

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

Copyright © 2004, 2011, Oracle and/or its affiliates. All rights reserved.

License Restrictions Warranty/Consequential Damages Disclaimer

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.

Warranty Disclaimer

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.

Restricted Rights Notice

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 RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Hazardous Applications Notice

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.

Trademark Notice

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 in the United States and other countries.

Third Party Content, Products, and Services Disclaimer

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, 2011, 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 RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. 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. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Contents

Preface	7
1 Replicating Data With EMC Symmetrix Remote Data Facility Software	11
Administering Data Replication in an SRDF Protection Group	12
Initial Configuration of SRDF Software	13
Setting the Path to the SRDF SYMCLI	14
Configuring Data Replication With SRDF Software on the Primary Cluster	15
Configuring Data Replication With SRDF Software on the Secondary Cluster	21
Configuring the Other Entities on the Secondary Cluster	22
2 Administering SRDF Protection Groups	29
Strategies for Creating SRDF Protection Groups	29
Creating a Protection Group While the Application Is Offline	30
Creating a Protection Group While the Application Is Online	30
Creating, Modifying, Validating, and Deleting an SRDF Protection Group	32
▼ How to Create and Configure an SRDF Protection Group	33
Requirements to Support Oracle Real Application Clusters With Data Replication Software	34
▼ How to Create a Protection Group for Oracle Real Application Clusters	35
How the Data Replication Subsystem Validates the Device Group	38
▼ How to Modify an SRDF Protection Group	39
Validating an SRDF Protection Group	39
▼ How to Delete an SRDF Protection Group	41
Administering SRDF Application Resource Groups	42
▼ How to Add an Application Resource Group to an SRDF Protection Group	42
▼ How to Delete an Application Resource Group From an SRDF Protection Group	44
Administering SRDF Data Replication Device Groups	45
▼ How to Add a Data Replication Device Group to an SRDF Protection Group	46

Validations Made by the Data Replication Subsystem	47
How the State of the SRDF Device Group Is Validated	47
▼ How to Modify an SRDF Data Replication Device Group	50
▼ How to Delete a Data Replication Device Group From an SRDF Protection Group	51
Replicating the SRDF Protection Group Configuration to a Partner Cluster	52
▼ How to Replicate the SRDF Protection Group Configuration to a Partner Cluster	52
Activating an SRDF Protection Group	53
▼ How to Activate an SRDF Protection Group	54
Deactivating an SRDF Protection Group	57
▼ How to Deactivate an SRDF Protection Group	58
Resynchronizing an SRDF Protection Group	60
▼ How to Resynchronize a Protection Group	61
Checking the Runtime Status of SRDF Data Replication	61
Displaying an SRDF Runtime Status Overview	62
Displaying a Detailed SRDF Runtime Status	63
3 Migrating Services That Use SRDF Data Replication	65
Detecting Cluster Failure on a System That Uses SRDF Data Replication	65
Detecting Primary Cluster Failure	65
Detecting Secondary Cluster Failure	66
Migrating Services That Use SRDF Data Replication With a Switchover	67
Validations That Occur Before a Switchover	67
Results of a Switchover From a Replication Perspective	68
▼ How to Switch Over an SRDF Protection Group From Primary to Secondary	68
Forcing a Takeover on a System That Uses SRDF Data Replication	69
Validations That Occur Before a Takeover	70
Results of a Takeover From a Replication Perspective	71
▼ How to Force Immediate Takeover of SRDF Services by a Secondary Cluster	71
Recovering Services to a Cluster on a System That Uses SRDF Replication	72
▼ How to Resynchronize and Revalidate the Protection Group Configuration	73
▼ How to Perform a Failback-Switchover on a System That Uses SRDF Replication	75
▼ How to Perform a Failback-Takeover on a System That Uses SRDF Replication	78
Recovering From a Switchover Failure on a System That Uses SRDF Replication	82
Switchover Failure Conditions	82
Recovering From Switchover Failure	83

- ▼ How to Make the Original Primary Cluster Primary for an SRDF Protection Group 84
 - ▼ How to Make the Original Secondary Cluster Primary for an SRDF Protection Group 85
- Recovering From an SRDF Data Replication Error 86
 - ▼ How to Detect Data Replication Errors 86
 - ▼ How to Recover From an SRDF Data Replication Error 87
- A Geographic Edition Properties for SRDF 89**
 - SRDF Properties 89
 - SRDF Properties That Must Not Be Changed 90
- Index93**

Preface

Oracle Solaris Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility provides procedures for administering EMC Symmetrix Remote Data Facility (SRDF) data replication with Oracle Solaris Cluster Geographic Edition (Geographic Edition) software. This document is intended for experienced system administrators with extensive knowledge of Oracle software and hardware. This document is not to be used as a planning or presales guide.

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

Related Books

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

Topic	Documentation
Overview	<i>Oracle Solaris Cluster Geographic Edition Overview</i> <i>Oracle Solaris Cluster Geographic Edition 3.3 5/11 Documentation Center</i>
Installation	<i>Oracle Solaris Cluster Geographic Edition Installation Guide</i>
Data Replication	<i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i> <i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy and Universal Replicator</i> <i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for Oracle Data Guard</i> <i>Oracle Solaris Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i>
System administration	<i>Oracle Solaris Cluster Geographic Edition System Administration Guide</i>

For a complete list of Geographic Edition documentation, see *Oracle Solaris Cluster Geographic Edition 3.3 5/11 Release Notes*.

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

Topic	Documentation
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Hardware installation and administration	<i>Oracle Solaris Cluster 3.3 Hardware Administration Manual</i> Individual hardware administration guides
Software installation	<i>Oracle Solaris Cluster Software Installation Guide</i>
Data service installation and administration	<i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i> Individual data service guides
Data service development	<i>Oracle Solaris Cluster Data Services Developer's Guide</i>
System administration	<i>Oracle Solaris Cluster System Administration Guide</i> <i>Oracle Solaris Cluster Quick Reference</i>
Software upgrade	<i>Oracle Solaris Cluster Upgrade Guide</i>
Error messages	<i>Oracle Solaris Cluster Error Messages Guide</i>
Command and function references	<i>Oracle Solaris Cluster Reference Manual</i> <i>Oracle Solaris Cluster Data Services Reference Manual</i> <i>Oracle Solaris Cluster Quorum Server Reference Manual</i>

Using UNIX Commands

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

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

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

Related Third-Party Web Site References

Third-party URLs are referenced in this document and provide additional, related information.

Note – Oracle is not responsible for the availability of third-party web sites mentioned in this document. Oracle does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Oracle will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation and Support

See the following web sites for additional resources:

- [Documentation](http://www.oracle.com/technetwork/indexes/documentation/index.html) (<http://www.oracle.com/technetwork/indexes/documentation/index.html>)
- [Support](http://www.oracle.com/us/support/systems/index.html) (<http://www.oracle.com/us/support/systems/index.html>)

Oracle Software Resources

Oracle Technology Network (<http://www.oracle.com/technetwork/index.html>) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the [Discussion Forums](http://forums.oracle.com) (<http://forums.oracle.com>).
- Get hands-on step-by-step tutorials with [Oracle By Example](http://www.oracle.com/technetwork/tutorials/index.html) (<http://www.oracle.com/technetwork/tutorials/index.html>).

Typographic Conventions

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

TABLE P-1 Typographic Conventions

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

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
AaBbCc123	What you type, contrasted with onscreen computer output	machine_name% su Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-2 Shell Prompts

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

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:

- “Administering Data Replication in an SRDF Protection Group” on page 12
- “Initial Configuration of SRDF Software” on page 13

For information about creating and deleting data replication device groups, see “Administering SRDF Data Replication Device Groups” on page 45. For information about obtaining a global and a detailed runtime status of replication, see “Checking the Runtime Status of SRDF Data Replication” on page 61.

Administering Data Replication in an SRDF Protection Group

This section summarizes the steps for configuring SRDF data replication in a protection group.

TABLE 1-1 Administration Tasks for SRDF Data Replication

Task	Description
Perform an initial configuration of the SRDF software.	See “Initial Configuration of SRDF Software” on page 13.
Create a protection group that is configured for SRDF data replication.	See “How to Create and Configure an SRDF Protection Group” on page 33 or “How to Create a Protection Group for Oracle Real Application Clusters” on page 35
Add a device group that is controlled by SRDF.	See “How to Add a Data Replication Device Group to an SRDF Protection Group” on page 46.
Add application resource groups to the protection group.	See “How to Add an Application Resource Group to an SRDF Protection Group” on page 42.
Replicate the protection group configuration to a secondary cluster.	See “How to Replicate the SRDF Protection Group Configuration to a Partner Cluster” on page 52.
Activate the protection group.	See “How to Activate an SRDF Protection Group” on page 54.
Verify the protection group configuration.	Perform a trial a switchover or takeover and test some simple failure scenarios before bringing your system online. See Chapter 3, “Migrating Services That Use SRDF Data Replication.” Note – You cannot perform personality swaps if you are running SRDF/Asynchronous data replication.
Check the runtime status of replication.	See “Checking the Runtime Status of SRDF Data Replication” on page 61.
Detect failure.	See “Detecting Cluster Failure on a System That Uses SRDF Data Replication” on page 65.

