

**Oracle® Solaris Cluster Geographic  
Edition Data Replication Guide for EMC  
Symmetrix Remote Data Facility**

**ORACLE®**

Part No: E39668  
July 2014, E39668-01



Copyright © 2004, 2014, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

Copyright © 2004, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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

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

# Contents

---

<b>Using This Documentation .....</b>	<b>7</b>
<b>1 Replicating Data With EMC Symmetrix Remote Data Facility Software .....</b>	<b>9</b>
Replicating Data in an SRDF Protection Group (Task Maps) .....	10
Initial Configuration of SRDF Software .....	11
Enabling the SRDF -symforce Option .....	11
Configuring Data Replication With SRDF Software on the Primary Cluster .....	11
Configuring Data Replication With SRDF Software on the Secondary Cluster .....	14
Configuring the Other Entities on the Secondary Cluster .....	15
<b>2 Creating SRDF Protection Groups .....</b>	<b>19</b>
Creating and Validating an SRDF Protection Group .....	19
Strategies for Creating SRDF Protection Groups .....	19
▼ How to Create and Configure an SRDF Protection Group .....	23
Creating an SRDF Protection Group to Support Oracle RAC .....	25
Validating an SRDF Protection Group .....	29
Adding Application Resource Groups to an SRDF Protection Group .....	31
Adding Data Replication Device Groups to an SRDF Protection Group .....	33
Validations Made by the Data Replication Subsystem .....	35
<b>3 Administering SRDF Protection Groups .....</b>	<b>39</b>
Administering SRDF Application Resource Groups .....	39
▼ How to Delete an Application Resource Group From an SRDF Protection Group .....	39
Administering SRDF Data Replication Device Groups .....	40
▼ How to Modify an SRDF Data Replication Device Group .....	41
▼ How to Delete a Data Replication Device Group From an SRDF Protection Group .....	42
Replicating the SRDF Protection Group Configuration to a Partner Cluster .....	43

▼ How to Replicate the SRDF Protection Group Configuration to a Partner Cluster .....	43
Checking the Runtime Status of SRDF Data Replication .....	45
Displaying an SRDF Runtime Status Overview .....	45
Displaying a Detailed SRDF Runtime Status .....	46
<b>4 Migrating Services That Use SRDF Data Replication .....</b>	<b>49</b>
Migrating Services That Use SRDF Data Replication With a Switchover .....	49
Validations That Occur Before a Switchover .....	50
Results of a Switchover From a Replication Perspective .....	50
Forcing a Takeover on a System That Uses SRDF Data Replication .....	51
Validations That Occur Before a Takeover .....	52
Results of a Takeover From a Replication Perspective .....	52
Recovering Services to a Cluster on a System That Uses SRDF Replication .....	53
▼ How to Perform a Failback-Switchover on a System That Uses SRDF Replication .....	54
▼ How to Perform a Failback-Takeover on a System That Uses SRDF Replication .....	57
Recovering From a Switchover Failure on a System That Uses SRDF Replication .....	61
Switchover Failure Conditions .....	62
Recovering From Switchover Failure .....	63
▼ How to Make the Original Primary Cluster Primary for an SRDF Protection Group .....	63
▼ How to Make the Original Secondary Cluster Primary for an SRDF Protection Group .....	64
Recovering From an SRDF Data Replication Error .....	65
▼ How to Detect Data Replication Errors .....	66
▼ How to Recover From an SRDF Data Replication Error .....	67
<b>A Geographic Edition Properties for SRDF .....</b>	<b>69</b>
SRDF Protection Group and Data Replication Component Properties .....	69
SRDF Cluster Resource Properties That Must Not Be Changed .....	70
<b>Index .....</b>	<b>73</b>

## Using This Documentation

---

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

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

## Product Documentation Library

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

## Access to Oracle Support

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

## Documentation Accessibility

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

## Feedback

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

## Replicating Data With EMC Symmetrix Remote Data Facility Software

---

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

The Geographic Edition software supports the use of certain modes of EMC Symmetrix Remote Data Facility (SRDF) software for data replication. The following modes are supported for disaster-recovery environments:

- Synchronous mode
- Asynchronous mode

---

**Note** - SRDF Adaptive mode is not supported for use with clustered data services. It does not guarantee data consistency in normal operations. The `domino` option, which can be set on SRDF device pairs, suspends writes to the primary site if writes to the secondary site fail. If you set the `domino` option on a device pair, any failure of secondary storage or of the communication link between the Geographic Edition partnered clusters might cause the complete loss of application data services. This is not a highly available configuration, and you should not configure the `domino` option on devices that will be used by clustered data services.

---

Before you can replicate data with SRDF software, you must be familiar with the SRDF documentation and have the SRDF product and the latest patches installed on your system. For information about installing the SRDF software, see the SRDF product documentation.

---

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

---

This chapter contains the following information for configuring and administering data replication with SRDF software for clusters using Geographic Edition software:

- [“Replicating Data in an SRDF Protection Group \(Task Maps\)” on page 10](#)
- [“Initial Configuration of SRDF Software” on page 11](#)

For information about creating and deleting data replication device groups, see [“Adding Data Replication Device Groups to an SRDF Protection Group” on page 33](#). For information

about obtaining a global and a detailed runtime status of replication, see [“Checking the Runtime Status of SRDF Data Replication”](#) on page 45.

## Replicating Data in an SRDF Protection Group (Task Maps)

This section summarizes the tasks for configuring and administering SRDF data replication in a protection group.

**TABLE 1-1** Configuration Tasks for SRDF Data Replication

Task	Description
Perform an initial configuration of the SRDF software.	See <a href="#">“Initial Configuration of SRDF Software”</a> on page 11.
Create a protection group that is configured for SRDF data replication.	See <a href="#">“How to Create and Configure an SRDF Protection Group”</a> on page 23 or <a href="#">“How to Create a Protection Group for Oracle Real Application Clusters”</a> on page 26
Add a device group that is controlled by SRDF.	See <a href="#">“How to Add a Data Replication Device Group to an SRDF Protection Group”</a> on page 33.
Add application resource groups to the protection group.	See <a href="#">“How to Add an Application Resource Group to an SRDF Protection Group”</a> on page 31.
Replicate the protection group configuration to a secondary cluster.	See <a href="#">“How to Replicate the SRDF Protection Group Configuration to a Partner Cluster”</a> on page 43.
Activate the protection group.	See <a href="#">“How to Activate a Protection Group”</a> in <a href="#">“Oracle Solaris Cluster Geographic Edition System Administration Guide”</a> .
Verify the protection group configuration.	Perform a trial a switchover or takeover and test some simple failure scenarios before bringing your system online. See <a href="#">Chapter 4, “Migrating Services That Use SRDF Data Replication”</a> . <b>Note</b> - You cannot perform personality swaps if you are running SRDF/Asynchronous data replication.
Check the runtime status of replication.	See <a href="#">“Checking the Runtime Status of SRDF Data Replication”</a> on page 45.

**TABLE 1-2** Administration Tasks for SRDF Data Replication

Task	Description
Detect failure.	See <a href="#">“Detecting Cluster Failure”</a> in <a href="#">“Oracle Solaris Cluster Geographic Edition System Administration Guide”</a> .
Migrate services by using a switchover.	See <a href="#">“Migrating Services That Use SRDF Data Replication With a Switchover”</a> on page 49. <b>Note</b> - You cannot perform personality swaps if you are running SRDF/Asynchronous data replication.
Migrate services by using a takeover.	See <a href="#">“Forcing a Takeover on a System That Uses SRDF Data Replication”</a> on page 51.
Recover data after forcing a takeover.	See <a href="#">“Recovering Services to a Cluster on a System That Uses SRDF Replication”</a> on page 53.

## Initial Configuration of SRDF Software

This section describes the tasks you need to perform to configure SRDF software on the primary and secondary clusters. It also includes information about the preconditions for creating SRDF protection groups.

- [“Enabling the SRDF -symforce Option” on page 11](#)
- [“Configuring Data Replication With SRDF Software on the Primary Cluster” on page 11](#)
- [“Configuring Data Replication With SRDF Software on the Secondary Cluster” on page 14](#)
- [“Configuring the Other Entities on the Secondary Cluster” on page 15](#)

Initial configuration of the primary and secondary clusters includes the following:

- Configuring an SRDF device group, `devgroup1`, with the required number of disks
- If using a raw-disk device group, configuring a raw-disk group, `rawdg`
- Configuring the file system, which includes creating the file system, creating mount points, and adding entries to the `/etc/vfstab` file
- Creating an application resource group, `apprg1`, which contains an `HASStoragePlus` resource

Geographic Edition software supports the hardware configurations that are supported by the Geographic Edition software. Contact your Oracle service representative for information about current supported Geographic Edition configurations.

## Enabling the SRDF -symforce Option

All nodes of both clusters must have the SRDF property `SYMAPI_ALLOW_RDF_SYMFORCE` enabled. This setting is required for proper function of certain `geopg` operations. Ensure that the SRDF / `var/symapi/config/options` file has the following entry:

```
SYMAPI_ALLOW_RDF_SYMFORCE=TRUE
```

See your EMC Symmetrix Remote Data Facility documentation for more information.

