

**Oracle® Solaris Cluster Geographic Edition  
Data Replication Guide for Hitachi TrueCopy  
and Universal Replicator**

Copyright © 2004, 2013, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

---

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

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

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

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

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

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

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

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

# Contents

---

<b>Preface</b> .....	7
<b>1 Replicating Data With Hitachi TrueCopy and Universal Replicator Software</b> .....	11
Administering Data Replication in a Hitachi TrueCopy or Universal Replicator Protection Group .....	12
Initial Configuration of Hitachi TrueCopy or Universal Replicator Software .....	13
Ensuring Data Consistency in Asynchronous Mode Replication .....	13
Overview of Initial Configuration Process .....	14
Configuration Requirements and Guidelines .....	14
Configuring Data Replication With Hitachi TrueCopy or Universal Replicator Software on the Primary Cluster .....	15
Configuring Data Replication With Hitachi TrueCopy or Universal Replicator Software on the Secondary Cluster .....	22
<b>2 Administering Hitachi TrueCopy and Universal Replicator Protection Groups</b> .....	29
Strategies for Creating Hitachi TrueCopy and Universal Replicator Protection Groups .....	29
Creating a Protection Group While the Application Is Offline .....	30
Creating a Protection Group While the Application Is Online .....	31
Creating, Modifying, Validating, and Deleting a Hitachi TrueCopy or Universal Replicator Protection Group .....	33
▼ How to Create and Configure a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters .....	34
Ensuring Data Consistency for Hitachi Universal Replicator in Asynchronous Mode .....	36
Requirements to Support Oracle Real Application Clusters With Data Replication Software .....	39
▼ How to Create a Protection Group for Oracle Real Application Clusters .....	40
How the Data Replication Subsystem Validates the Device Group .....	43
▼ How to Modify a Hitachi TrueCopy or Universal Replicator Protection Group .....	44
Validating a Hitachi TrueCopy or Universal Replicator Protection Group .....	45

▼ How to Delete a Hitachi TrueCopy or Universal Replicator Protection Group .....	46
Administering Hitachi TrueCopy and Universal Replicator Application Resource Groups ....	47
▼ How to Add an Application Resource Group to a Hitachi TrueCopy or Universal Replicator Protection Group .....	48
▼ How to Delete an Application Resource Group From a Hitachi TrueCopy or Universal Replicator Protection Group .....	49
Administering Hitachi TrueCopy and Universal Replicator Data Replication Device Groups .....	50
▼ How to Add a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator Protection Group .....	51
Validations Made by the Data Replication Subsystem .....	52
How the State of the Hitachi TrueCopy or Universal Replicator Device Group Is Validated .....	53
▼ How to Modify a Hitachi TrueCopy or Universal Replicator Data Replication Device Group .....	57
▼ How to Delete a Data Replication Device Group From a Hitachi TrueCopy or Universal Replicator Protection Group .....	57
Replicating the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster .....	59
▼ How to Replicate the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster .....	59
Activating a Hitachi TrueCopy or Universal Replicator Protection Group .....	60
▼ How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group .....	63
Deactivating a Hitachi TrueCopy or Universal Replicator Protection Group .....	65
▼ How to Deactivate a Hitachi TrueCopy or Universal Replicator Protection Group .....	66
Resynchronizing a Hitachi TrueCopy or Universal Replicator Protection Group .....	69
▼ How to Resynchronize a Protection Group .....	69
Checking the Runtime Status of Hitachi TrueCopy and Universal Replicator Data Replication .....	70
Displaying a Hitachi TrueCopy or Universal Replicator Runtime Status Overview .....	70
Displaying a Detailed Hitachi TrueCopy or Universal Replicator Runtime Status .....	71
<b>3 Migrating Services That Use Hitachi TrueCopy and Universal Replicator Data Replication ...</b>	<b>73</b>
Detecting Cluster Failure on a System That Uses Hitachi TrueCopy or Universal Replicator Data Replication .....	73
Detecting Primary Cluster Failure .....	74
Detecting Secondary Cluster Failure .....	74
Migrating Services That Use Hitachi TrueCopy or Universal Replicator Data Replication With	