TABLE 1-1 Administration Tasks for SRDF Data Replication (Continued)

Task	Description
Migrate services by using a switchover.	See “ Migrating Services That Use SRDF Data Replication With a Switchover ” on page 67. Note – You cannot perform personality swaps if you are running SRDF/Asynchronous data replication.
Migrate services by using a takeover.	See “ Forcing a Takeover on a System That Uses SRDF Data Replication ” on page 69.
Recover data after forcing a takeover.	See “ Recovering Services to a Cluster on a System That Uses SRDF Replication ” on page 72.

Initial Configuration of SRDF Software

This section describes the steps 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.

- “[Setting the Path to the SRDF SYMCLI](#)” on page 14
- “[Configuring Data Replication With SRDF Software on the Primary Cluster](#)” on page 15
- “[Configuring Data Replication With SRDF Software on the Secondary Cluster](#)” on page 21
- “[Configuring the Other Entities on the Secondary Cluster](#)” on page 22

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, `rawdgroup`
- If using ZFS, configuring a `zpool`
- If using Veritas Volume Manager (VxVM):
 - Configuring the VxVM disk group, `dg1`
 - Configuring the VxVM volume, `vol1`
 - Configuring the Oracle Solaris Cluster device group for the VxVM volume
- 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 Oracle Solaris Cluster software. Contact your Oracle service representative for information about current supported Oracle Solaris Cluster configurations.

The Geographic Edition software installation process on a single-node cluster creates the `/var/cluster/rgm/physnode_affinities` file. Its existence causes positive and negative

resource group affinities to be enforced at the level of the physical node, as they are in all multi-node clusters. Without this file, a single-node cluster uses resource group affinities at the level of the zone-node. The absence of this file can cause the malfunction of clustered applications, so do not remove it unless you clearly understand the potential consequences of its removal.

TABLE 1-2 Task Map: Steps in Configuring SRDF Data Replication for Geographic Edition Systems

Task	Instructions
Setting the path to the correct version of SRDF	“Setting the Path to the SRDF SYMCLI” on page 14
Configuring the SRDF device group	“Setting Up SRDF Device Groups” on page 15
Configuring a raw-disk device group	“How to Set Up Raw-Disk Device Groups for Geographic Edition Systems” on page 16
Configuring a VxVM device group	Perform both of the following procedures: <ul style="list-style-type: none"> ▪ “How to Configure VxVM Volumes for Use With SRDF Replication” on page 18 ▪ “How to Configure the Oracle Solaris Cluster Device Group for a VxVM Disk Group” on page 18
Configuring the file system and creating the application resource group	“How to Configure a Highly Available Local File System for SRDF Replication” on page 20

Setting the Path to the SRDF SYMCLI

To ensure that the Geographic Edition infrastructure uses a current, supported version of SRDF, you must manually set the location of the correct SYMCLI on all nodes of all clusters in the partnership.

▼ How to Set the Path to the SRDF SYMCLI

Perform this procedure on each cluster node, in each partner cluster.

- **On each node of the cluster, create a symbolic link to the SYMCLI of the SRDF software to be used by the Geographic Edition infrastructure.**

```
# ln -s /opt/emc/SYMCLI/srdfversion /opt/emc/SYMCLI/scgeo_default
```

If `/opt/emc/SYMCLI/scgeo_default` is not found, Geographic Edition software uses the SYMCLI of the latest version of SRDF software that is currently installed on the node and that is supported by Geographic Edition software.

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.

- [“Setting Up SRDF Device Groups” on page 15](#)
- [“How to Set Up Raw-Disk Device Groups for Geographic Edition Systems” on page 16](#)
- [“How to Configure VxVM Volumes for Use With SRDF Replication” on page 18](#)
- [“How to Configure the Oracle Solaris Cluster Device Group for a VxVM Disk Group” on page 18](#)
- [“How to Configure a Highly Available Local File System with ZFS for SRDF Replication” on page 19](#)
- [“How to Configure a Highly Available Local File System for SRDF Replication” on page 20](#)

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
```

Creating 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 000000003264.

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
```

▼ How to Set Up Raw-Disk Device Groups for Geographic Edition Systems

Geographic Edition supports the use of raw-disk device groups in addition to various volume managers. When you initially configure Oracle Solaris Cluster, device groups are automatically configured for each raw device in the cluster. Use this procedure to reconfigure these automatically created device groups for use with Geographic Edition.

1 For the devices that you want to use, unconfigure the predefined device groups.

The following commands remove the predefined device groups for d7 and d8.

```
phys-paris-1# cldevicegroup disable dsk/d7 dsk/d8
phys-paris-1# cldevicegroup offline dsk/d7 dsk/d8
phys-paris-1# cldevicegroup delete dsk/d7 dsk/d8
```

2 Create the new raw-disk device group, including the desired devices.

Ensure that the new DID does not contain any slashes. The following command creates a global device group, rawdg, which contains d7 and d8.

```
phys-paris-1# cldevicegroup create -n phys-paris-1,phys-paris-2 \
-t rawdisk -d d7,d8 rawdg
```

Example 1–1 Configuring a Raw-Disk Device Group

This example illustrates configuring the device group on the primary cluster, configuring the same device group on the partner cluster, and adding the group to an EMC Symmetrix protection group. Geographic Edition requires that the same Oracle Solaris Cluster device group, in this example rawdg, exists on both clusters..

Remove the automatically created device groups from the primary cluster.

```
phys-paris-1# cldevicegroup disable dsk/d7 dsk/d8
phys-paris-1# cldevicegroup offline dsk/d7 dsk/d8
phys-paris-1# cldevicegroup delete dsk/d7 dsk/d8
```

Create the raw-disk device group on the primary cluster.

```
phys-paris-1# cldevicegroup create -n phys-paris-1,phys-paris-2 \
-t rawdisk -d d7,d8 rawdg
```

Remove the automatically created device groups from the partner cluster.

```
phys-newyork-1# cldevicegroup disable dsk/d5 dsk/d6
phys-newyork-1# cldevicegroup offline dsk/d5 dsk/d6
phys-newyork-1# cldevicegroup delete dsk/d5 dsk/d6
```

Create the raw-disk device group on the partner cluster.

```
phys-newyork-1# cldevicegroup create -n phys-newyork-1,phys-newyork-2 \
-t rawdisk -d d5,d6 rawdg
```

Add the raw-disk device group to the protection group rawpg.

```
phys-paris-1# geopg create -d srdf -p Nodelist=phys-paris1,phys-paris-2 \
-o Primary -p cluster_dgs=rawdg -s paris-newyork-ps rawpg
```

Next Steps Create a raw-disk device group on the partner cluster with the same name as the one you created here. See [“How to Replicate the Configuration Information From the Primary Cluster, When Using Raw-Disk Device Groups”](#) on page 24 for the instructions about this task.

After you have configured the device group on both clusters, you can use the device group name wherever one is required in Geographic Edition commands such as geopg.

▼ How to Configure VxVM Volumes for Use With SRDF Replication

SRDF data replication is supported with VxVM volumes and raw-disk device groups. If you are using VxVM, you must configure VxVM volumes on the disks you selected for your SRDF device group.

- 1 **On `cluster-paris`, create VxVM disk groups on shared disks that will be replicated to the partner cluster `cluster-newyork`.**

For example, the d1 and d2 disks are configured as part of a VxVM disk group which is called dg1, by using commands such as `vxdiskadm` and `vxvg`. These disks are the ones that will be replicated to the partner cluster.

- 2 **After configuration is complete, verify that the disk group was created by using the `vxvg list` command.**

This command should list dg1 as a disk group.

- 3 **Create the VxVM volume.**

For example, a volume that is called vol1 is created in the dg1 disk group. The appropriate VxVM commands, such as `vxassist`, are used to configure the volume.

Next Steps Perform the steps in “[How to Configure the Oracle Solaris Cluster Device Group for a VxVM Disk Group](#)” on page 18 to configure the VxVM volume as an Oracle Solaris Cluster device group.

▼ How to Configure the Oracle Solaris Cluster Device Group for a VxVM Disk Group

- 1 **Register the VxVM disk group that you configured in the previous procedure with Oracle Solaris Cluster.**

Use the Oracle Solaris Cluster commands `clsetup` or `cldevice` and `cldevicegroup`.

For more information about these commands, refer to the [clsetup\(1CL\)](#) man page or the [cldevice\(1CL\)](#) and [cldevicegroup\(1CL\)](#) man pages.

- 2 **Synchronize the VxVM configuration with Oracle Solaris Cluster software, again by using the `clsetup` or `cldevice` and `cldevicegroup` commands.**

- 3 **After configuration is complete, verify the disk group registration.**

```
phys-paris-1# cldevicegroup show devicegroupname
```

The VxVM disk group, dg1, should be displayed in the output.

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

▼ How to Configure a Highly Available Local File System with ZFS for SRDF Replication

Follow this procedure to configure a highly available local file system that uses a ZFS zpool.

Before You Begin Ensure that the Oracle Solaris Cluster application resource group has already been configured.