## Configuring Data Replication With SRDF Software on the Primary Cluster

This section describes the steps you must perform on the primary cluster before you can configure SRDF data replication with Geographic Edition software. The following information is in this section:

- [“Setting Up SRDF Device Groups” on page 12](#)
- [“Checking the Configuration of SRDF Devices” on page 13](#)
- [“How to Create an RDF1 Device Group” on page 13](#)

## Setting Up SRDF Device Groups

SRDF devices are configured in pairs. The mirroring relationship between the pairs becomes operational as soon as the SRDF links are online. If you have dynamic SRDF available, you have the capability to change relationships between R1 and R2 volumes in your device pairings on the fly without requiring a BIN file configuration change.

---

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

---

The EMC Symmetrix database file on each host stores configuration information about the EMC Symmetrix units attached to the host. The EMC Symmetrix global memory stores information about the pair state of operating EMC SRDF devices.

EMC SRDF device groups are the entities that you add to Geographic Edition protection groups to enable the Geographic Edition software to manage EMC Symmetrix pairs.

The SRDF device group can hold one of two types of devices:

- RDF1 source device, which acts as the primary
- RDF2 target device, which acts as the secondary

As a result, you can create two types of SRDF device group, RDF1 and RDF2. An SRDF device can be moved to another device group only if the source and destination groups are of the same group type.

You can create RDF1 device groups on a host attached to the EMC Symmetrix software that contains the RDF1 devices. You can create RDF2 device groups on a host attached to the EMC Symmetrix software that contains the RDF2 devices. You can perform the same SRDF operations from the primary or secondary cluster, using the device group that was built on that side.

When you add remote data facility devices to a device group, all of the devices must adhere to the following restrictions:

- The device must be an SRDF device.
- The device must be either an RDF1 or RDF2 type device, as specified by the device group type.
- The device must belong to the same SRDF group number.

- The SRDF device group configuration must be the same on all nodes of both the primary and secondary clusters. For example, if you have a device group DG1, which is configured as RDF1, on node1 of clusterA, then node2 of clusterA should also have a device group called DG1 with the same disk set. Also, clusterB should have an SRDF device group called DG1, which is configured as RDF2, defined on all nodes.

## Checking the Configuration of SRDF Devices

Before adding SRDF devices to a device group, use the `symrdf list` command to list the EMC Symmetrix devices configured on the EMC Symmetrix units attached to your host.

```
# symrdf list
```

By default, the command displays devices by their EMC Symmetrix device name, a hexadecimal number that the EMC Symmetrix software assigns to each physical device. To display devices by their physical host name, use the `pd` argument with the `symrdf` command.

```
# symrdf list pd
```

## ▼ How to Create an RDF1 Device Group

The following steps create a device group of type RDF1 and add an RDF1 EMC Symmetrix device to the group.

1. **Create a device group named `devgroup1`.**

```
phys-paris-1# symdg create devgroup1 -type rdf1
```

2. **Add an RDF1 device, with the EMC Symmetrix device name of `085`, to the device group on the EMC Symmetrix storage unit identified by the number `00000003264`.**

A default logical name of the form `DEV001` is assigned to the RDF1 device.

```
phys-paris-1# symld -g devgroup1 -sid 3264 add dev 085
```

**Next Steps** Create the Geographic Edition device groups, file systems, or ZFS storage pools you want to use, specifying the LUNs in the SRDF device group. You also need to create an `HASStoragePlus` resource for the device group, file system, or ZFS storage pool you use.

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

- 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 more information about creating device groups, file systems, and ZFS storage pools in a cluster configuration, see [“Oracle Solaris Cluster System Administration Guide”](#). For information about creating an HAStoragePlus resource, see [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

## Configuring Data Replication With SRDF Software on the Secondary Cluster

This section describes the steps you must complete on the secondary cluster before you can configure SRDF data replication in Geographic Edition software.

### ▼ How to Create the RDF2 Device Group on the Secondary Cluster

**Before You Begin** Before you can issue the SRDF commands on the secondary cluster, you need to create a RDF2 type device group on the secondary cluster that contains the same definitions as the RDF1 device group.

---

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

---

1. **Use the `symdg export` command to create a text file, `devgroup1.txt`, that contains the RDF1 group definitions.**

```
phys-paris-1# symdg export devgroup -f devgroup.txt -rdf
```

2. **Use the `rcp` or `ftp` command to transfer the file to the secondary cluster.**

```
phys-paris-1# rcp devgroup1.txt phys-newyork-1:/.
phys-paris-1# rcp devgroup1.txt phys-newyork-2:/.
```

3. **On the secondary cluster, use the `symdg import` command to create the RDF2 device group by using the definitions from the text file.**

Run the following command on each node in the newyork cluster.

```
# symdg import devgroup1 -f devgroup1.txt
```

```
Adding standard device 054 as DEV001...
Adding standard device 055 as DEV002...
```

## Configuring the Other Entities on the Secondary Cluster

Next, you need to configure any volume manager, the Geographic Edition device groups, and the highly available cluster file system.

### ▼ How to Replicate the Configuration Information From the Primary Cluster, When Using Raw-Disk Device Groups

1. **On the primary cluster, start replication for the devgroup1 device group.**

```
phys-paris-1# symrdf -g devgroup1 -noprompt establish
```

```
An RDF 'Incremental Establish' operation execution is in progress for device group
'devgroup1'. Please wait...
Write Disable device(s) on RA at target (R2).....Done.
Suspend RDF link(s).....Done.
Mark target (R2) devices to refresh from source (R1).....Started.
Device: 054 ..... Marked.
Mark target (R2) devices to refresh from source (R1).....Done.
Suspend RDF link(s).....Done.
Merge device track tables between source and target.....Started.
Device: 09C ..... Merged.
Merge device track tables between source and target.....Done.
Resume RDF link(s).....Done.
```

```
The RDF 'Incremental Establish' operation successfully initiated for device group
'devgroup1'.
```

2. **On the primary cluster, confirm that the state of the SRDF pair is synchronized.**

```
phys-newyork-1# symrdf -g devgroup1 verify
```

```
All devices in the RDF group 'devgroup1' are in the 'Synchronized' state.
```

3. **On the primary cluster, split the pair by using the symrdf split command.**

```
phys-paris-1# symrdf -g devgroup1 -noprompt split
```

```
An RDF 'Split' operation execution is in progress for device group 'devgroup1'.
Please wait...
```

```
Suspend RDF link(s).....Done.
Read/Write Enable device(s) on RA at target (R2).....Done.
```

The RDF 'Split' operation device group 'devgroup1'.

**4. Map the EMC disk drive to the corresponding DID numbers.**

You use these mappings when you create the raw-disk device group.

**a. Use the `symrdf` command to find devices in the SRDF device group.**

```
phys-paris-1# symrdf -g devgroup1 query
...
DEV001 00DD RW      0      3 NR 00DD RW      0      0 S.. Split
DEV002 00DE RW      0      3 NR 00DE RW      0      0 S.. Split
...
```

**b. Display detailed information about all devices.**

```
phys-paris-1# symdev show 00DD
...
Symmetrix ID: 000187990182

Device Physical Name      : /dev/rdisk/c6t5006048ACCC81DD0d18s2
Device Symmetrix Name     : 00DD
```

**c. Once you know the ctd label, use the `cldevice` command to see more information about that device.**

```
phys-paris-1# cldevice show c6t5006048ACCC81DD0d18

=== DID Device Instances ===

DID Device Name:                /dev/did/rdisk/d5
Full Device Path:
pemc3:/dev/rdisk/c8t5006048ACCC81DEFd18
Full Device Path:
pemc3:/dev/rdisk/c6t5006048ACCC81DD0d18
Full Device Path:
pemc4:/dev/rdisk/c6t5006048ACCC81DD0d18
Full Device Path:
pemc4:/dev/rdisk/c8t5006048ACCC81DEFd18
Replication:                    none
default_fencing:                global
```

In this example, you see that the ctd label `c6t5006048ACCC81DD0d18` maps to `/dev/did/rdisk/d5`.

**d. Repeat steps as needed for each of the disks in the device group and on each cluster.**

**5. Create the device group, file system, or ZFS storage pool you want to use.**

Use the LUNs in the SRDF device group.

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

- 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 more information, see [“Oracle Solaris Cluster System Administration Guide”](#).

**6. Create an HAStoragePlus resource for the device group, file system, or ZFS storage pool you will use.**

For more information, see [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

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

```
phys-newyork-1# clresourcegroup online -emM apprg1  
phys-newyork-1# clresourcegroup offline apprg1
```

**8. Unmount the file system.**

```
phys-newyork-1# umount /mounts/sample
```

**9. Take the Geographic Edition device group offline.**

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

**10. Reestablish the SRDF pair.**

```
phys-newyork-1# symrdf -g devgroup1 -noprompt establish
```

Initial configuration on the secondary cluster is now complete.



## Creating SRDF Protection Groups

---

This chapter contains information about for configuring data replication with SRDF software.

### Creating and Validating an SRDF Protection Group

This section contains the following topics:

- [“Strategies for Creating SRDF Protection Groups” on page 19](#)
- [“How to Create and Configure an SRDF Protection Group” on page 23](#)
- [“Creating an SRDF Protection Group to Support Oracle RAC” on page 25](#)
- [“Validating an SRDF Protection Group” on page 29](#)
- [“Adding Application Resource Groups to an SRDF Protection Group” on page 31](#)
- [“Adding Data Replication Device Groups to an SRDF Protection Group” on page 33](#)
- [“Validations Made by the Data Replication Subsystem” on page 35](#)

---

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

For more information, see [“Creating a Protection Group That Does Not Require Data Replication”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

---

### Strategies for Creating SRDF Protection Groups

Before you begin creating protection groups, consider which of the following strategies is best for you:

- Creating the protection group while the application remains online.
- Taking the application offline before creating the protection group.