a Switchover .....	75
Validations That Occur Before a Switchover .....	75
Results of a Switchover From a Replication Perspective .....	76
▼ How to Switch Over a Hitachi TrueCopy or Universal Replicator Protection Group From Primary to Secondary .....	76
Forcing a Takeover on a System That Uses Hitachi TrueCopy or Universal Replicator Data Replication .....	77
Validations That Occur Before a Takeover .....	78
Results of a Takeover From a Replication Perspective .....	79
▼ How to Force Immediate Takeover of Hitachi TrueCopy or Universal Replicator Services by a Secondary Cluster .....	80
Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy or Universal Replicator Replication .....	81
▼ How to Resynchronize and Revalidate the Protection Group Configuration .....	81
▼ How to Perform a Failback-Switchover on a System That Uses Hitachi TrueCopy or Universal Replicator Replication .....	83
▼ How to Perform a Failback-Takeover on a System That Uses Hitachi TrueCopy or Universal Replicator Replication .....	86
Recovering From a Switchover Failure on a System That Uses Hitachi TrueCopy or Universal Replicator Replication .....	91
Switchover Failure Conditions .....	91
Recovering From Switchover Failure .....	92
▼ How to Make the Original Primary Cluster Primary for a Hitachi TrueCopy or Universal Replicator Protection Group .....	93
▼ How to Make the Original Secondary Cluster Primary for a Hitachi TrueCopy or Universal Replicator Protection Group .....	93
Recovering From a Hitachi TrueCopy or Universal Replicator Data Replication Error .....	94
How to Detect Data Replication Errors .....	94
▼ How to Recover From a Hitachi TrueCopy or Universal Replicator Data Replication Error .....	95
<b>A Geographic Edition Properties for Hitachi TrueCopy and Universal Replicator .....</b>	<b>97</b>
Hitachi TrueCopy and Universal Replicator Properties .....	97
Hitachi TrueCopy and Universal Replicator Properties That Must Not Be Changed .....	99
<b>Index .....</b>	<b>101</b>



# Preface

---

The *Oracle Solaris Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy and Universal Replicator* provides procedures for administering Hitachi TrueCopy and Universal Replicator data replication with Oracle Solaris Cluster Geographic Edition (Geographic Edition) software.

---

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

---

This document is intended for experienced system administrators with extensive knowledge of Oracle software and hardware. Do not use this document as a presales guide. You should have already determined your system requirements and purchased the appropriate equipment and software before reading this document.

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

## 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

## Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Description	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
<b>AaBbCc123</b>	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<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. <b>Note:</b> Some emphasized items appear bold online.

## Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

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



## Related Books

Information about related Geographic Edition topics is available in the documentation that is listed in the following table. All Geographic Edition documentation is available at [http://www.oracle.com/technetwork/indexes/documentation/index.html#sys\\_sw](http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw).

Topic	Documentation
Overview	<i>Oracle Solaris Cluster Geographic Edition Overview</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> <i>Oracle Solaris Cluster Geographic Edition Remote Replication Guide for Sun ZFS Storage Appliance</i>
System administration	<i>Oracle Solaris Cluster Geographic Edition System Administration Guide</i>

Topic	Documentation
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Hardware installation and administration	<i>Oracle Solaris Cluster 3.3 3/13 Hardware Administration Manual</i> and 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> and 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>

Topic	Documentation
Command and function references	<i>Oracle Solaris Cluster Reference Manual</i> <i>Oracle Solaris Cluster Data Services Reference Manual</i>

---

## Access to Oracle Support

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

## Getting Help

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

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

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

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

---

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

# Replicating Data With Hitachi TrueCopy and Universal Replicator 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.

Oracle Solaris Cluster Geographic Edition (Geographic Edition) software supports the use of Hitachi TrueCopy and Universal Replicator software for data replication. Before you start replicating data with Hitachi TrueCopy or Universal Replicator software, you must be familiar with the Hitachi TrueCopy and Universal Replicator documentation, have the Hitachi TrueCopy or Universal Replicator product, and have the latest Hitachi TrueCopy or Universal Replicator patches installed on your system. For information about installing the Hitachi TrueCopy or Universal Replicator software, see the Hitachi TrueCopy and Universal Replicator product documentation.