Observe the following requirements and restrictions for ZFS:

- Mirrored and unmirrored ZFS zpools are supported.
- ZFS zpool spares are not supported with storage-based replication in a Geographic Edition configuration. The information about the spare that is stored in the zpool results in the zpool 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.

1 Create a ZFS zpool.

```
# zpool create appdataz mirror cNtXdY cNtAdB
```

`create appdataz` Specifies the name of the zpool to create.

`mirror cNtXdY cNtAdB` Specifies the LUNs to replicate with SRDF.

2 Add an HAStoragePlus resource to the application resource group, *app-rg*.

```
# clresource create -g app-rg \
-t HAStoragePlus \
-p zpools=appdataz \
hasp4appdataz
```

`-g app-rg` Specifies the application resource group.

`-p zpools=appdataz` Specifies the zpool

`hasp4appdataz` Specifies the name of the HAStoragePlus resource to create.

Example 1–2 Configuring a Highly Available Local File System With ZFS for SRDF Replication

This example creates a locally mounted file system, with HAStoragePlus using a ZFS zpool. The file system created in this example is mounted locally every time the resource is brought online.

This example assumes that the `app-rg1` resource group already exists

1. Create the zpool `appdata1`.

```
# zpool create appdata1 mirror c6t600604800018790002353594D313137d0
c6t600604800018790002353594D313143d0
```

2. Add the HAStoragePlus resource `hasp4appdata-rs` to the application resource group `app-rg1`.

```
# clresource create -g app-rg1 \
-t HASStoragePlus \
-p zpool=appdata1 \
hasp4appdata-rs
```

▼ How to Configure a Highly Available Local File System for SRDF Replication

Note – If you want to create a highly available local file system that uses a ZFS zpool, do not perform this procedure. Instead, go to [“How to Configure a Highly Available Local File System with ZFS for SRDF Replication”](#) on page 19.

Before You Begin Before you configure the file system on `cluster-paris`, ensure that the Oracle Solaris Cluster entities you require, such as application resource groups, device groups, and volumes, have already been configured.

- 1 **Create the required file system on the `vol1` volume at the command line.**
- 2 **On each node in the cluster, create mount points for the file system you just created.**

```
# mkdir -p /mounts/sample
/mounts/sample    Your mount point.
```

- 3 **Add an entry to the `/etc/vfstab` file that contains information such as the mount location.**

Whether the file system is to be mounted locally or globally depends on various factors, such as your performance requirements, or the type of application resource group you are using.

Note – You must set the `mount at boot` field in this file to `no`. This value prevents the file system from mounting on the secondary cluster at cluster startup. Instead, the Oracle Solaris Cluster software and the Geographic Edition framework handle mounting the file system by using the `HASStoragePlus` resource when the application is brought online on the primary cluster.

- 4 **Add the `HASStoragePlus` resource to the application resource group, `apprg1`.**

Adding the resource to the application resource group ensures that the necessary file systems are mounted before the application is brought online.

For more information about the `HASStoragePlus` resource type, refer to [“Enabling Highly Available Local File Systems”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

5 Verify that the device group was registered properly.

The following command should display the device group dg1.

```
phys-paris-1# cldevicegroup show dg1
```

Example 1-3 Configuring a Highly Available Local File System

This example creates a locally mounted file system, with HAS storagePlus. The file system created in this example is mounted locally every time the resource is brought online.

This example assumes that the following already exist:

- The apprg1 resource group
- The dg1 VxVM device group
- The vol1 VxVM volume

1. Create a UNIX file system (UFS).

```
phys-paris-1# newfs dev/vx/dsk/dg1/vol1
```

2. On each node in the cluster, create mount points for the file system.

```
phys-paris-1# mkdir -p /mounts/sample
phys-paris-2# mkdir -p /mounts/sample
```

3. Create mount points on all cluster paris nodes.

```
phys-paris-1# mkdir /mounts/sample
```

4. Add the following entry to the /etc/vfstab file:

```
phys-paris-1# /dev/vs/dsk/dg1/vol1 /dev/vx/rdisk/dg1/vol1 /mounts/sample \
ufs 2 no logging
```

5. Add the HASStoragePlus resource type.

```
phys-paris-1# clresource create -g apprg1 -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=/mounts/sample -p Affinityon=TRUE \
-p GlobalDevicePaths=dg1 rs-hasp
```

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 `syndg export` command to create a text file, `devgroup1.txt`, that contains the RDF1 group definitions.

```
phys-paris-1# syndg 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 `syndg 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.

```
# syndg 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 Oracle Solaris Cluster device groups, and the highly available cluster file system. This process is slightly different depending on whether you are using VxVM or raw-disk device groups. The following procedures provide instructions:

- [“How to Replicate the VxVM Configuration Information From the Primary Cluster”](#) on page 22
- [“How to Replicate the Configuration Information From the Primary Cluster, When Using Raw-Disk Device Groups”](#) on page 24
- [“How to Replicate the Configuration Information From the Primary Cluster, When Using ZFS zpools”](#) on page 27

▼ How to Replicate the VxVM Configuration Information From the Primary Cluster

- 1 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 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 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 Enable all the volumes to be scanned.

```
phys-newyork-1# vxdctl enable
```

5 Import the VxVM disk group, dg1.

```
phys-newyork-1# vxdg -C import dg1
```

6 Verify that the VxVM disk group was successfully imported.

```
phys-newyork-1# vxdg list
```

7 Enable the VxVM volume.

```
phys-newyork-1# /usr/sbin/vxrecover -g dg1 -s -b
```

8 Verify that the VxVM volumes are recognized and enabled.

```
phys-newyork-1# vxprint
```

9 Create the VxVM disk group, dg1, in Oracle Solaris Cluster software.

```
phys-newyork-1# cldevicegroup create -n phys-newyork-1,phys-newyork-2 \
-t vxvm dg1
```

10 Add an entry to the /etc/vfstab file on phys-newyork-1.

```
/dev/vx/dsk/dg1/vol1 /dev/vx/rdisk/dg1/vol1 /mounts/sample ufs 2 no logging
```

11 Create a mount directory on newyork.

```
phys-newyork-1# mkdir -p /mounts/sample
phys-newyork-2# mkdir -p /mounts/sample
```

12 Create an application resource group, apprg1, by using the clresourcegroup command.

```
phys-newyork-1# clresourcegroup create apprg1
```

13 Create the HASStoragePlus resource in apprg1.

```
phys-newyork-1# clresource create -g apprg1 -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=/mounts/sample -p AffinityOn=TRUE \
-p GlobalDevicePaths=dg1 rs-hasp
```

This HASStoragePlus resource is required for Geographic Edition systems because the software relies on the resource to bring the device groups and file systems online when the protection group starts on the primary cluster.

14 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
```

15 Unmount the file system.

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

16 Take the Oracle Solaris Cluster device group offline.

```
phys-newyork-1# cldevicegroup offline dg1
```

17 Verify that the VxVM disk group was deported.

```
phys-newyork-1# vxdg list
```

18 Reestablish the SRDF pair.

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

Initial configuration on the secondary cluster is now complete.

▼ 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. Use the `powermt` command to write detailed information about all devices into a temporary file.

```
phys-paris-1# /etc/powermt display dev=all > /tmp/file
```

c. Open the temporary file and look for the `ctd` label that applies to the appropriate device.

```

Logical device ID=00DD
state=alive; policy=BasicFailover; priority=0; queued-IOs=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths  Interf.  Mode  State  Q-IOs  Errors
=====
3073 pci@1d/SUNW,qlc@1      c6t5006048ACCC81DD0d18s0 FA 1dA  active  alive
      0 0
3075 pci@1d/SUNW,qlc@2      c8t5006048ACCC81DEFd18s0 FA 16cB  unlic  alive
      0 0

```

In this example, you see that the logical device ID 00DD maps to the ctd label c6t5006048ACCC81DD0d18.

d. 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.

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

5 Create a raw-disk device group on the partner cluster.

Use the same device group name as you used for the one on the primary cluster.

In the following command, the newyork cluster is the partner of the paris cluster.

```
phys-newyork-1# cldevicegroup disable dsk/d5 dsk/d6
phys-newyork-1# cldevicegroup offline dsk/d5 dsk/d6
phys-newyork-1# cldevicegroup delete dsk/d5 dsk/d6
phys-newyork-1# cldevicegroup create -n phys-newyork-1,phys-newyork-2 \
-t rawdisk -d d5,d6 rawdg
```

6 Verify that the device group rawdg was created.

```
phys-newyork-1# cldevicegroup show rawdg
```

7 Add an entry to the `/etc/vfstab` file on `phys-newyork-1`.

```
/dev/global/dsk/d5s2 /dev/global/rdisk/d5s2 /mounts/sample ufs 2 no logging
```

8 Create a mount directory on newyork.

```
phys-newyork-1# mkdir -p /mounts/sample
phys-newyork-2# mkdir -p /mounts/sample
```

9 Make a file system for the new device.