The following sections describe the tasks for each strategy:

- [“Creating a Protection Group While the Application Is Offline” on page 20](#)
- [“Creating a Protection Group While the Application Is Online” on page 20](#)
- [“Example: Creating an SRDF Protection Group While the Application Remains Online” on page 21](#)

## Creating a Protection Group While the Application Is Offline

To create a protection group while the application resource groups is offline, complete the following tasks.

1. Create the protection group from a node on one cluster.  
For more information, see [“How to Create and Configure an SRDF Protection Group” on page 23](#).
2. Add the data replication device group to the protection group.  
For more information, see [“How to Add a Data Replication Device Group to an SRDF Protection Group” on page 33](#).
3. Take the application resource group offline.
4. Add the application resource group to the protection group.  
For more information, see [“How to Add an Application Resource Group to an SRDF Protection Group” on page 31](#).
5. On the other cluster, retrieve the protection group configuration.  
For more information, see [“How to Replicate the SRDF Protection Group Configuration to a Partner Cluster” on page 43](#).
6. From either cluster, start the protection group globally.  
For more information, see [“How to Activate a Protection Group” in “Oracle Solaris Cluster Geographic Edition System Administration Guide ”](#).

## Creating a Protection Group While the Application Is Online

This strategy allows you to create a protection group without any application outage. To add an existing application resource group to a new protection group without taking the application offline, complete the following tasks on the cluster where the application resource group is online.

1. Create the protection group from a cluster node.  
For more information, see [“How to Create and Configure an SRDF Protection Group” on page 23](#).
2. Add the data replication device group to the protection group.

For more information, see [“How to Add a Data Replication Device Group to an SRDF Protection Group” on page 33](#).

3. Start the protection group locally.

For more information, see [“How to Activate a Protection Group” in “Oracle Solaris Cluster Geographic Edition System Administration Guide ”](#).

4. Add the application resource group to the protection group.

For more information, see [“How to Add an Application Resource Group to an SRDF Protection Group” on page 31](#).

Next, complete the following tasks on the other cluster.

- Retrieve the protection group configuration.

For more information, see [“How to Replicate the SRDF Protection Group Configuration to a Partner Cluster” on page 43](#).

- Activate the protection group locally.

For more information, see [“How to Activate a Protection Group” in “Oracle Solaris Cluster Geographic Edition System Administration Guide ”](#).

## Example: Creating an SRDF Protection Group While the Application Remains Online

This example creates a protection group without taking the application offline.

In this example, the `apprg1` resource group is online on the `cluster-paris` cluster.

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

```
phys-paris-1# geopg create -d srdf -p Cluster_dgs=dg1 \
-o Primary -s paris-newyork-ps srdjpg
Protection group "srdjpg" has been successfully created
```

2. Add the device group, `devgroup1`, to the protection group.

```
phys-paris-1# geopg add-device-group devgroup1 rdjpg
```

3. Activate the protection group locally.

```
phys-paris-1# geopg start -e local srdjpg
Processing operation.... this may take a while...
Protection group "srdjpg" successfully started.
```

This command starts data replication.

4. Add an application resource group that is already online to the protection group.

```
phys-paris-1# geopg add-resource-group apprg1 srdjpg
Following resource groups were successfully inserted:
```

"apprg1"

5. Verify that the application resource group was added successfully.

```
phys-paris-1# geoadm list srdfpg
Cluster: cluster-paris
  Partnership "paris-newyork-ps"      : OK
  Partner clusters                    : newyork
  Synchronization                    : OK
  ICRM Connection                    : OK

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

  Protection group "srdfpg"          : Degraded
  Partnership                        : paris-newyork-ps
  Synchronization                    : OK

  Cluster cluster-paris              : Degraded
  Role                               : Primary
  Configuration                      : OK
  Data replication                   : Degraded
  Resource groups                    : OK

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

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

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

7. Activate the protection group locally on the partner cluster.

```
phys-newyork-1# geopg start -e local srdfpg
Processing operation... this may take a while...
Protection group "srdfpg" successfully started.
```

8. Verify that the protection group was successfully created and activated.

Running the `geoadm status` command on `cluster-paris` produces the following output:

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

Partnership "paris-newyork-ps"      : OK
```

```

Partner clusters                : newyork
Synchronization                : OK
ICRM Connection                : OK

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

Protection group "srdpg"       : OK
Partnership                   : paris-newyork-ps
Synchronization               : OK

Cluster cluster-paris         : OK
Role                           : Primary
Configuration                  : OK
Data replication               : OK
Resource groups                : OK

Cluster cluster-newyork       : OK
Role                           : Secondary
Configuration                  : OK
Data Replication               : OK
Resource Groups                : OK

```

## ▼ How to Create and Configure an SRDF Protection Group

---

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

---

**Before You Begin** Before you create a protection group, ensure that the following conditions are met:

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

---

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

---

You can also replicate the existing configuration of a protection group from a remote cluster to the local cluster. For more information, see [“Replicating the SRDF Protection Group Configuration to a Partner Cluster”](#) on page 43.

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

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

**2. Create a new protection group that uses SRDF replication.**

This command creates a protection group on all nodes of the local cluster.

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

*-s partnership*

Specifies the name of the partnership.

*-o local-role*

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

*-d srdf*

Specifies that the protection group data is replicated by the SRDF software.

*-p property-setting*

Specifies the properties of the protection group.

You can specify the following properties:

- **Description** – Describes the protection group.
- **Timeout** – Specifies the time-out period for the protection group in seconds.
- **NodeList** – Lists the host names of the machines that can be primary for the replication subsystem.
- **Cluster\_dgs** – Lists the device groups where the data is written. The Geographic Edition device groups must exist and have the same name on both the primary cluster and the secondary cluster.

To specify Solaris Volume Manager disk sets when a disk set name is not the same on both clusters, use the following syntax which includes the cluster names:

```
Cluster_dgs=local-cluster[disk-set]:remote-cluster[disk-set]
```

Include the brackets ([ ]) around all disk set names for each cluster. Separate multiple disk set names with a comma. For example:

```
Cluster_dgs=paris[dg1,dg2]:newyork[dg10,dg11]
```

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”.

*protection-group*

Specifies the name of the protection group.

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

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

**Example 2-1** Creating and Configuring an SRDF Protection Group

This example creates an SRDF protection group on `cluster-paris`, which is set as the primary cluster.

```
# geopg create -s paris-newyork-ps -o primary -d srdf \
-p Cluster_dgs="dg1:dg5,dg8:dg10" srdffg
```

**Example 2-2** Creating an SRDF Protection Group for Application Resource Groups That Are Online

This example creates an SRDF protection group, `srdffg`, for an application resource group, `resourcegroup1`, that is currently online on `cluster-newyork`.

1. Create the protection group without the application resource group.

```
# geopg create -s paris-newyork-ps -o primary -d srdf \
-p nodelist=phys-paris-1,phys-paris-2 srdffg
```

2. Activate the protection group.

```
# geopg start -e local srdffg
```

3. Add the application resource group.

```
# geopg add-resource-group resourcegroup1 srdffg
```

## Creating an SRDF Protection Group to Support Oracle RAC

Geographic Edition software supports Oracle Real Application Clusters (Oracle RAC) with SRDF software. This section contains the following information:

- [“Requirements to Support Oracle RAC With Data Replication Software” on page 26](#)
- [“How to Create a Protection Group for Oracle Real Application Clusters” on page 26](#)

## Requirements to Support Oracle RAC With Data Replication Software

Observe the following requirements when you configure Oracle RAC:

- Each Oracle Clusterware OCR and Voting Disk Location must be in its own device group on each cluster and cannot be replicated.
- Static data such as Oracle Clusterware and database binaries are not required to be replicated. But this data must be accessible from all nodes of both clusters.
- You must create storage resources for dynamic database files that are replicated in their own resource groups. These storage resources must be separate from the resource group that holds the storage resource for Oracle Clusterware.
- To be able to leave Oracle RAC infrastructure resource groups outside of Geographic Edition control, you must run Geographic Edition binaries on both cluster partners and set the Oracle RAC protection group `External_Dependency_Allowed` property to `true`.
- Do not add the Oracle Clusterware OCR and Voting Disk device group to the protection group's `cluster_dgs` property.
- Do not add Oracle RAC infrastructure resource groups to the protection group. Only add the `rac_server_proxy` resource group and resource groups for device groups that are replicated to the protection group. Also, you must set to `False` the `Auto_start_on_new_cluster` resource group property for the `rac_server_proxy` resource group and resource groups and for device groups that are replicated.
- When you use a cluster file system for an Oracle RAC file system, such as a flash recovery area, alert, or trace log files, you must manually create on both clusters a separate resource group that uses the `HASStoragePlus` resource to bring online these corresponding file systems. You must set a strong resource dependency from non-Oracle Clusterware `SUNW.ScalDeviceGroup` resources to this `HASStoragePlus` resource. Then add this `HASStoragePlus` resource group to the Oracle RAC protection group.

## ▼ How to Create a Protection Group for Oracle Real Application Clusters

**Before You Begin** Before you create a protection group for Oracle Real Application Clusters (Oracle RAC), ensure that the following conditions are met:

- Read [“Requirements to Support Oracle RAC With Data Replication Software” on page 26](#).
- The node list of the protection group must be the same as the node list of Oracle RAC framework resource group.