This chapter contains the procedures for configuring and administering data replication with Hitachi TrueCopy and Universal Replicator software. The chapter contains the following sections:

- [“Administering Data Replication in a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 12
- [“Initial Configuration of Hitachi TrueCopy or Universal Replicator Software”](#) on page 13

For information about creating and deleting data replication device groups, see [“Administering Hitachi TrueCopy and Universal Replicator Data Replication Device Groups”](#) on page 50. For information about obtaining a global and a detailed runtime status of replication, see [“Checking the Runtime Status of Hitachi TrueCopy and Universal Replicator Data Replication”](#) on page 70.

# Administering Data Replication in a Hitachi TrueCopy or Universal Replicator Protection Group

This section summarizes the steps for configuring Hitachi TrueCopy and Universal Replicator data replication in a protection group.

TABLE 1-1 Administration Tasks for Hitachi TrueCopy and Universal Replicator Data Replication

Task	Description
Review configuration requirements and guidelines, and perform an initial configuration of the Hitachi TrueCopy or Universal Replicator software.	See “Initial Configuration of Hitachi TrueCopy or Universal Replicator Software” on page 13.
Create a protection group that is configured for Hitachi TrueCopy or Universal Replicator data replication.	See “How to Create and Configure a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters” on page 34 or “How to Create a Protection Group for Oracle Real Application Clusters” on page 40.
Add a device group that is controlled by Hitachi TrueCopy or Universal Replicator.	See “How to Add a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator Protection Group” on page 51.
Add an application resource group to the protection group.	See “How to Add an Application Resource Group to a Hitachi TrueCopy or Universal Replicator Protection Group” on page 48.
Replicate the protection group configuration to a secondary cluster.	See “How to Replicate the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster” on page 59.
Test the configured partnership and protection groups to validate the setup.	Perform a trial switchover or takeover and test some simple failure scenarios. See Chapter 3, “Migrating Services That Use Hitachi TrueCopy and Universal Replicator Data Replication.”
Activate the protection group.	See “How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group” on page 63.
Check the runtime status of replication.	See “Checking the Runtime Status of Hitachi TrueCopy and Universal Replicator Data Replication” on page 70.
Detect failure.	See “Detecting Cluster Failure on a System That Uses Hitachi TrueCopy or Universal Replicator Data Replication” on page 73.
Migrate services by using a switchover.	See “Migrating Services That Use Hitachi TrueCopy or Universal Replicator Data Replication With a Switchover” on page 75.

**TABLE 1-1** Administration Tasks for Hitachi TrueCopy and Universal Replicator Data Replication  
(Continued)

Task	Description
Migrate services by using a takeover.	See “ <a href="#">Forcing a Takeover on a System That Uses Hitachi TrueCopy or Universal Replicator Data Replication</a> ” on page 77.
Recover data after forcing a takeover.	See “ <a href="#">Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy or Universal Replicator Replication</a> ” on page 81.
Detect and recover from a data replication error.	See “ <a href="#">Recovering From a Hitachi TrueCopy or Universal Replicator Data Replication Error</a> ” on page 94.

## Initial Configuration of Hitachi TrueCopy or Universal Replicator Software

This section describes how to configure Hitachi TrueCopy or Universal Replicator software on the primary and secondary cluster. It also includes information about the preconditions for creating Hitachi TrueCopy and Universal Replicator protection groups. This section provides the following information:

- “[Ensuring Data Consistency in Asynchronous Mode Replication](#)” on page 13
- “[Overview of Initial Configuration Process](#)” on page 14
- “[Configuration Requirements and Guidelines](#)” on page 14
- “[Configuring Data Replication With Hitachi TrueCopy or Universal Replicator Software on the Primary Cluster](#)” on page 15
- “[Configuring Data Replication With Hitachi TrueCopy or Universal Replicator Software on the Secondary Cluster](#)” on page 22

## Ensuring Data Consistency in Asynchronous Mode Replication

Hitachi Universal Replicator can provide guaranteed data consistency in asynchronous mode replication, in which the replication fence level is set to `async`. Asynchronous mode replication is commonly used between a primary data center and a distant disaster recovery site. Guaranteed data consistency in asynchronous mode is therefore critical to the functioning of a disaster recovery system.