```
phys-newyork-1# newfs /dev/global/rdisk/d5s2
phys-newyork-1# mount /mounts/sample
```

- 10 Create an application resource group, `apprg1`, by using the `clresourcegroup` command.**

```
phys-newyork-1# clresourcegroup create apprg1
```

- 11 Create the `HASStoragePlus` resource in `apprg1`.**

```
phys-newyork-1# clresource create -g apprg1 -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=/mounts/sample -p AffinityOn=TRUE \
-p GlobalDevicePaths=rawdg rs-hasp
```

This `HASStoragePlus` resource is required for Geographic Edition systems, because the software relies on the resource to bring the device groups and file systems online when the protection group starts on the primary cluster.

- 12 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
```

- 13 Unmount the file system.**

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

- 14 Take the Oracle Solaris Cluster device group offline.**

```
phys-newyork-1# cldevicegroup offline rawdg
```

- 15 Reestablish the SRDF pair.**

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

Initial configuration on the secondary cluster is now complete.

▼ How to Replicate the Configuration Information From the Primary Cluster, When Using ZFS zpools

- 1 From one node of either cluster, ensure that the SRDF device group that replicates the LUNs in the zpool is in a synchronized state.**

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

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

- 2 From the primary cluster, export the zpool**

- a. From one node of the primary cluster, disable the `HASStoragePlus` resource that manages the zpool.**

```
phys-paris-1# clresource disable hasp4appdataz
```

- b. On each node of the primary cluster, verify that the zpool is exported.**

```
phys-paris# zpool list
```

3 From one node of either 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 From one node of the secondary cluster, import the zpool.

```
phys-newyork-1# zpool import appdataz
```

5 From one node of the secondary cluster, add the HAStoragePlus resource for the zpool.

```
phys-newyork-1# clresource create -g app-rg \  
-t HAStoragePlus \  
-p zpools=appdataz \  
hasp4appdataz
```

6 From one node of the secondary cluster, unmanage the resource group where the HAStoragePlus resource was created.

This step exports the zpool on the secondary cluster.

```
phys-newyork-1# clresource disable -g app-rg +  
phys-newyork-1# clresourcegroup offline app-rg  
phys-newyork-1# clresourcegroup unmanage app-rg
```

7 From one node of either cluster, reestablish the SRDF pair.

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

Initial configuration of the zpool on the secondary cluster is now complete.

Administering SRDF Protection Groups

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

- “Strategies for Creating SRDF Protection Groups” on page 29
- “Creating, Modifying, Validating, and Deleting an SRDF Protection Group” on page 32
- “Administering SRDF Application Resource Groups” on page 42
- “Administering SRDF Data Replication Device Groups” on page 45
- “Replicating the SRDF Protection Group Configuration to a Partner Cluster” on page 52
- “Activating an SRDF Protection Group” on page 53
- “Deactivating an SRDF Protection Group” on page 57
- “Resynchronizing an SRDF Protection Group” on page 60
- “Checking the Runtime Status of SRDF Data Replication” on page 61

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.
This strategy allows you to create a protection group without any application outage.
- Taking the application offline before creating the protection group.

The following sections describe the steps for each strategy:

- “Creating a Protection Group While the Application Is Offline” on page 30
- “Creating a Protection Group While the Application Is Online” on page 30

Creating a Protection Group While the Application Is Offline

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

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 33.](#)
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 46.](#)
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 42.](#)
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 52.](#)
6. From either cluster, start the protection group globally.
For more information, see [“How to Activate an SRDF Protection Group” on page 54.](#)

Creating a Protection Group While the Application Is Online

To add an existing application resource group to a new protection group without taking the application offline, complete the following steps 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 33.](#)
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 46.](#)
3. Start the protection group locally.
For more information, see [“How to Activate an SRDF Protection Group” on page 54.](#)
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 42.

Complete the following steps 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 52.
- Activate the protection group locally.
For more information, see [“How to Activate an SRDF Protection Group”](#) on page 54.

EXAMPLE 2-1 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 rdfpg
```

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 appr1 srdjpg
Following resource groups were successfully inserted:
"appr1"
```

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

```
phys-paris-1# geoadm list srdjpg
```

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

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

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

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

EXAMPLE 2-1 Creating an SRDF Protection Group While the Application Remains Online
(Continued)

```

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
    
```

Creating, Modifying, Validating, and Deleting an SRDF Protection Group

This section contains the following topics:

- [“How to Create and Configure an SRDF Protection Group” on page 33](#)
- [“Requirements to Support Oracle Real Application Clusters With Data Replication Software” on page 34](#)
- [“How to Create a Protection Group for Oracle Real Application Clusters” on page 35](#)
- [“How the Data Replication Subsystem Validates the Device Group” on page 38](#)
- [“How to Modify an SRDF Protection Group” on page 39](#)
- [“Validating an SRDF Protection Group” on page 39](#)
- [“How to Delete an SRDF Protection Group” on page 41](#)

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 datareplicationtype` option when you use the `geopg` command. The `geoadm status` command shows a state for these protection groups of Degraded.

For more information, see “Creating a Protection Group That Does Not Require Data Replication” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

▼ How to Create and Configure an SRDF Protection Group

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 52.

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 “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Create a new protection group that uses SRDF replication by using the `geopg create` command.

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

```
# geopg create -s partnershipname -o localrole -d srdf [-p property [-p...]] \  
protectiongroupname
```

- | | |
|--|--|
| <code>-s <i>partnershipname</i></code> | Specifies the name of the partnership. |
| <code>-o <i>localrole</i></code> | Specifies the role of this protection group on the local cluster as either primary or secondary. |
| <code>-d <i>srdf</i></code> | Specifies that the protection group data is replicated by the SRDF software. |
| <code>-p <i>propertysetting</i></code> | 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 Oracle Solaris Cluster device groups must exist and have the same name on both the primary cluster and the secondary cluster.

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*.

protectiongroupname 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 2–2 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 srdffg
```

Requirements to Support Oracle Real Application Clusters With Data Replication Software

Geographic Edition software supports Oracle Real Application Clusters (Oracle RAC) with SRDF software. Observe the following requirements when you configure Oracle RAC:

- Each CRS OCR and Voting Disk Location must be in its own device group on each cluster and cannot be replicated.
- Static data such as CRS and database binaries are not required to be replicated. But this data must be accessible from all nodes of both clusters.

- You must create a `SUNW.ScalDeviceGroup` resource in its own resource group for the device group that holds dynamic database files. This resource group must be separate from the resource group that holds the clusterware `SUNW.ScalDeviceGroup` resource.
- To be able to leave RAC infrastructure resource groups outside of Geographic Edition control, you must run Geographic Edition binaries on both cluster partners and set the RAC protection group `External_Dependency_Allowed` property to `true`.
- Do not add the CRS OCR and Voting Disk device group to the protection group's `cluster_dgs` property.
- Do not add 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 nonClusterware `SUNW.ScalDeviceGroup` resources to this `HASStoragePlus` resource. Then add this `HASStoragePlus` resource group to the 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 Real Application Clusters With Data Replication Software](#)” on page 34.
- 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.
- *If you are using the VxVM cluster feature to manage data*, you must specify the cluster feature disk group and Oracle Solaris Cluster device groups for other data volumes in the `cluster_dgs` property.

When a cluster and the VxVM cluster feature software restart, the Oracle RAC framework automatically tries to import all cluster feature device groups that were imported already before cluster went down. Therefore, the attempt to import the device groups to the original primary fails.

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 “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration 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 partnershipname -o localrole -d srdf \
-p External_Dependency_Allowed=true [-p property [-p...]] protectiongroupname
```

- s *partnershipname* Specifies the name of the partnership.
- o *localrole* 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 *propertysetting* 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 RAC, setting 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 VxVM cluster feature disk 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*.

- protectiongroupname* 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...]] protectiongroupname
```

-p propertysetting 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*.

protectiongroupname 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 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 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 resourcegroup protectiongroupname
```

resourcegroup 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.

protectiongroupname Specifies the name of the protection group.

Example 2–3 Creating a Protection Group for Oracle RAC

This example creates the protection group `pg1` which uses Oracle RAC and the cluster feature.

A cluster feature disk group `oracle-dg` controls the data which is replicated by the SRDF device group `DG01`. The node list of the Oracle RAC framework resource group is set to all nodes of the cluster.

1. Create the protection group on the primary cluster with the cluster feature disk group `racdbdg`.

```
# geopg create -s pts1 -o PRIMARY -d srdf -p cluster_dgs=racdbdg \  
-p external_dependency_allowed=true pg1  
Protection group "pg1" successfully created.
```

2. Add the SRDF device group `DG01` to protection group `pg1`.

```
# geopg add-device-group DG01 pg1  
Device group "DG01" successfully added to the protection group "pg1".
```

3. Add the `rac_server_proxy_rg` resource group and the replicated device-group resource groups, `hasp4rac-rg` and `sca1dbdg-rg`, to the protection group.

```
# geopg add-resource-group rac_server_proxy-rg,hasp4rac-rg,sca1dbdg-rg pg1
```

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 Oracle Solaris Cluster 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 Oracle Solaris Cluster 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 Oracle Solaris Cluster commands could result in unrecoverable failure.

▼ How to Modify an SRDF Protection Group

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

1 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 “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Modify the configuration of the protection group.