- If one cluster is running Oracle RAC on a different number of nodes than another cluster, ensure that all nodes on both clusters have the same resource groups defined.

### 1. Log in to a cluster node on the primary cluster.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

This command creates a protection group on all nodes of the local cluster.

```
# geopg create -s partnership \
-o local-role \
-d srdf \
-p External_Dependency_Allowed=true \
[-p property [-p...]] \
protection-group
```

`-s partnership`

Specifies the name of the partnership.

`-o local-role`

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

`-d srdf`

Specifies that the protection group data is replicated by the SRDF software.

`-p property-setting`

Specifies the properties of the protection group.

You can specify the following properties:

- `Description` – Describes the protection group.
- `External_Dependency_Allowed` - Specifies whether to allow any dependencies between resource groups and resources that belong to this protection group and resource groups and resources that do not belong to this protection group. For Oracle RAC, set this property to `true`.
- `Timeout` – Specifies the timeout period for the protection group in seconds.
- `NodeList` – Lists the host names of the machines that can be primary for the replication subsystem.
- `Cluster_dgs` – Specifies the device group where the data is written.

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”](#).

*protection-group*

Specifies the name of the protection group.

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

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

### 3. Add an SRDF device group to the protection group.

```
# geopg add-device-group [-p property [-p...]] protection-group
```

*-p property-setting*

Specifies the properties 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”.

*protection-group*

Specifies the name of the protection group.

### 4. Add to the protection group only the `rac_server_proxy` resource group and the resource groups for device groups that are replicated.

---

**Note** - Do not add the Oracle RAC framework resource group to the protection group. This ensures that, if the protection group becomes secondary on the node, the framework resource group does not become unmanaged. In addition, multiple Oracle RAC databases can be on the cluster, and the databases can be under Geographic Edition control or not under its control.

---

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

*resource-group*

Specifies a comma-separated list of resource groups to add to or delete from the protection group. The specified resource groups must already be defined.

The protection group must be online before you add a resource group. The `geopg add-resource-group` command fails when a protection group is offline and the resource group that is being added is online.

---

**Note** - If a protection group has already been started at the time that you add a resource group, the resource group remains unmanaged. You must start the resource group manually by running the `geopg start` command.

---

*protection-group*

Specifies the name of the protection group.

## Validating an SRDF Protection Group

When the `geoadm status` output displays that the Configuration status of a protection group is Error, you can validate the configuration by using the `geogg validate` command. This command checks the current state of the protection group and its entities.

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

During protection group validation, the SRDF data replication layer of the Geographic Edition software validates the following:

- The SYMCLI is installed on each of the nodes in the `NodeList` property.
- The specified device group is a valid Geographic Edition device group. The data replication layer uses the `cldevicegroup list` command if the `Cluster_dgs` property is specified. The data replication layer also verifies that the device group is of a valid type.
- The properties are valid for each SRDF device group that has been added to the protection group.

This section contains the following information:

- [“How the Data Replication Subsystem Validates the Device Group” on page 29](#)
- [“How to Validate an SRDF Protection Group” on page 30](#)

## How the Data Replication Subsystem Validates the Device Group

The Geographic Edition data replication layer validates the protection group's replication role against the configuration of the SRDF RDF1 and RDF2 devices. If the configurations do not match, the validation returns an error.

If the `Cluster_dgs` property is specified, then the data replication layer verifies that the device group specified is a valid Geographic Edition device group. The data replication layer also verifies that the device group is of a valid type.

---

**Note** - The device groups that are specified in the `Cluster_dgs` property must be written to only by applications that belong to the protection group. This property must not specify device groups that receive information from applications outside the protection group.

---

An Geographic Edition replication resource group is automatically created when the protection group is created.



---

**Caution** - Do not change, remove, or bring 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 Geographic Edition commands might result in unrecoverable failure.

---

## ▼ How to Validate an SRDF Protection Group

---

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

---

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

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

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

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

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

```
protection-group
```

Specifies a unique name that identifies a single protection group

**Example 2-3** Validating the Configuration of a Protection Group

This example validates the protection group `srdfpg`.

```
# geopg validate srdfpg
```

## Adding Application Resource Groups to an SRDF Protection Group

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

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

### ▼ How to Add an Application Resource Group to an SRDF Protection Group

**Before You Begin** You can add an existing resource group to the list of application resource groups for a protection group. Before you add an application resource group to a protection group, ensure that the following conditions are met:

- The protection group is defined.
- The resource group exists on both clusters and is in an appropriate state.
  - On the primary cluster, if the protection group is active, the application resource group is online.
  - On the secondary cluster, the application resource group is in the Unmanaged state.
- The HASToragePlus resource must exist in the application resource group, so that it can bring online the devices and mount the file systems.

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.

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

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

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

If necessary, change the property value to False.

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

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

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

If necessary, change the property value to True.

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

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

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

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

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

---

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

---

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

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

*resource-group-list*

Specifies the name or names of the application resource group.

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

*protection-group*

Specifies the name of the protection group.

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

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

If the application resource group configuration is OK on the local cluster, the application resource group gets added to the protection group on the local and remote cluster. If the subsequent configuration validation on the remote cluster does not result in a status of OK on the partner, the status is set to Error on the partner cluster.

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

**Example 2-4** Adding an Application Resource Group to a Protection Group

This example adds two application resource groups, apprg1 and apprg2, to srdffg.

```
# georg add-resource-group apprg1,apprg2 srdffg
```

## Adding Data Replication Device Groups to an SRDF Protection Group

This section provides procedures for adding an SRDF data replication device group. For details about configuring an SRDF data replication protection group, see [“How to Create and Configure an SRDF Protection Group”](#) on page 23.

### ▼ How to Add a Data Replication Device Group to an SRDF Protection Group

---

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

---

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

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

**2. Create a data replication device group in the protection group.**

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

```
# geogg add-device-group -p property [-p...] device-group protection-group
```

*-p property*

Specifies the properties of the data replication device group.

You can specify the following SRDF properties:

- **DG\_or\_CG** – Specifies that the device group is an SRDF device group.

You must set this property to DG.

- **R1SID** – Specifies the primary (RDF1) EMC Symmetrix ID of the EMC Symmetrix devices.

The data replication layer automatically sets the value of this property. You should specify the primary (RDF1) EMC Symmetrix ID of the EMC Symmetrix devices if you change the settings of the EMC Symmetrix host.

- **R2SID** – Specifies the secondary (RDF2) EMC Symmetrix ID of the EMC Symmetrix devices.

The data replication layer automatically sets the value of this property. You should specify the secondary (RDF2) EMC Symmetrix ID of the EMC Symmetrix devices if you change the settings of the EMC Symmetrix host.

*device-group*

Specifies the name of the new data replication device group.

*protection-group*

Specifies the name of the protection group that will contain the new data replication device group.

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

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

#### **Example 2-5** Adding a Data Replication Device Group to an SRDF Protection Group

This example adds the SRDF data replication device group to the `srdfpg` protection group.

```
# geogg add-device-group devgroup1 srdfpg
```

## Validations Made by the Data Replication Subsystem

When the SRDF device group is added to a protection group, the data replication layer makes the following validations.

- The specified device group name exists in the SRDF configuration.
- The replication role matches the SRDF protection group role.
- The EMC Symmetrix source, R1SID, and the EMC Symmetrix target, R2SID, can be reached.

When an SRDF device group is added to a protection group, a Geographic Edition data replication resource is automatically created by this command. This resource monitors data replication state. The name of the resource is `sc_geo_dr-SRDF-protection-group-device-group`, where *protection-group* is the name of the protection group and *device-group* is the name of the device group. This resource is placed in the corresponding Geographic Edition resource group, which is named `sc_geo_dr-SRDF-protection-group`.




---

**Caution** - Do not change, remove, or bring 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 Geographic Edition commands might result in unrecoverable failure.

---

This section provides the following information:

- [“How the State of the SRDF Device Group Is Validated” on page 35](#)
- [“Determining the State of an Individual SRDF Device Group” on page 35](#)
- [“Determining the Aggregate SRDF Device Group State” on page 36](#)
- [“Determining the SRDF Pair State” on page 37](#)

## How the State of the SRDF Device Group Is Validated

The state of each SRDF device group is mapped to the Geographic Edition resource group state. The `symrdf -g device-group query` command returns this state.

The remainder of this section describes the individual device group states and how these states are validated against the local role of the protection group.

## Determining the State of an Individual SRDF Device Group

An individual SRDF device group can be in one of the following states:

- Synchronized
- SynInProgress
- Failedover
- R1 Updated
- R1 UpdInProgress
- Split
- Suspended
- Partitioned
- Invalid

## Determining the Aggregate SRDF Device Group State

If a protection group contains only one SRDF device group, then the aggregate device group state is the same as the individual device group state.

When a protection group contains multiple SRDF device groups, the aggregate device group state is obtained as described in the following table.