Guaranteed data consistency in asynchronous replication mode requires the following:

- You must run Hitachi Universal Replicator. Hitachi TrueCopy cannot always guarantee data consistency in asynchronous mode.

- On both clusters of the Geographic Edition partnership, you must have Hitachi storage arrays that are supported for use with Hitachi Universal Replicator. Talk to your Oracle representative for a list of currently supported hardware.
- You must configure journal volumes on the Hitachi storage arrays at both sites. For instructions, see the Hitachi documentation for your array.
- A journal volume must be associated with each asynchronously replicated paired device in the `/etc/horcm.conf` file. You configure this association in `/etc/horcm.conf` as a property of the parameter `HORCM_LDEV`. You cannot use the property `HORCM_DEV`. For details, see [“Configuration of the `/etc/horcm.conf` File” on page 16](#) and [“Journal Volumes” on page 16](#).
- Each asynchronously replicated Hitachi device group that is used by one particular service or application must be assigned the same consistency group ID (CTGID) as the protection group that manages it. To do so, you can complete the following steps:
  1. Create the protection group with the CTGID that you want to use.
  2. Add uninitialized Hitachi device groups to the protection group.
  3. Start the protection group.

For details, see [“Ensuring Data Consistency for Hitachi Universal Replicator in Asynchronous Mode” on page 36](#).

## Overview of Initial Configuration Process

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

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

## Configuration Requirements and Guidelines

Observe the following requirements and guidelines:

- If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Geographic Edition software does not support using a Hitachi TrueCopy or Universal Replicator S-VOL and Command Device as a Oracle Solaris Cluster quorum device. See [“Using Storage-Based Data Replication Within a Cluster” in \*Oracle Solaris Cluster System Administration Guide\*](#) for more information.
- All Hitachi TrueCopy or Universal Replicator device groups with the same consistency group ID (CTGID) must be added to the same protection group.

- Geographic Edition software uses the default CCI instance to manage the Hitachi TrueCopy or Universal Replicator devices. Geographic Edition software starts the default CCI instance whenever a TrueCopy device group is managed by Geographic Edition software. Applications that are not under the control of Geographic Edition software can also use the default CCI instance or any other instances without risk to Geographic Edition or application processes or data.
- 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 Oracle Solaris Cluster device groups that are listed in the `cluster_dgs` protection group property must exist and have the same device group name on both the primary cluster and the secondary cluster.
- The Oracle Solaris Cluster 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 might cause the malfunction of clustered applications. Do not remove the file unless you understand the potential consequences.

## Configuring Data Replication With Hitachi TrueCopy or Universal Replicator Software on the Primary Cluster

This section describes the tasks that you must perform on the primary cluster before you can configure Hitachi TrueCopy or Universal Replicator data replication in the Geographic Edition software.

- [“Configuration of the `/etc/horcm.conf` File” on page 16](#)
- [“Journal Volumes” on page 16](#)
- [“Configuring the `/etc/horcm.conf` File on the Nodes of the Primary Cluster” on page 17](#)
- [“How to Set Up Raw-Disk Device Groups for Geographic Edition Systems” on page 18](#)
- [“How to Configure a Highly Available Local File System With ZFS for Hitachi Universal Replicator Replication” on page 19](#)
- [“How to Configure a Highly Available Local File System for Hitachi TrueCopy or Universal Replicator Replication” on page 21](#)

In all examples in this document, the “primary” cluster is the cluster on which the application data service is started during routine operations. The partner cluster is “secondary.” The primary cluster is named `cluster-paris`, and the secondary cluster is named `cluster-newyork`. The `cluster-paris` cluster consists of two nodes, `phys-paris-1` and `phys-paris-2`. The `cluster-newyork` cluster also consists of two nodes, `phys-newyork-1` and

phys - newyork - 2. Two device groups are configured on each cluster. The devgroup1 device group contains the paired devices pair1 and pair2. The devgroup2 device group contains the paired devices pair3 and pair4.