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

```
# geopg set-prop -p property [-p...] protectiongroupname
```

-p propertysetting 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*.

protectiongroupname 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 2–4 Modifying the Configuration of a Protection Group

This example modifies the `Timeout` property of the protection group that was created in [Example 2–2](#).

```
# geopg set-prop -p Timeout=2700 srdpfg
```

Validating an SRDF Protection Group

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 Oracle Solaris Cluster device group. The data replication layer uses the `scstat -D` 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.

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.

▼ How to Validate an SRDF Protection Group

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

1 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 “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration 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.

```
# geogg validate protectiongroupname
protectiongroupname    Specifies a unique name that identifies a single protection group
```

Example 2-5 Validating the Configuration of a Protection Group

This example validates a protection group.

```
# geogg validate protectiongroupname
```


▼ How to Delete an SRDF Protection Group

Before You Begin If you want to delete the protection group everywhere, you must run the `geogg delete` command on each cluster where the protection group exists.

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

- The protection group you want to delete exists locally.
- The protection group is offline on all clusters from which you want to delete it.

Note – You must remove the application resource groups from the protection group in order to keep the application resource groups online while deleting the protection group. See [Example 2–9](#) for examples of this procedure.

1 Log in to one of the nodes on the primary cluster, `cluster-paris`.

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

2 Delete the protection group.

This command deletes the configuration of the protection group from the local cluster. The command also removes the replication resource group for each SRDF device group in the protection group. This command does not alter the pair state of the SRDF device group.

```
# geogg delete protectiongroupname
protectiongroupname    Specifies the name of the protection group
```

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

Example 2–6 Deleting a Protection Group

This example deletes a protection group from both partner clusters. The protection group is offline on both partner clusters.

In this example, `cluster-paris` is the primary cluster. For a reminder of the sample cluster configuration, see “[Example Geographic Edition Cluster Configuration](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

```
# rlogin phys-paris-1 -l root
phys-paris-1# geogg delete srdpfg
# rlogin phys-newyork-1 -l root
phys-newyork-1# geogg delete srdpfg
```

Example 2-7 Deleting an SRDF Protection Group While Keeping Application Resource Groups Online

This example keeps online two application resource groups, `apprg1` and `apprg2`, while deleting their protection group, `srdfpg` from both partner clusters. Remove the application resource groups from the protection group, then delete the protection group.

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

Administering SRDF Application Resource Groups

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.

This section contains information about the following tasks:

- [“How to Add an Application Resource Group to an SRDF Protection Group” on page 42](#)
- [“How to Delete an Application Resource Group From an SRDF Protection Group” on page 44](#)

▼ 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.
- The `Auto_start_on_new_cluster` property of the resource group is set to `False`. You can view this property by using the `clresourcegroup` command.

```
# clresourcegroup show -p auto_start_on_new_cluster apprg
```

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

When the protection group is activated, application resource groups need to be online only on the primary cluster.

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

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

- The application resource group does not have dependencies on resource groups and resources outside of this protection group unless the `External_Dependency_Allowed` protection group property is set to `TRUE`. To add several application resource groups that share dependencies while the `External_Dependency_Allowed` protection group property is set to `FALSE`, you need to add all the application resource groups that share dependencies to the protection group in a single operation. If you add the application resource groups separately, the operation fails.
- The `HASStoragePlus` 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 `Offline`.

If the resource group is `Offline` and the protection group is `Active` after the configuration of the protection group has changed, the local state of the protection group becomes `Degraded`.

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 activate 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 [“Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*](#).

2 Add an application resource group to the protection group.

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.

```
# geogg add-resource-group resourcegrouplist protectiongroup
```

resourcegrouplist Specifies the name or names of the application resource group.

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

protectiongroup 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 2-8 Adding an Application Resource Group to a Protection Group

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

```
# geopg add-resource-group apprg1,apprg2 srdfpg
```

▼ 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.

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 [“Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration 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 `geopg remove-resource-group` command.

```
# geopg remove-resource-group resourcegrouplist protectiongroup
```

resourcegrouplist Specifies the list of application resource groups.

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

protectiongroup 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 2-9 Deleting an Application Resource Group From a Protection Group

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

```
# geopg remove-resource-group apprg1,apprg2 srdpfg
```

Administering SRDF Data Replication Device Groups

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

- “How to Add a Data Replication Device Group to an SRDF Protection Group” on page 46
- “Validations Made by the Data Replication Subsystem” on page 47
- “How the State of the SRDF Device Group Is Validated” on page 47
- “How to Modify an SRDF Data Replication Device Group” on page 50
- “How to Delete a Data Replication Device Group From an SRDF Protection Group” on page 51

For details about configuring a SRDF data replication protection group, see “How to Create and Configure an SRDF Protection Group” on page 33.

▼ How to Add a Data Replication Device Group to an SRDF 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 “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration 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.

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

-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.

devicegroupname Specifies the name of the new data replication device group.

protectiongroupname 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 `geopg` command, refer to the [geopg\(1M\)](#) man page.

Example 2–10 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.

```
# geopg 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, an Oracle Solaris Cluster data replication resource is automatically created by this command. This resource monitors data replication state. The name of the resource is `sc_ggeo_dr-SRDFprotectiongroupname-devicegroupname`. This resource is placed in the corresponding Oracle Solaris Cluster resource group, which is named `sc_geo_dr-SRDFprotectiongroupname`.



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 Oracle Solaris Cluster commands could result in unrecoverable failure.

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 dname 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

TABLE 2-1 Conditions That Determine the Aggregate Device Group State (Continued)

Condition	Aggregate Device Group State
One or more of the individual device groups states are SynInProg and none of the individual device group states is Invalid, Partitioned, Suspended, Split, R1 UpdInProg, R1 Updated, or Failedover.	SynInProg
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 <i>Invalid</i> and the pair state is not <i>Incorrect Role</i> .	<i>Faulted</i>	<i>Invalid state</i>
The RDF pair state is <i>Partitioned</i> and the pair state is not <i>Incorrect Role</i> , or <i>Invalid</i> .	<i>Faulted</i>	<i>Partitioned</i>
The RDF pair state is <i>Suspended</i> and the pair state is not <i>Incorrect Role</i> , <i>Invalid</i> , or <i>Partitioned</i> .	<i>Faulted</i>	<i>Suspended</i>
The RDF pair state is <i>SyncInProg</i> and the pair state is not <i>Incorrect Role</i> , <i>Invalid</i> , <i>Partitioned</i> , or <i>Suspended</i> .	<i>Degraded</i>	<i>SyncInProg</i>
The RDF pair state is <i>R1 UpdInProg</i> and the pair state is not <i>Incorrect Role</i> , <i>Invalid</i> , <i>Partitioned</i> , <i>Suspended</i> , or <i>SyncInProg</i> .	<i>Faulted</i>	<i>R1 UpdInProg</i>
The RDF pair state is <i>Split</i> and the pair state is not <i>Incorrect Role</i> , <i>Invalid</i> , <i>Partitioned</i> , <i>Suspended</i> , <i>SyncInProg</i> , or <i>R1 UpdInProg</i> .	<i>Faulted</i>	<i>Split</i>
The RDF pair state is <i>Failed over</i> and the pair state is not <i>Incorrect Role</i> , <i>Invalid</i> , <i>Partitioned</i> , <i>Suspended</i> , <i>SyncInProg</i> , <i>R1 UpdInProg</i> , or <i>Split</i> .	<i>Faulted</i>	<i>Failed over</i>
The RDF pair state is <i>R1 Updated</i> and the pair state is not <i>Incorrect Role</i> , <i>Invalid</i> , <i>Partitioned</i> , <i>Suspended</i> , <i>SyncInProg</i> , <i>R1 UpdInProg</i> , <i>Split</i> , or <i>Failed over</i> .	<i>Faulted</i>	<i>Replicating with role change</i>

TABLE 2-2 Mapping From the RDF Pair State to the Resource Status (Continued)

Condition	Resource Status	Status Message
The RDF pair state is Synchronized.	Online	Replicating

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.

▼ How to Modify an SRDF Data Replication Device 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 “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration 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...] srdfdevicegroupname protectiongroupname
-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*.

```
srdfdevicegroupname       Specifies the name of the new data replication device group.
```

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

Example 2-11 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

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 [“How to Delete an SRDF Protection Group” on page 41](#). 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 44](#).

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 [“Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration 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.

```
# geogg remove-device-group devicegroupname protectiongroupname
devicegroupname      Specifies the name of the data replication device group
protectiongroupname 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-protectiongroupname-devicegroupname`, 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 2–12 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.

```
# geogg 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

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 Oracle Solaris Cluster 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 “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration 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 partnershipname protectiongroup
```

-s partnershipname 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.

protectiongroup 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. For information about how to replicate Oracle Solaris Cluster entities, see “Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

Example 2–13 Replicating the SRDF Protection Group to a Partner Cluster

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

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

Activating an SRDF Protection Group

When you activate a protection group, the protection group assumes the role that you assigned to it during configuration. You can activate a protection group in the following ways:

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

Activating an SRDF protection group on a cluster has the following effect on the data replication layer:

- The data replication configuration of the protection group is validated. During validation, the current local role of a protection group is compared with the configuration of the SRDF device groups.

If the SRDF device group is not in a `Failedover` state, the local role of the protection group should match the role of the SRDF device group.