**TABLE 2-1** Conditions That Determine the Aggregate Device Group State

Condition	Aggregate Device Group State
Any of the individual device group states are Invalid.	Invalid
Any of the individual device groups states are Partitioned and none of the individual device group states is Invalid.	Partitioned
One or more of the individual device groups states are Suspended and none of the individual device group states is Invalid, or Partitioned.	Suspended
One or more of the individual device groups states are Split and none of the individual device group states is Invalid, Partitioned, or Suspended.	Split
One or more of the individual device groups states are R1 UpdInProgress and none of the individual device group states is Invalid, Partitioned, Suspended, or Split.	R1 UpdInProgress
One or more of the individual device groups states are R1 Updated and none of the individual device group states is Invalid, Partitioned, Suspended, Split, or R1 UpdInProgress.	R1 Updated
One or more of the individual device groups states are Failedover and none of the individual device group states is Invalid, Partitioned, Suspended, Split, R1 UpdInProgress, or R1 Updated.	Failedover
One or more of the individual device groups states are SynInProgress and none of the individual device group states is Invalid, Partitioned, Suspended, Split, R1 UpdInProgress, R1 Updated, or Failedover.	SynInProgress
All of the individual device group states are Synchronized.	Synchronized

## Determining the SRDF Pair State

The resource status message reflects the role and state of the RDF pair. For example, the resource status and status message of `Faulted Split`, is reported when the RDF pair is in a `Split` state.

The RDF pair state is mapped to the associated resource status as described in the following table.

**TABLE 2-2** Mapping From the RDF Pair State to the Resource Status

Condition	Resource Status	Status Message
The RDF pair state is <code>Invalid</code> and the pair state is not <code>Incorrect Role</code> .	<code>Faulted</code>	<code>Invalid state</code>
The RDF pair state is <code>Partitioned</code> and the pair state is not <code>Incorrect Role</code> , or <code>Invalid</code> .	<code>Faulted</code>	<code>Partitioned</code>
The RDF pair state is <code>Suspended</code> and the pair state is not <code>Incorrect Role</code> , <code>Invalid</code> , or <code>Partitioned</code> .	<code>Faulted</code>	<code>Suspended</code>
The RDF pair state is <code>SyncInProgress</code> and the pair state is not <code>Incorrect Role</code> , <code>Invalid</code> , <code>Partitioned</code> , or <code>Suspended</code> .	<code>Degraded</code>	<code>SyncInProgress</code>
The RDF pair state is <code>R1 UpdInProgress</code> and the pair state is not <code>Incorrect Role</code> , <code>Invalid</code> , <code>Partitioned</code> , <code>Suspended</code> , or <code>SyncInProgress</code> .	<code>Faulted</code>	<code>R1 UpdInProgress</code>
The RDF pair state is <code>Split</code> and the pair state is not <code>Incorrect Role</code> , <code>Invalid</code> , <code>Partitioned</code> , <code>Suspended</code> , <code>SyncInProgress</code> , or <code>R1 UpdInProgress</code> .	<code>Faulted</code>	<code>Split</code>
The RDF pair state is <code>Failed over</code> and the pair state is not <code>Incorrect Role</code> , <code>Invalid</code> , <code>Partitioned</code> , <code>Suspended</code> , <code>SyncInProgress</code> , <code>R1 UpdInProgress</code> , or <code>Split</code> .	<code>Faulted</code>	<code>Failed over</code>
The RDF pair state is <code>R1 Updated</code> and the pair state is not <code>Incorrect Role</code> , <code>Invalid</code> , <code>Partitioned</code> , <code>Suspended</code> , <code>SyncInProgress</code> , <code>R1 UpdInProgress</code> , <code>Split</code> , or <code>Failed over</code> .	<code>Faulted</code>	<code>Replicating with role change</code>
The RDF pair state is <code>Synchronized</code> .	<code>Online</code>	<code>Replicating</code>

The state of the RDF pair determines the availability of consistent data in the partnership. When the state of the RDF resource on the primary or secondary cluster is `Degraded` or `Faulted`, the data volumes might not be synchronized even if the application can still write data from the primary volume to the secondary volume. The RDF pair will be in a `Partitioned` state and the invalid entries will be logged as the data is written to the primary volume. Manual recovery operations are required to resolve the error and resynchronize the data.



## Administering SRDF Protection Groups

---

This chapter contains the procedures for configuring and administering data replication with SRDF software. The chapter contains the following sections:

- [“Adding Application Resource Groups to an SRDF Protection Group” on page 31](#)
- [“Adding Data Replication Device Groups to an SRDF Protection Group” on page 33](#)
- [“Replicating the SRDF Protection Group Configuration to a Partner Cluster” on page 43](#)
- [“Checking the Runtime Status of SRDF Data Replication” on page 45](#)

### Administering SRDF Application Resource Groups

This section contains information about deleting an application resource group from an SRDF protection group.

#### ▼ How to Delete an Application Resource Group From an SRDF Protection Group

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

---

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

---

**Before You Begin** 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. Log in to a cluster node.**

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

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

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

If resource groups in the protection group have dependencies between them, you must remove all affected resource groups in the same `geogg remove-resource-group` command.

```
# geogg remove-resource-group resource-group-list protection-group
```

*resource-group-list*

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

*protection-group*

Specifies the name of the protection group.

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.

**Example 3-1** Deleting an Application Resource Group From a Protection Group

This example removes two application resource groups, `apprg1` and `apprg2`, from `srd fpg`.

```
# geogg remove-resource-group apprg1,apprg2 srd fpg
```

## Administering SRDF Data Replication Device Groups

This section provides the following information about administering SRDF data replication device groups:

- [“How the State of the SRDF Device Group Is Validated”](#) on page 35
- [“How to Modify an SRDF Data Replication Device Group”](#) on page 41

- [“How to Delete a Data Replication Device Group From an SRDF Protection Group” on page 42](#)

## ▼ How to Modify an SRDF Data Replication Device Group

---

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

---

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

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

### 2. Modify the device group.

This command modifies the properties of a device group 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 of the same name.

```
# geopg modify-device-group -p property [-p...] srdf-device-group protection-group
```

*-p property*

Specifies the properties of the data replication device 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”](#).

*srdf-device-group*

Specifies the name of the new data replication device group.

*protection-group*

Specifies the name of the protection group that will contain the new data replication device group.

#### **Example 3-2** Modifying the Properties of an SRDF Data Replication Device Group

This example modifies the R1SID property of a data replication device group that is part of an SRDF protection group.

```
# geopg modify-device-group -p R1SID=215 srdfdg srdfpg
```

## ▼ How to Delete a Data Replication Device Group From an SRDF Protection Group

---

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

---

**Before You Begin** You might delete a data replication device group from a protection group if you added a data replication device group to a protection group. Normally, after an application is configured to write to a set of disks, you would not change the disks.

Deleting a data replication device group does not stop replication or change the replication status of the data replication device group.

For information about deleting protection groups, refer to [“Deleting Protection Groups and Data Replication Components”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#). For information about deleting application resource groups from a protection group, refer to [“How to Delete an Application Resource Group From an SRDF Protection Group”](#) on page 39.

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

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

### 2. Remove the device group.

This command removes a device group from 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 of the same name.

```
# geopg remove-device-group device-group protection-group
```

*device-group*

Specifies the name of the data replication device group

*protection-group*

Specifies the name of the protection group

When a device group is deleted from an SRDF protection group, the corresponding Oracle Solaris Cluster resource, `sc_geo_dr-SRDF-protection-group-device-group`, is removed from the replication resource group. As a result, the deleted device group is no longer monitored. The replication resource group is removed when the protection group is deleted.

**Example 3-3** Deleting a Replication Device Group From an SRDF Protection Group

This example removes an SRDF data replication device group `srdfdg` from the `srdfpg` protection group.

```
# geopg remove-device-group srdfdg srdfpg
```

## Replicating the SRDF Protection Group Configuration to a Partner Cluster

After you have configured data 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 SRDF Protection Group Configuration to a Partner Cluster

---

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

---

**Before You Begin** Before you replicate the configuration of an SRDF 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 device groups in the protection group on the remote cluster exist on the local cluster.
- The system files on all nodes that can master the application have been updated for the application.
- The application resource groups in the protection group on the remote cluster exist on the local cluster.
- The `Auto_start_on_new_cluster` property of the application resource group is set to `False`. You can view this property by using the `clresourcegroup` command.

```
# clresourcegroup show -p auto_start_on_new_cluster apprg1
```

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

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

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

```
# clresourcegroup set -p Auto_start_on_new_cluster=False apprg1
```

- The `HASStoragePlus` resource exists in the application resource group to enable devices and mount file systems.

#### 1. Log in to `phys-newyork-1`.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see [“Securing Geographic Edition Software”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

`phys-newyork-1` 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”](#).

#### 2. Replicate the protection group configuration to the partner cluster by using the `geopg get` command.

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

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

`-s partnership`

Specifies the name of the partnership from which the protection group configuration information should be retrieved and the name of the partnership where the protection will be created locally.

`protection-group`

Specifies the name of the protection group.

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

---

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

---

**Example 3-4** Replicating the SRDF Protection Group to a Partner Cluster

This example replicates the configuration of `srdffpg` from `cluster-paris` to `cluster-newyork`.

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

## Checking the Runtime Status of SRDF Data Replication

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

- [“Displaying an SRDF Runtime Status Overview”](#) on page 45
- [“Displaying a Detailed SRDF Runtime Status”](#) on page 46

### Displaying an SRDF Runtime Status Overview

The status of each SRDF data replication resource indicates the status of replication on a particular device group. 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 ”](#).

