specifies a unique name that identifies a single protection group.

For more information, see "How to Validate a Sun ZFS Storage Appliance Protection Group" on page 26.

3 **On `cluster-paris`, activate each protection group.**

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

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

`-e global`
Specifies the scope of the command.

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

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

pg-name
Specifies the name of the protection group.

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

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

For more information about the `geopg start` command, see "How to Activate a Sun ZFS Storage Appliance Protection Group" on page 43.

4 **Confirm that the protection group configuration is OK.**

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

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

The protection group has a local state of `OK` when the Sun ZFS Storage Appliance components on `cluster-newyork` have a `Synchronized` pair state.

Refer to the `Protection Group` section of the output.

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

```
phys-newyork-1# clresource status -g zfssa-rep-rg
```

▼ How to Perform a Failback-Switchover on a System That Uses Sun 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 Sun 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 pg-name
```

pg-name
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 ps-name
```

ps-name
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 pg-name
```

For more information about synchronizing protection groups, see "Resynchronizing a Sun ZFS Storage Appliance Protection Group" on page 47.

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 pg-name
```

pg-name
Specifies a unique name that identifies a single protection group

For more information, see "How to Validate a Sun ZFS Storage Appliance Protection Group" on page 26.

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 pg-name
```

`-e global`
Specifies the scope of the command. By specifying a `global` scope, the command operates on both clusters.

pg-name
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 Sun ZFS Storage Appliance Protection Group" on page 43.

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 pg-name
```

For more information, see "How to Switch Over Sun ZFS Storage Appliance Remote Replication From the Primary Cluster to the Secondary Cluster" on page 53.

`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 pg-name
```

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

For more information about the runtime status of replication, see "Checking the Runtime Status of Sun ZFS Storage Appliance Remote Replication" on page 48.

▼ How to Perform a Failback-Takeover on a System That Uses Sun 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 Sun 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 pg-name
```

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

pg-name 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 pg-name
  ```

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

  ```
  phys-newyork-1# geopg update pg-name
  ```

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 pg-name
```

For more information, see "How to Validate a Sun ZFS Storage Appliance Protection Group" on page 26.

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

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

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 Sun ZFS Storage Appliance Protection Group" on page 43.

8 **Remove obsolete projects from the `paris` 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 a Sun 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 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" on page 66
- "How to Recover From a Sun ZFS Storage Appliance Remote Replication Error" on page 67

▼ 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-rep-rs
```

zfssa-rep-rs Specifies the name of the Sun ZFS Storage Appliance resource.

For information about how different `Resource status` values map to actual replication pair states, see Table 1–3.

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

```
...
-- Resources --

            Resource Name       Node Name           State     Status Message
            -------------       ---------           -----     --------------
  Resource: zfssa-rep-rs         phys-paris-1        Online    Faulted  - The
most recent replication update was canceled by an administrator.
  Resource: zfssa-rep-rs         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 Sun 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 a Sun 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 Sun ZFS Storage Appliance documentation to determine the causes of the `Faulted` state.**

2 **Recover from the `Faulted` state by using the Sun 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 pg-name
```

pg-name
Specifies the name of the Sun 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 pg-name
```

5 **Review the runtime status of the protection group.**

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

Index