If the SRDF device group is in a `Failedover` state, then the local role of the protection group becomes secondary while the role of the SRDF device group remains primary.

- Data replication is started on the data replication device groups that are configured for the protection group, no matter whether the activation occurs on a primary or secondary cluster. Data is always replicated from the cluster on which the local role of the protection group is primary to the cluster on which the local role of the protection group is secondary.

Application handling proceeds only after data replication has been started successfully.

Activating a protection group has the following effect on the application layer:

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

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

- When a protection group is activated on the secondary cluster, the application resource groups are *not* started. The resource groups are put into the unmanaged state.

▼ How to Activate an SRDF Protection Group

When you activate a protection group using the command in this procedure, application resource groups in the protection group are also brought online. For details about how the `-e` (scope) option affects resource groups in the protection group, see the [geopg\(1M\)](#) man page.

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 “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Activate the protection group.

```
# geopg start -e scope [-n] protectiongroupname
```

-e scope Specifies the scope of the command.

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

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

-n Prevents the start of data replication at protection group startup.

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

protectiongroupname Specifies the name of the protection group.

The `geopg start` command uses the `scswitch -Z -g resourcegroupname` command to bring resource groups and resources online. For more information about using this command, see the [scswitch\(1M\)](#) man page.

Example 2–14 How the Geographic Edition Software Issues the Command to Start Replication

In this example, the Geographic Edition software starts data replication of an SRDF device group.

First, the SRDF protection group is created.

```
phys-paris-1# geopg create -s paris-newyork-ps -o primary -d srdf srdfpg
```

The device group `devgroup1` is added to the protection group.

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

The current RDF pair state of an SRDF device group, `devgroup1`, is returned in the output of the `symrdf query` command as follows:

```
phys-paris-1# symrdf -g devgroup1 query
Device Group (DG) Name      : devgroup1
DG's Type                   : RDF1
DG's Symmetrix ID          : 000187401215
```

	Source (R1) View		Target (R2) View	MODES
Standard	ST A	LI N	ST A	

Logical Device	Dev	T E	R1 Inv Tracks	R2 Inv Tracks	K S	Dev	T E	R1 Inv Tracks	R2 Inv Tracks	MDA	RDF Pair STATE
DEV001	00E4	RW	0	36	NR	00E4	RW	36	0	S..	Split
DEV002	00E5	RW	0	36	NR	00E5	RW	36	0	S..	Split
DEV003	00E6	RW	0	36	NR	00E6	RW	36	0	S..	Split

The aggregate device group state is Split.

Next, the protection group, `srdfpg`, is activated by using the `geopg start` command.

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

The Geographic Edition software runs the `symrdf -g devgroup1 establish` command at the data replication level. If the command is successful, the state of `devgroup1` is returned in the output of the `symrdf query` command as follows:

```
phys-paris-1# symrdf -g devgroup1 query
Device Group (DG) Name      : devgroup1
DG's Type                   : RDF1
DG's Symmetrix ID          : 000187401215
```

Source (R1) View					Target (R2) View					MODES			
Standard	Logical Device	T E	R1 Inv Tracks	R2 Inv Tracks	LI N	Standard	Logical Device	T E	R1 Inv Tracks	R2 Inv Tracks	MDA	RDF Pair STATE	
	DEV001	00E4	RW	0	0	RW	DEV001	00E4	WD	0	0	S..	Synchronized
	DEV002	00E5	RW	0	0	RW	DEV002	00E5	WD	0	0	S..	Synchronized
	DEV003	00E6	RW	0	0	RW	DEV003	00E6	WD	0	0	S..	Synchronized

Example 2–15 Activating an SRDF Protection Group Globally

This example activates a protection group globally.

```
# geopg start -e global srdfpg
```

The protection group, `srdfpg`, is activated on both clusters where the protection group is configured. The application resource groups in `srdfpg` are brought online on the primary cluster.

Example 2–16 Activating an SRDF Protection Group Locally

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


```
# geopg start -e local srdpfg
```

If the local cluster is the primary cluster, the resource groups in `srdpfg` also brought online on that cluster.

Deactivating an SRDF Protection Group

You can deactivate a protection group on the following levels:

- Globally – Deactivates a protection group on both clusters where the protection group is configured.
- On the primary cluster only – Secondary cluster remains active.
- On the secondary cluster only – Primary cluster remains active.

Deactivating an SRDF protection group on a cluster has the following effect on the data replication layer:

- The data replication configuration of the protection group is validated. During validation, the current local role of the protection group is compared with the aggregate device group state. If validation is successful, data replication is stopped.
- Data replication is stopped on the data replication device groups that are configured for the protection group, whether the deactivation occurs on a primary or secondary cluster.

Deactivating a protection group has the following effect on the application layer:

- When a protection group is deactivated on the primary cluster, all of the application resource groups configured for the protection group are stopped and unmanaged.
- When a protection group is deactivated on the secondary cluster, the resource groups on the secondary cluster are not affected. Application resource groups that are configured for the protection group might remain active on the primary cluster, depending on the activation state of the primary cluster.

The SRDF command that is used to stop data replication depends on the RDF state of the SRDF device group.

The following table describes the SRDF command that is used to stop data replication for each of the possible combinations of factors.

TABLE 2-3 Commands Used to Stop SRDF Data Replication

Aggregate Device Group State	Valid Local Protection Group Role	SRDF Command
Split, Suspended, Partitioned, or Failover	primary or secondary	No command is run because no data is being replicated.

TABLE 2-3 Commands Used to Stop SRDF Data Replication (Continued)

Aggregate Device Group State	Valid Local Protection Group Role	SRDF Command
Synchronized or R1Updated	primary or secondary	The <code>symrdf split</code> command is run.

▼ How to Deactivate an SRDF 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 “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Deactivate the protection group.

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

```
# geogg stop -e scope [-D] protectiongroupname
```

`-e scope`

Specifies the scope of the command.

If the scope is `Local`, then the command operates on the local cluster only. If the scope is `Global`, the command operates on both clusters where the protection group is deployed.

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

`-D`

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

If you omit this option, the data replication subsystem and the protection group are both stopped.

`protectiongroupname`

Specifies the name of the protection group.

Example 2-17 How the Geographic Edition Software Issues the Command to Stop Replication

This example illustrates how the Geographic Edition software determines the SRDF command that is used to stop data replication.

The current state of the SRDF device group, `devgroup1`, is returned in the output of the `symrdf` query command as follows:

```
phys-paris-1# symrdf -g devgroup1 query
Device Group (DG) Name      : devgroup1
DG's Type                   : RDF1
DG's Symmetrix ID          : 000187401215
```

Source (R1) View					Target (R2) View				MODES		
Standard	LI	ST			LI	ST					
Logical	T	R1 Inv	R2 Inv	K	T	R1 Inv	R2 Inv		RDF Pair		
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MDA	STATE
DEV001	00E4	RW	0	0	RW	00E4	WD	0	0	S..	Synchronized
DEV002	00E5	RW	0	0	RW	00E5	WD	0	0	S..	Synchronized
DEV003	00E6	RW	0	0	RW	00E6	WD	0	0	S..	Synchronized
DEV004	00E7	RW	0	0	RW	00E7	WD	0	0	S..	Synchronized
DEV005	00E8	RW	0	0	RW	00E8	WD	0	0	S..	Synchronized
DEV006	00E9	RW	0	0	RW	00E9	WD	0	0	S..	Synchronized

A device group, devgroup1, is added to the protection group as follows:

```
phys-paris-1# geopg add-device-group -p DG_or_CG=DG devgroup1 srdfpg
```

Next, the protection group, srdfpg, is deactivated by using the geopg stop command.

```
phys-paris-1# geopg stop -s local srdfpg
```

The Geographic Edition software runs the symrdf -g devgroup1 split command at the data replication level.

If the command is successful, the state of devgroup1 is returned in the output of the symrdf query command as follows:

```
phys-paris-1# symrdf -g devgroup1 query
Device Group (DG) Name      : devgroup1
DG's Type                   : RDF1
DG's Symmetrix ID          : 000187401215
```

Source (R1) View					Target (R2) View				MODES		
Standard	LI	ST			LI	ST					
Logical	T	R1 Inv	R2 Inv	K	T	R1 Inv	R2 Inv		RDF Pair		
Device	Dev	E	Tracks	Tracks	S	Dev	E	Tracks	Tracks	MDA	STATE
DEV001	00E4	RW	0	0	NR	00E4	RW	0	0	S..	Split
DEV002	00E5	RW	0	0	NR	00E5	RW	0	0	S..	Split
DEV003	00E6	RW	0	0	NR	00E6	RW	0	0	S..	Split
DEV004	00E7	RW	0	0	NR	00E7	RW	0	0	S..	Split
DEV005	00E8	RW	0	0	NR	00E8	RW	0	0	S..	Split
DEV006	00E9	RW	0	0	NR	00E9	RW	0	0	S..	Split

Example 2–18 Deactivating a Protection Group on All Clusters

This example deactivates a protection group on all clusters.

```
# geopg stop -e global srdpfg
```

Example 2–19 Deactivating a Protection Group on a Local Cluster

This example deactivates a protection group on the local cluster.

```
# geopg stop -e local srdpfg
```

Example 2–20 Stopping Data Replication While Leaving the Protection Group Online

This example stops only data replication on both partner clusters.

```
# geopg stop -e local -D srdpfg
```

If the administrator decides later to deactivate both the protection group and its underlying data replication subsystem, the administrator can rerun the command without the `-D` option:

```
# geopg stop -e local srdpfg
```

Example 2–21 Deactivating an SRDF Protection Group While Keeping Application Resource Groups Online

This example keeps two application resource groups, `apprg1` and `apprg2`, online while deactivating their protection group, `srdpfg`, on both clusters.