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

## ▼ How to Check the Overall Runtime Status of Replication

---

**Note** - You can also accomplish this procedure by using the Oracle Solaris Cluster Manager GUI. Click Partnerships, then navigate to the components whose status you want to check. For more information about Oracle Solaris Cluster Manager, see [Chapter 13, “Using the Oracle Solaris Cluster GUI,”](#) in “Oracle Solaris Cluster System Administration Guide”.

---

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

You must be assigned the Basic Solaris User RBAC rights profile to complete this procedure. For more information about RBAC, see “[Securing Geographic Edition Software](#)” in “[Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide](#)”.

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

```
# geoadm status
```

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

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

**3. Check the runtime status of data replication for each SRDF device group.**

```
# clresourcegroup status sc_geo_dr-SRDF-protection-group  
# clresource status sc_geo_dr-SRDF-protection-group-device-group
```

Refer to the Status and Status Message fields for the data replication device group you want to check.

**See Also** For more information about these fields, see [Table 3-1](#).

## Displaying a Detailed SRDF Runtime Status

The Geographic Edition software internally creates and maintains one replication resource group for each protection group. The name of the replication resource group has the following format:

```
sc_geo_dr-SRDF-protection-group
```

If you add an SRDF device group to a protection group, Geographic Edition software creates a resource for each device group. This resource monitors the status of replication for its device group. The name of each resource has the following format, where *protection-group* is the name of the protection group and *device-group* is the name of the device group:

```
sc_geo_dr-SRDF-protection-group-device-group
```

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

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

**TABLE 3-1** Status and Status Messages of an Online SRDF Replication Resource Group

Status	Status Message
Online	Replicating
Degraded	Suspended
Degraded	SyncInProgress
Faulted	Incorrect role
Faulted	Invalid state
Faulted	Partitioned
Faulted	R1 UpdInProgress
Faulted	Split
Faulted	Failed over

For more information about these values, refer to the SRDF documentation.

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



## Migrating Services That Use SRDF Data Replication

---

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

- [“Migrating Services That Use SRDF Data Replication With a Switchover” on page 49](#)
- [“Forcing a Takeover on a System That Uses SRDF Data Replication” on page 51](#)
- [“Recovering Services to a Cluster on a System That Uses SRDF Replication” on page 53](#)
- [“Recovering From a Switchover Failure on a System That Uses SRDF Replication” on page 61](#)
- [“Recovering From an SRDF Data Replication Error” on page 65](#)

### Migrating Services That Use SRDF Data Replication With a Switchover

Perform a switchover of an SRDF protection group when you want to migrate services to the partner cluster in an orderly fashion. Basic Geographic Edition operations such as `geogg switchover`, perform a `symrdf swap` operation. The `symrdf swap` operation requires significantly more time for static RDF than dynamic RDF. Therefore, you might need to increase the value of the timeout property of the protection group when using static RDF.

A switchover consists of the following:

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

- Application services are brought online on the new primary cluster, `cluster-newyork`.

---

**Note** - You cannot perform personality swaps if you are running SRDF/Asynchronous data replication.

---

This section contains information about the following topics:

- [“Validations That Occur Before a Switchover” on page 50](#)
- [“Results of a Switchover From a Replication Perspective” on page 50](#)

## Validations That Occur Before a Switchover

When a switchover is initiated by using the `geopg switchover` command, the data replication subsystem runs several validations on both clusters. The switchover is performed only if the validation step succeeds on both clusters.

First, the replication subsystem checks that the SRDF device group is in a valid aggregate RDF pair state. Then, it checks that the local device group type on the target primary cluster, `cluster-newyork`, is RDF2. The `symrdf -g device-group-name -query` command returns the local device group's state. These values correspond to a RDF1 or RDF2 state. The following table describes the SRDF command that is run on the new primary cluster, `cluster-newyork`.

**TABLE 4-1** SRDF Switchover Validations on the New Primary Cluster

RDF Pair State	SRDF Switchover Command That Is Run on <code>cluster-newyork</code>
Synchronized	Suspends the RDF link.
R1Updated, Failedover, Suspended	The <code>symrdf swap</code> command switches the role.
Other RDF pair states	No command is run.

## Results of a Switchover From a Replication Perspective

After a successful switchover, at the data replication level the roles of the primary and secondary volumes have been switched. The pre-switchover RDF1 volumes become the RDF2 volumes. The pre-switchover RDF2 volumes become the RDF1 volumes. Data replication continues from the new RDF1 volumes to the new RDF2 volumes.

The `Local-role` property of the protection group is also switched regardless of whether the application could be brought online on the new primary cluster as part of the switchover

operation. On the cluster on which the protection group had a Local role of Secondary, the Local-role property of the protection group becomes Primary. On the cluster on which the protection group had a Local-role of Primary, the Local-role property of the protection group becomes Secondary.

## Forcing a Takeover on a System That Uses SRDF Data 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 volume and the secondary volume. The information in this section assumes that the protection group has been started.

The following steps occur after a takeover is initiated:

- If the former primary cluster, `cluster-paris`, can be reached and the protection group is not locked for notification handling or some other reason, the application services are taken offline on the former primary cluster.

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

- Data volumes of the former primary cluster, `cluster-paris`, are taken over by the new primary cluster, `cluster-newyork`.

---

**Note** - This data might be inconsistent with the original primary volumes. After the takeover, data replication from the new primary cluster, `cluster-newyork`, to the former primary cluster, `cluster-paris`, is stopped.

---

- Application services are brought online on the new primary cluster, `cluster-newyork`.

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 takeover, see [Appendix D, “Takeover Postconditions,”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

The following sections describe the steps you must perform to force a takeover by a secondary cluster.

- [“Validations That Occur Before a Takeover” on page 52](#)
- [“Results of a Takeover From a Replication Perspective” on page 52](#)

## Validations That Occur Before a Takeover

When a takeover is initiated by using the `geopg takeover` command, the data replication subsystem runs several validations on both clusters. These steps are conducted on the original primary cluster only if the primary cluster can be reached. If validation on the original primary cluster fails, the takeover still occurs.

First, the replication subsystem checks that the SRDF device group is in a valid aggregate RDF pair state. The SRDF commands that are used for the takeover are described in the following table.

**TABLE 4-2** SRDF Takeover Validations on the New Primary Cluster

Aggregate RDF Pair State	Protection Group Local Role	SRDF Takeover Commands That Are Run on <code>cluster-newyork</code>
FailedOver	Primary	<pre>symrdf \$option \$dg write_disable r2 symrdf -g dg suspend symrdf \$option \$dg rw_enable r1</pre>
FailedOver	Secondary	No command is run.
Synchronized, Suspended, R1 Updated, Partitioned	All	<code>symrdf -g dg failover</code>

## Results of a Takeover From a Replication Perspective

From a replication perspective, after a successful takeover, the `Local-role` property of the protection group is changed to reflect the new role, regardless of whether the application could be brought online on the new primary cluster as part of the takeover operation. On `cluster-newyork`, where the protection group had a `Local-role` of `Secondary`, the `Local-role` property of the protection group becomes `Primary`. On `cluster-paris`, where the protection group had a `Local-role` of `Primary`, the following might occur:

- If the cluster can be reached, the `Local-role` property of the protection group becomes `Secondary`.

- If the cluster cannot be reached, the `Local - role` property of the protection group remains `Primary`.

If the takeover is successful, the applications are brought online. You do not need to run a separate `geopg start` command.




---

**Caution** - After a successful takeover, data replication between the new primary cluster, `cluster-newyork`, and the old primary cluster, `cluster-paris`, is stopped. If you want to run a `geopg start` command, you must use the `-n` option to prevent replication from resuming.

---

## Recovering Services to a Cluster on a System That Uses SRDF Replication

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

Geographic Edition software supports the following two kinds of failback:

- **Failback-switchover.** During a failback-switchover, applications are brought online again on the original primary cluster, `cluster-paris`, after the data of the original primary cluster was resynchronized with the data on the secondary cluster, `cluster-newyork`.

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

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

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

This section contains information about the following topics:

- [“How to Perform a Failback-Switchover on a System That Uses SRDF Replication”](#) on page 54

- [“How to Perform a Failback-Takeover on a System That Uses SRDF Replication” on page 57](#)

## ▼ How to Perform a Failback-Switchover on a System That Uses SRDF 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`. The clusters have the following roles:

- If the original primary cluster, `cluster-paris`, has been down, confirm that the cluster is booted and that the Geographic Edition infrastructure is enabled on the cluster. For more information about booting a cluster, see [“Booting a Cluster” in “Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).
- The protection group on `cluster-newyork` has the primary role.
- The protection group on `cluster-paris` has either the primary role or secondary role, depending on whether `cluster-paris` can be reached during the takeover from `cluster-newyork`.

### 1. Ensure that the RDF1 role is not in the `split` state on `cluster-paris`.

This task is necessary to finish recovery if the cluster had experienced a complete site failure.