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

```
# geopg remove-resource-group apprg1,apprg2 srdpfg
```

2. Deactivate the protection group.

```
# geopg stop -e global srdpfg
```

Resynchronizing an SRDF Protection Group

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

For example, you might need to resynchronize protection groups after booting the cluster. For more information, see [“Booting a Cluster”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Resynchronizing a protection group updates only entities that are related to Geographic Edition software. For information about how to update Oracle Solaris Cluster entities, see [“Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

▼ How to Resynchronize a Protection Group

Before You Begin The protection group must be deactivated on the cluster where you are running the `geopg update` command. For information about deactivating a protection group, see [“Deactivating an SRDF Protection Group”](#) on page 57.

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 [“Geographic Edition Software and RBAC”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Resynchronize the protection group.

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

```
# geopg update protectiongroupname
protectiongroupname    Specifies the name of the protection group
```

Example 2–22 Resynchronizing a Protection Group

This example resynchronizes a protection group.

```
# geopg update srdpfg
```

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 62
- [“Displaying a Detailed SRDF Runtime Status”](#) on page 63

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

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

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

2 Check the runtime status of replication.

```
# geoadm status
```

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

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

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

```
# clresourcegroup status sc_geo_dr-SRDF-protectiongroupname
# clresource status sc_geo_dr-SRDF-protectiongroupname-srdf dname
```

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 2–4](#).

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-protectiongroupname
```

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:

```
# sc_geo_dr-SRDFprotectiongroupname-srdfdevicegroupname
```

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 2-4 Status and Status Messages of an Online SRDF Replication Resource Group

Status	Status Message
Online	Replicating
Degraded	Suspended
Degraded	SyncInProg
Faulted	Incorrect role
Faulted	Invalid state
Faulted	Partitioned
Faulted	R1 UpdInProg
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:

- “Detecting Cluster Failure on a System That Uses SRDF Data Replication” on page 65
- “Migrating Services That Use SRDF Data Replication With a Switchover” on page 67
- “Forcing a Takeover on a System That Uses SRDF Data Replication” on page 69
- “Recovering Services to a Cluster on a System That Uses SRDF Replication” on page 72
- “Recovering From a Switchover Failure on a System That Uses SRDF Replication” on page 82
- “Recovering From an SRDF Data Replication Error” on page 86

Detecting Cluster Failure on a System That Uses SRDF Data Replication

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

- “Detecting Primary Cluster Failure” on page 65
- “Detecting Secondary Cluster Failure” on page 66

Detecting Primary Cluster Failure

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

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

- Heartbeat failure is detected by a partner cluster.

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

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

For more information about logging, see “[Viewing the Geographic Edition Log Messages](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

- If the partnership is configured for heartbeat-loss notification, then one or both of the following actions occurs:
 - An email is sent to the address that is configured by the `Notification_emailaddr` property.
 - The script defined in `Notification_actioncmd` is executed.

For more information about configuring heartbeat-loss notification, see “[Configuring Heartbeat-Loss Notification](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

Detecting Secondary Cluster Failure

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

During failure detection, the following actions take place:

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

Migrating Services That Use 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 `geopg 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 67
- “[Results of a Switchover From a Replication Perspective](#)” on page 68
- “[How to Switch Over an SRDF Protection Group From Primary to Secondary](#)” on page 68

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 3-1 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`.

▼ How to Switch Over an SRDF Protection Group From Primary to Secondary

Before You Begin For a successful switchover, data replication must be active between the primary and the secondary clusters and data volumes on the two clusters must be synchronized.

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

- The Geographic Edition software is up and running on the both clusters.
- The secondary cluster is a member of a partnership.
- Both cluster partners can be reached.
- The protection group is in the OK state.



Caution – If you have configured the `Cluster_dgs` property, only applications that belong to the protection group can write to the device groups specified in the `Cluster_dgs` property.

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 “[Geographic Edition Software and RBAC](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Initiate the switchover.

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

```
# geogg switchover [-f] -m newprimarycluster protectiongroupname
```

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

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

protectiongroupname Specifies the name of the protection group

Example 3–1 Forcing a Switchover From Primary to Secondary

This example performs a switchover to the secondary cluster.

```
# geogg switchover -f -m cluster-newyork srdpg
```

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 70](#)
- [“Results of a Takeover From a Replication Perspective” on page 71](#)
- [“How to Force Immediate Takeover of SRDF Services by a Secondary Cluster” on page 71](#)

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 3-2 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	<code>symrdf \$option \$dg write_disable r2</code> <code>symrdf -g dg suspend</code> <code>symrdf \$option \$dg rw_enable r1</code>
FailedOver	Secondary	No command is run.

TABLE 3-2 SRDF Takeover Validations on the New Primary Cluster (Continued)

Aggregate RDF Pair State	Protection Group Local Role	SRDF Takeover Commands That Are Run on cluster-newyork
Synchronized, Suspended, R1 Updated, Partitioned	All	symrdf -g dg failover

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 `geogg 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 `geogg start` command, you must use the `-n` option to prevent replication from resuming.

▼ How to Force Immediate Takeover of SRDF Services by a Secondary Cluster

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

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

1 Log in to a node in the secondary cluster.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see “Geographic Edition Software and RBAC” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

2 Initiate the takeover.

```
# geopg takeover [-f] protectiongroupname
```

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

protectiongroupname Specifies the name of the protection group

Example 3-2 Forcing a Takeover by a Secondary Cluster

This example forces the takeover of `srdfpg` by the secondary cluster `cluster-newyork`.

The `phys-newyork-1` cluster is the first node of the secondary cluster. For a reminder of which node is `phys-newyork-1`, see “[Example Geographic Edition Cluster Configuration](#)” in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

```
phys-newyork-1# geopg takeover -f srdfpg
```

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

Recovering Services to a Cluster on a System That Uses 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 Resynchronize and Revalidate the Protection Group Configuration” on page 73
- “How to Perform a Failback-Switchover on a System That Uses SRDF Replication” on page 75
- “How to Perform a Failback-Takeover on a System That Uses SRDF Replication” on page 78

▼ How to Resynchronize and Revalidate the Protection Group Configuration

Use this procedure to resynchronize and revalidate data on the original primary cluster, `cluster-paris`, with the data on the current primary cluster, `cluster-newyork`.

Before You Begin Before you resynchronize and revalidate the protection group configuration, a takeover has occurred on `cluster-newyork`. The clusters now 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 the protection group could be reached during the takeover.

1 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 partnershipname
partnershipname    Specifies the name of the partnership
```

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

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

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

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

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

protectiongroupname Specifies the name of the protection group

For more information about synchronizing protection groups, see [“Resynchronizing an SRDF Protection Group” on page 60.](#)

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

```
phys-paris-1# geopg validate protectiongroupname
```

protectiongroupname Specifies a unique name that identifies a single protection group

For more information, see [“How to Validate an SRDF Protection Group” on page 40.](#)

3 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 -n -e local protectiongroupname
```

`-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.

protectiongroupname Specifies the name of the protection group.

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 an SRDF Protection Group” on page 54.](#)

4 Confirm that the protection group configuration is OK.

First, confirm that 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 SRDF pair state.

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

Refer to the Protection Group section of the output.

Next, confirm that all resources in the replication resource group, *protectiongroupname-rep-rg*, report a status of OK.

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

▼ 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 devicegroup query
```

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

```
phys-paris-1# symrdf -g devicegroup 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 partnershipname
partnershipname    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 protectiongroupname
```

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 protectiongroupname
protectiongroupname    Specifies the name of the protection group
```

For more information about synchronizing protection groups, see [“Resynchronizing an SRDF Protection Group” on page 60](#).

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 protectiongroupname
protectiongroupname    Specifies a unique name that identifies a single protection group
```

For more information, see [“How to Validate an SRDF Protection Group” on page 40](#).

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 protectiongroupname
```

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

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

protectiongroupname 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 an SRDF Protection Group” on page 54](#).

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.

```
# geopg switchover [-f] -m cluster-paris protectiongroupname
```

For more information, see [“How to Switch Over an SRDF Protection Group From Primary to Secondary” on page 68](#).

`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 protectiongroupname
```

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 61.

▼ 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 `geopg 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 devicegroup query
```

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

```
phys-paris-1# symrdf -g devicegroup 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 partnershipname
partnershipname    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 protectiongroupname
```

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 protectiongroupname
```

protectiongroupname Specifies the name of the protection group

For more information about resynchronizing protection groups, see [“How to Resynchronize a Protection Group”](#) on page 61.

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 protectiongroupname
```

protectiongroupname Specifies a unique name that identifies a single protection group

For more information, see [“How to Validate an SRDF Protection Group”](#) on page 40.

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 protectiongroupname
```

`-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.

protectiongroupname Specifies the name of the protection group

For more information, see [“How to Activate an SRDF Protection Group”](#) on page 54.

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] protectiongroupname
```

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

protectiongroupname Specifies the name of the protection group

For more information about the `geopg takeover` command, see [“How to Force Immediate Takeover of SRDF Services by a Secondary Cluster”](#) on page 71.

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] protectiongroupname
```

`-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.

protectiongroupname Specifies the name of the protection group.

For more information about the `geopg start` command, see [“How to Activate an SRDF Protection Group”](#) on page 54.

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 protectiongroupname
```

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 61.

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

Basic Geographic Edition operations such as `geogg 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 RDE than dynamic RDE. Therefore, you might need to increase the value of the `timeout` property of the protection group when using static RDE.

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 82
- “[Recovering From Switchover Failure](#)” on page 83
- “[How to Make the Original Primary Cluster Primary for an SRDF Protection Group](#)” on page 84
- “[How to Make the Original Secondary Cluster Primary for an SRDF Protection Group](#)” on page 85

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# geogg 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:
    srdfpg role: Secondary
cluster-newyork:
    srdfpg role: Secondary
```

```
phys-newyork-1# symdg list
```

Name	Type	Valid	Symmetrix ID	Devs	Number of		
					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# geogg 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# geogg update srdpfg
```

After the `geogg update` command run successfully, `srdpfg` has the following role on each cluster:

```
cluster-paris:
  srdpfg role: Primary
cluster-newyork:
  srdpfg 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								
Name	Type	Valid	Symmetrix ID	Devs	Number of			
					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# geogg start -e Global srdpfg
```

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# geogg update srdpfg
```

After the `geogg update` command runs successfully, `srdpfg` has the following role on each cluster:

```
cluster-paris:
  srdpfg role: Secondary
cluster-newyork:
  srdpfg 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			
Name	Type	Valid	Symmetrix ID	Devs	Number of		
					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# geogg start -e Global srdpfg
```

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 86
- “How to Recover From an SRDF Data Replication Error” on page 87

▼ How to Detect Data Replication Errors

- 1 Check the status of the replication resources by using the `scstat -g` command.

```
# clresource status -v sc_geo_dr-SRDF-protectiongroupname-srdf dgroupname
```

For information about how different Resource status values map to actual replication pair states, see [Table 2–4](#).

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 2–4](#) 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 "srdfpg" : 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.

- 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 protectiongroupname
```

protectiongroupname 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 protectiongroupname
```

protectiongroupname 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 device groups.

This appendix contains the following sections:

- “SRDF Properties” on page 89
- “SRDF Properties That Must Not Be Changed” on page 90

SRDF Properties

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

TABLE A-1 SRDF Properties

Property	Description
Data Replication Property: Cluster_dgs (string array)	<p>Lists the Oracle Solaris Cluster device groups where the data is written. The list is comma delimited. 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.</p> <p>Tuning recommendations: This property can only be tuned when the protection group is offline.</p> <p>Category: Optional</p> <p>Default: Empty</p>

TABLE A-1 SRDF Properties (Continued)

Property	Description
Data Replication Property: NodeList (string array)	Lists the host names of the machines that can be primary for the replication mechanism. This list is comma delimited. Tuning recommendations: This property can be tuned at any time. Category: Optional Default: All nodes in the cluster
Device Group Property: DG_or_CG (string)	Specifies if the device group is an SRDF device group or a SRDF consistency group. Tuning recommendations: This property must be set to DG. Category: Required Default: DG
Device Group Property: R1_SID (string)	Specifies the primary (RDF1) EMC Symmetrix ID of the EMC Symmetrix devices. Tuning recommendations: You can tune this property at any time. Category: Required Default: None, until you add a Symmetrix Remote Data Facility device group.
Device Group Property: R2_SID (string)	Specifies the secondary (RDF2) EMC Symmetrix ID of the EMC Symmetrix devices. Tuning recommendations: You can tune this property at any time. Category: Required Default: None, until you add a Symmetrix Remote Data Facility device group.

SRDF Properties That Must Not Be Changed

The Geographic Edition software internally changes some properties for the SUNWscgprepdrf resource. Therefore, you must not edit these properties manually.

For SRDF, do not edit the following properties:

- DG_or_CG – Defines the SRDF device group that contains the volumes that are being replicated.
- R1_SID – Defines the local data replication role.
- R2_SID – Defines the local data replication role.
- SRDF_group

- `Replication_role`

Index

A

- activating protection groups, 53–57
- administering
 - data replication with, 65–88
 - device groups, 45–51
- application resource groups
 - administering, 42–45
 - creating, 42–44
 - removing, 44–45
- Asynchronous mode, 11

C

- configuring
 - local file system, 20–21
 - protection groups
 - instructions, 33–34
 - SRDF software
 - on secondary cluster, 21–22
 - SYMCLI location, 14
 - volume
 - on primary cluster, 18
 - ZFS zpools, 19–20
- creating
 - application resource group, 42–44
 - protection groups
 - instructions, 33–34
 - replication device group, 46–47

D

- data recovery, 72–82
 - failback-switchover, 75–78
 - failback-takeover, 78–82
- deactivating protection groups, 57–60
- deleting
 - application resource group, 44–45
 - protection groups, 41–42
 - replication device group, 51
- detecting failure, 65–66
- device groups
 - adding to protection group, 46–47
 - administering, 45–51
 - aggregate state of, 48–49
 - configuring, 18
 - individual state, 48
 - modifying, 50
 - property validations, 47
 - removing, 51
 - state validations, 47–50
 - subsystem validations, 47
- device ID, with raw-disk device group, 16–17

F

- failback-switchover, 75–78
- failback-takeover, 78–82
- failure
 - detecting, 65–66
 - primary cluster, 65–66
 - secondary cluster, 66

H

- HAStoragePlus resource
 - configuring, 20–21
 - configuring with ZFS, 19–20

L

- local file-system configuration, 20–21

M

- modes, 11
- modifying
 - protection groups, 39
 - replication device group, 50

P

- primary cluster
 - data recovery, 72–82
 - failure detection, 65–66
 - switchover, 67–69
- properties, SRDF, 89–90
- protection groups
 - activating, 53–57
 - adding application resource group to, 42–44
 - adding device group to, 46–47
 - configuring, 33–34
 - creating, 33–34
 - deactivating, 57–60
 - deleting, 41–42
 - modifying, 39
 - modifying device group from, 50
 - removing application resource group, 44–45
 - removing device group from, 51
 - replicating configuration of, 52–53
 - resynchronizing, 60–61
 - validating, 40

R

- raw-disk device groups, 16–17
- RDF pair state, mapping to resource status, 49–50
- RDF1 role, 75, 79
- recovery
 - See* data recovery
 - from replication error, 86–88
 - from switchover failure, 82–86
- replication
 - adding device group, 46–47
 - modifying device group, 50
 - protection group configuration, 52–53
 - recovering from errors, 86–88
 - removing device group, 51
 - runtime status details, 63
 - runtime status overview, 62
 - SRDF, 65–88
 - switchover failure, 82–86
- resource groups
 - application, 42–45
 - replication status, 63
- resource status, mapping from RDF pair state, 49–50
- resynchronizing, protection groups, 60–61
- runtime status
 - replication, 61–63
 - state and status messages, 63

S

- secondary cluster
 - failure detection, 66
 - switchover, 67–69
- SRDF
 - activating protection groups, 53–57
 - administering data replication with, 65–88
 - application resource groups
 - adding to protection group, 42–44
 - administering, 42–45
 - removing, 44–45
 - configuring secondary cluster, 21–22
 - data recovery
 - failback-switchover, 75–78
 - failback-takeover, 78–82
 - deactivating protection groups, 57–60

SRDF (Continued)

- detecting failure, 65–66
 - primary cluster, 65–66
 - secondary cluster, 66
- device groups
 - administering, 45–51
 - configuring, 18
 - modifying, 50
 - removing, 51
- local file-system configuration, 20–21
- properties of, 89–90
- protection groups
 - activating, 53–57
 - creating, 33–34
 - deactivating, 57–60
 - deleting, 41–42
 - modifying, 39
 - replicating configuration of, 52–53
 - resynchronizing, 60–61
 - validating, 40
- recovering from errors, 86–88
- recovering from switchover failure, 82–86
- runtime status, 61–63
 - detailed, 63
 - overall, 62
 - state and status messages, 63
- start commands, 53–57
- switchover, 68–69
- SYMCLI, 14
- takeover, 69–72
- volume set on primary cluster, 18
- ZFS zpool configuration, 19–20
- switchover, 67–69
 - primary to secondary, 68–69
 - results of, 68
 - validations, 67–68
- switchover failure, recovering from, 82–86
- SYMCLI, setting the location, 14
- Synchronous mode, 11

T

- takeover, 69–72
 - failback-switchover, 75–78

takeover (Continued)

- failback-takeover, 78–82
- forcing, 71–72
- results of, 71
- validations, 70–71

V

- validating, protection groups, 40
- volume set, configuring, 18
- VxVM, 18

Z

- ZFS zpool configuration, 19–20