- a. On `cluster-paris`, the original primary, display the role and state of the data replication.

```
phys-paris-1# symrdf -g device-group query
```

- b. If the role is `RDF1` and is in the `split` state, fail over the device group.

```
phys-paris-1# symrdf -g device-group failover
```

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

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

**a. On cluster-paris, resynchronize the partnership.**

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

```
partnership
```

Specifies the name of the partnership

---

**Note** - You need to 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. Determine whether the protection group on the original primary cluster, cluster-paris, is active.**

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

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

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

**d. Verify that the protection group is stopped.**

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

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

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

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

```
protection-group
```

Specifies the name of the protection group

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

**3. 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 an error state.

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

*protection-group*

Specifies a unique name that identifies a single protection group

For more information, see [“Validating a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

**4. 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 local protection-group
```

`-e local`

Specifies the scope of the command.

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

*protection-group*

Specifies the name of the protection group.

---

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

---

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

For more information about the `geopg start` command, see [“How to Activate a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

**5. 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 SRDF device groups on `cluster-newyork` have a Synchronized RDF pair state.

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.

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

```
# geoadm status
```

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

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

For more information, see [“How to Switch Over Replication From the Primary Cluster to the Secondary Cluster”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

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

8. **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 data replication for each SRDF protection group.

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

Refer to the `Status` and `Status Message` fields that are presented for the data replication device group you want to check. For more information about these fields, see [Table 2-1](#).

For more information about the runtime status of data replication, see [“Checking the Runtime Status of SRDF Data Replication”](#) on page 45.

## ▼ How to Perform a Failback-Takeover on a System That Uses SRDF Replication

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

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

---

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

---

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

- If the original primary cluster, `cluster-paris`, has been down, confirm that the cluster is booted and that the Geographic Edition infrastructure is enabled on the cluster. For more information about booting a cluster, see [“Booting a Cluster”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).
- The protection group on `cluster-newyork` has the primary role.
- The protection group on `cluster-paris` has either the primary role or secondary role, depending on whether `cluster-paris` can be reached during the takeover from `cluster-newyork`.

**1. Ensure that the RDF1 role is not in the split state on cluster-paris.**

This task is necessary to finish recovery if the cluster had experienced a complete site failure.

**a. On cluster-paris, display the role and state of the data replication.**

```
phys-paris-1# symrdf -g device-group query
```

**b. If the role is RDF1 and is in the split state, fail over the device group.**

```
phys-paris-1# symrdf -g device-group failover
```

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

`cluster-paris` forfeits its own configuration and replicates the `cluster-newyork` configuration locally.

**a. On cluster-paris, resynchronize the partnership.**

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

```
partnership
```

Specifies the name of the partnership

---

**Note** - You need to perform this step only once per partnership, even if you are performing a failback-takeover for multiple protection groups in the partnership.

---

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

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

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

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

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

**d. Verify that the protection group is stopped.**

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

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

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

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

```
protection-group
```

Specifies the name of the protection group

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

**3. 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 an error state.

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

```
protection-group
```

Specifies a unique name that identifies a single protection group

For more information, see [“Validating a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition Installation and Configuration Guide”](#).

**4. On cluster-paris, activate each protection group in the secondary role without data replication.**

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

---

**Note** - You must use the `-n` option which specifies that data replication should not be used for this protection group. If this option is omitted, data replication starts at the same time as the protection group.

---

```
phys-paris-1# geopg start -e local -n protection-group
```

`-e local`

Specifies the scope of the command.

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

`-n`

Specifies that data replication should not be used for this protection group. If this option is omitted, data replication starts at the same time as the protection group.

*protection-group*

Specifies the name of the protection group

For more information, see [“How to Activate a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

Replication from `cluster-newyork` to `cluster-paris` is not started because the `-n` option is used on `cluster-paris`.

#### 5. On `cluster-paris`, initiate a takeover for each protection group.

```
phys-paris-1# geopg takeover [-f] protection-group
```

`-f`

Forces the command to perform the operation without your confirmation

*protection-group*

Specifies the name of the protection group

For more information about the `geopg takeover` command, see [“How to Force Immediate Takeover of a Protection Group by a Secondary Cluster”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

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`.

#### 6. On `cluster-newyork`, activate each protection group.

At the end of step 4, the local state of the protection group on `cluster-newyork` is `Offline`. To start monitoring the local state of the protection group, you must activate the protection group on `cluster-newyork`.

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

```
phys-newyork-1# geopg start -e local [-n] protection-group
```

`-e local`

Specifies the scope of the command.

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

`-n`

Prevents the start of data replication at protection group startup.

If you omit this option, the data replication subsystem starts at the same time as the protection group.

*protection-group*

Specifies the name of the protection group.

For more information about the `geopg start` command, see [“How to Activate a Protection Group”](#) in [“Oracle Solaris Cluster Geographic Edition System Administration Guide”](#).

## 7. Ensure that the takeover 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
```

---

**Note** - If you used the `-n` option in step 5 to prevent data replication from starting, the “Data replication” status will not be in the OK state.

---

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

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

Refer to the `Status` and `Status Message` fields that are presented for the data replication device group you want to check. For more information about these fields, see [Table 2-1](#).

For more information about the runtime status of data replication, see [“Checking the Runtime Status of SRDF Data Replication”](#) on page 45.

## Recovering From a Switchover Failure on a System That Uses SRDF Replication

Basic Geographic Edition operations such as `geopg switchover`, perform a `symrdf swap` operation at the SRDF data replication level. In SRDF terminology, a switchover is called a *swap*. The `symrdf swap` operation requires significantly more time for static RDF than dynamic

RDF. Therefore, you might need to increase the value of the timeout property of the protection group when using static RDF.

If all of the SRDF commands return a value of 0, the switchover is successful. In some cases, a command might return an error code (a value other than 0). These cases are considered switchover failures.

If a switchover failure occurs, the secondary volumes might not be fully synchronized with the primary volumes. Geographic Edition software does not start the applications on the new intended primary cluster in a switchover failure scenario.

The remainder of this section describes the initial conditions that lead to a switchover failure and how to recover from a switchover failure.

This section contains information about the following topics:

- [“Switchover Failure Conditions” on page 62](#)
- [“Recovering From Switchover Failure” on page 63](#)
- [“How to Make the Original Primary Cluster Primary for an SRDF Protection Group” on page 63](#)
- [“How to Make the Original Secondary Cluster Primary for an SRDF Protection Group” on page 64](#)

## Switchover Failure Conditions

This section describes a switchover failure scenario. In this scenario, `cluster-paris` is the original primary cluster and `cluster-newyork` is the original secondary cluster.

A switchover switches the services from `cluster-paris` to `cluster-newyork` as follows:

```
phys-newyork-1# geopg switchover -f -m cluster-newyork srdpfg
```

While processing the `geopg switchover` command, the `symrdf swap` command runs and returns errors for the SRDF device group, `devgroup1`. As a result, the `geopg switchover` command returns the following failure message:

```
Processing operation... this may take a while ....  
"Switchover" failed for the following reason:  
Switchover failed for SRDF DG devgroup1
```

After this failure message has been issued, the two clusters are in the following states:

```
cluster-paris:  
srdpfg role: Secondary  
cluster-newyork:  
srdpfg role: Secondary  
  
phys-newyork-1# symdg list
```

D E V I C E      G R O U P S							
Number of Name	Type	Valid	Symmetrix ID	Devs	GKs	BCVs	VDEVs
devgroup1	RDF1	Yes	000187401215	2	0	0	0
devgroup2	RDF2	Yes	000187401215	6	0	0	0

## Recovering From Switchover Failure

This section describes procedures to recover from the failure scenario described in the previous section. These procedures bring the application online on the appropriate cluster.

1. Place the SRDF device group, `devgroup1`, in the `Split` state.

Use the `symrdf split` commands to place the device groups that are in the protection group on both `cluster-paris` and `cluster-newyork` in the `Split` state.

```
phys-newyork-1# symrdf -g devgroup1 split
```

2. Make one of the clusters `Primary` for the protection group.

Make the original primary cluster, `cluster-paris`, `Primary` for the protection group if you intend to start the application on the original primary cluster. The application uses the current data on the original primary cluster.

Make the original secondary cluster, `cluster-newyork`, `Primary` for the protection group if you intend to start the application on the original secondary cluster. The application uses the current data on the original secondary cluster.




---

**Caution** - Because the `symrdf swap` command did not perform a swap, the data volumes on `cluster-newyork` might not be synchronized with the data volumes on `cluster-paris`. If you intend to start the application with the same data as appears on the original primary cluster, you must not make the original secondary cluster `Primary`.

---

## ▼ How to Make the Original Primary Cluster Primary for an SRDF Protection Group

1. Deactivate the protection group on the original primary cluster.

```
phys-paris-1# geopg stop -e Local srdpfg
```

2. Resynchronize the configuration of the protection group.

This command updates the configuration of the protection group on `cluster-paris` with the configuration information of the protection group on `cluster-newyork`.

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

After the `geopg update` command run successfully, `srdfpg` has the following role on each cluster:

```
cluster-paris:
srdfpg role: Primary
cluster-newyork:
srdfpg role: secondary
```

3. **Determine whether the device group has the RDF1 role on the original primary cluster.**

```
phys-paris-1# symdg list | grep devgroup1
```

4. **If the device group does not have the RDF1 role on the original primary cluster, run the `symrdf swap` command so that the device group, `devgroup1`, resumes the RDF1 role.**

```
phys-paris-1# symrdf -g devgroup1 failover
phys-paris-1# symrdf -g devgroup1 swap
```

Confirm that the swap was successful by using the `symrdf list` command to view the device group information.

```
phys-paris-1# symdg list
D E V I C E      G R O U P S
```

Number of Name	Type	Valid	Symmetrix ID	Devs	GKs	BCVs	VDEVs
devgroup1	RDF1	Yes	000187401215	6	0	0	0
devgroup2	RDF1	Yes	000187401215	2	0	0	0

5. **Activate the protection group on both clusters in the partnership.**

```
phys-paris-1# geopg start -e Global srdfpg
```

This command starts the application on `cluster-paris`. Data replication starts from `cluster-paris` to `cluster-newyork`.

## ▼ How to Make the Original Secondary Cluster Primary for an SRDF Protection Group

1. **Resynchronize the configuration of the protection group.**

This command updates the configuration of the protection group on `cluster-newyork` with the configuration information of the protection group on `cluster-paris`.

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

After the `geopg update` command runs successfully, `srdfpg` has the following role on each cluster:

```
cluster-paris:
srdfpg role: Secondary
cluster-newyork:
srdfpg role: Primary
```

2. **Run the `symrdf swap` command so that the device group, `devgroup2`, has the RDF2 role.**

```
phys-paris-1# symrdf -g devgroup2 failover
```

```
phys-paris-1# symrdf -g devgroup2 swap
```

Confirm that the swap was successful by using the `symrdf list` command to view the device group information.

```
phys-paris-1# symdg list
```

```
D E V I C E      G R O U P S

Number of
Name           Type      Valid  Symmetrix ID  Devs  GKs  BCVs  VDEVs
devgroup1      RDF2     Yes    000187401215  6     0    0
devgroup2      RDF2     Yes    000187401215  2     0    0     0
```

3. **Activate the protection group on both clusters in the partnership.**

```
phys-newyork-1# geopg start -e Global srdfpg
```

This command starts the application on `cluster-newyork`. Data replication starts from `cluster-newyork` to `cluster-paris`.




---

**Caution** - This command overwrites the data on `cluster-paris`.

---

## Recovering From an SRDF Data Replication Error

When an error occurs at the data replication level, the error is reflected in the status of the resource in the replication resource group of the relevant device group. This changed status

appears in the Data Replication status field in the output of the `geoadm status` command for that protection group.

This section contains information about the following topics:

- [“How to Detect Data Replication Errors” on page 66](#)
- [“How to Recover From an SRDF Data Replication Error” on page 67](#)

## ▼ How to Detect Data Replication Errors

---

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

---

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

```
# clresource status -v sc_geo_dr-SRDF-protection-group-device-group
```

For information about how different Resource status values map to actual replication pair states, see [Table 3-1](#).

Running the `clresource status` command might return the following:

```
...
-- Resources --

Resource Name                               Node Name   State      Status Message
-----
Resource: sc_geo_dr-SRDF-srdfpg-devgroup1  pemc1      Online    Online - Partitioned
Resource: sc_geo_dr-SRDF-srdfpg-devgroup1  pemc2      Offline   Offline
...
```

2. **Display the aggregate resource status for all device groups 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 SRDF device group, `devgroup1`, is in the Suspended state on `cluster-paris`. [Table 3-1](#) indicates that the Suspended state corresponds to a resource status of `FAULTED`. So, the data replication state of the protection group is also `FAULTED`. This state is reflected in the output of the `geoadm status` command, which displays the state of the protection group as `Error`.

```
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 "srdffg" : Error
Partnership           : paris-newyork-ps
Synchronization       : OK

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

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

```

## ▼ How to Recover From an SRDF Data Replication Error

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

---

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

---

1. **Use the procedures in the SRDF documentation to determine the causes of the FAULTED state.**
2. **Recover from the faulted state by using the SRDF procedures.**  
If the recovery procedures change the state of the device group, this state is automatically detected by the resource and is reported as a new protection group state.
3. **Revalidate the protection group configuration.**

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

```
protection-group
```

Specifies the name of the SRDF protection group.

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

**4. Review the status of the protection group configuration.**

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

```
protection-group
```

Specifies the name of the SRDF protection group

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

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

## Geographic Edition Properties for SRDF

---

This appendix provides the properties of Geographic Edition data replication components for SRDF.

This appendix contains the following sections:

- “SRDF Protection Group and Data Replication Component Properties” on page 69
- “SRDF Cluster Resource Properties That Must Not Be Changed” on page 70

### SRDF Protection Group and Data Replication Component Properties

The following table describes the SRDF properties that the Geographic Edition software defines.

**TABLE A-1** SRDF Properties

Property	Description
Protection Group Property: Cluster_dgs (string array)	Lists the Oracle Solaris Cluster device groups where the data is written. The list is comma delimited. Solaris Volume Manager disk sets are specified as <i>local-cluster[disk-set]:remote-cluster[disk-set]</i> . Only applications that belong to the protection group should write to these device groups. The Oracle Solaris Cluster device groups must exist and have the same name on both the primary cluster and the secondary cluster.  <b>Tuning recommendations:</b> This property can only be tuned when the protection group is offline.  <b>Category:</b> Optional  <b>Default:</b> Empty
Data Replication Component Property: DG_or_CG (string)	Specifies the type of SRDF component: device group (DG) or SRDF consistency group (CG).  <b>Tuning recommendations:</b> This property must be set to DG.  <b>Category:</b> Required

Property	Description
	<b>Default:</b> DG
Protection Group Property: NodeList (string array)	Lists the host names of the machines that can be primary for the replication mechanism. This list is comma delimited.  <b>Tuning recommendations:</b> This property can be tuned at any time.  <b>Category:</b> Optional  <b>Default:</b> All nodes in the cluster
Data Replication Component Property: R1_SID (string)	Specifies the primary (RDF1) EMC Symmetrix ID of the EMC Symmetrix devices.  <b>Tuning recommendations:</b> You can tune this property at any time.  <b>Category:</b> Required  <b>Default:</b> None, until you add a Symmetrix Remote Data Facility data replication component.
Data Replication Component Property: R2_SID (string)	Specifies the secondary (RDF2) EMC Symmetrix ID of the EMC Symmetrix devices.  <b>Tuning recommendations:</b> You can tune this property at any time.  <b>Category:</b> Required  <b>Default:</b> None, until you add a Symmetrix Remote Data Facility data replication component.

## SRDF Cluster Resource Properties That Must Not Be Changed

The Geographic Edition software automatically creates a cluster resource using the `SUNW.GeoCtLSRDF` resource type and using the naming convention `sc_geo_dr-SRDF-protection-group-device-group`. The Geographic Edition framework internally changes some properties for this `SUNW.GeoCtLSRDF` cluster resource. Altering the configuration or state of these entities directly with Oracle Solaris Cluster commands might result in unrecoverable failure.

Therefore, you must not edit the following properties when they are configured in a `SUNW.GeoCtLSRDF` cluster resource of the name `sc_geo_dr-SRDF-protection-group-device-group`:

- `DG_or_CG` – Defines the type of SRDF component: device group (DG) or consistency group (CG).
- `R1_SID` – Defines the local data replication role.
- `R2_SID` – Defines the local data replication role.
- `SRDF_group` – Defines the SRDF device group that contains the volumes that are being replicated.

- `Replication_role` – Defines the SRDF role attribute.



# Index

---

## A

- administering
  - data replication with, 49
  - device groups, 33, 40
- application resource groups
  - administering, 31, 39
  - creating, 31
  - removing, 39
- Asynchronous mode, 9

## C

- configuring
  - protection groups
    - instructions, 23
  - SRDF software
    - on secondary cluster, 14
- creating
  - application resource group, 31
  - protection groups
    - instructions, 23
  - replication device group, 33

## D

- data recovery, 53
  - failback-switchover, 54
  - failback-takeover, 57
- deleting
  - application resource group, 39
  - replication device group, 42
- device groups
  - adding to protection group, 33
  - administering, 33, 40
  - aggregate state of, 36
  - individual state, 35

- modifying, 41
- property validations, 35
- removing, 42
- state validations, 35
- subsystem validations, 35

## F

- failback-switchover, 54
- failback-takeover, 57

## M

- modes, 9
- modifying
  - replication device group, 41

## P

- primary cluster
  - data recovery, 53
  - switchover, 49
- properties
  - SRDF, 69
- protection groups
  - adding application resource group to, 31
  - adding device group to, 33
  - configuring, 23
  - creating, 23
  - creating when application resource group online, 25
  - modifying device group from, 41
  - removing application resource group, 39
  - removing device group from, 42
  - replicating configuration of, 43

validating, 30

## R

RDF pair state

mapping to resource status, 37

RDF1 role, 54, 58

recovery *See* data recovery

from replication error, 65

from switchover failure, 61

replication

adding device group, 33

modifying device group, 41

protection group configuration, 43

recovering from errors, 65

removing device group, 42

runtime status details, 46

runtime status overview, 45

SRDF, 49

switchover failure, 61

resource groups

application, 31, 39

replication status, 47

resource status

mapping from RDF pair state, 37

runtime status

replication, 45

state and status messages, 47

## S

secondary cluster

switchover, 49

SRDF

administering data replication with, 49

application resource groups

adding to protection group, 31

administering, 31, 39

removing, 39

configuring secondary cluster, 14

data recovery

failback-switchover, 54

failback-takeover, 57

device groups

administering, 33, 40

modifying, 41

removing, 42

properties of, 69

protection groups

creating, 23

replicating configuration of, 43

validating, 30

recovering from errors, 65

recovering from switchover failure, 61

runtime status, 45

detailed, 46

overall, 45

state and status messages, 47

takeover, 51

switchover, 49

results of, 50

validations, 50

switchover failure

recovering from, 61

Synchronous mode, 9

## T

takeover, 51

failback-switchover, 54

failback-takeover, 57

results of, 52

validations, 52

## V

validating

protection groups, 30