## Configuration of the /etc/horcm.conf File

As used with the Geographic Edition configuration, a Hitachi TrueCopy or Universal Replicator device group is a named entity consisting of sets of paired Logical Unit Numbers (LUNs). One member of each pair of LUNs is located in local storage on the primary cluster and the other member is located in local storage on a Geographic Edition partner cluster. Data is written to one member of a pair of LUNs in local storage on the primary cluster and replicated to the other member of the pair on local storage on the secondary cluster. Each LUN in a pair is assigned the same name as the name that is assigned to the other LUN in the pair. Thus, data that is written to the LUN assigned the pair1 device name on the primary cluster is replicated to the LUN assigned the pair1 device name on the secondary cluster. Data that is written to the LUN assigned the pair2 device name on the primary cluster is replicated to the LUN assigned the pair2 device name on the secondary cluster.

On each storage-attached node of each cluster, pairs are given names and assigned to a device group in the /etc/horcm.conf file. Additionally, in this file, each device group is assigned a name that is the same on all storage-attached nodes of all clusters that are participating in a Geographic Edition partnership.

In the /etc/horcm.conf file, you configure each Hitachi TrueCopy or Universal Replicator device group as a property of either the HORCM\_DEV parameter or the HORCM\_LDEV parameter. Depending on their intended use, you might configure one device group in the /etc/horcm.conf file as a property of HORCM\_DEV and another device group as a property of HORCM\_LDEV. However, a single device group can only be configured as a property of HORCM\_DEV or of HORCM\_LDEV. For any one device group, the selected parameter, HORCM\_DEV or HORCM\_LDEV, must be consistent on all storage-attached nodes of all clusters that are participating in the Geographic Edition partnership.

Of the parameters that are configured in the /etc/horcm.conf file, only HORCM\_DEV and HORCM\_LDEV have requirements that are specific to the Geographic Edition configuration. For information about configuring other parameters in the /etc/horcm.conf file, see the documentation for Hitachi TrueCopy and Universal Replicator.

## Journal Volumes

Entries in the /etc/horcm.conf file for Hitachi Universal Replicator device groups can associate journal volumes with data LUNs. Journal volumes are specially configured LUNs on the storage system array. On both the primary and secondary arrays, local journal volumes store data that has been written to application data storage on the primary cluster, but not yet replicated to application data storage on the secondary cluster. Journal volumes thereby enable Hitachi Universal Replicator to maintain the consistency of data even if the connection between the paired clusters in a Geographic Edition partnership temporarily fails. A journal volume can



be used by more than one device group on the local cluster, but typically is assigned to just one device group. Hitachi TrueCopy does not support journaling.

If you want to implement journaling, you must configure Hitachi Universal Replicator device groups as properties of the `HORCM_LDEV` parameter because only that parameter supports the association of data LUNs with journal volumes in the Geographic Edition Hitachi Universal Replicator module. If you configure Hitachi Universal Replicator device groups by using the `HORCM_DEV` parameter, no journaling occurs, and Hitachi Universal Replicator has no greater functionality than does Hitachi TrueCopy.

## Configuring the `/etc/horcm.conf` File on the Nodes of the Primary Cluster

On each storage-attached node of the primary cluster, you configure Hitachi TrueCopy and Universal Replicator device groups as properties of the `HORCM_DEV` or `HORCM_LDEV` parameter in the `/etc/horcm.conf` file, and associate them with LUNs and, if appropriate, journal volumes. All devices that are configured in this file, including journal volumes, must be in locally attached storage. The `/etc/horcm.conf` file is read by the `HORCM` daemon when it starts, which occurs during reboot or when the Geographic Edition software is started. If you change the `/etc/horcm.conf` file on any node after the Geographic Edition software is started, and you do not anticipate rebooting, you must restart the `HORCM` daemon on that node by using the commands:

```
phys-paris-1# horcm-installation-directory/usr/bin/horcmshutdown.sh
phys-paris-1# horcm-installation-directory/usr/bin/horcmstart.sh
```

Table 1-2 shows the configuration of one journaling Hitachi Universal Replicator device group in the `/etc/horcm.conf` file as a property of the `HORCM_LDEV` parameter. Each LUN in the device group is described on a single line consisting of four space-delimited entries. The LUNs in the `devgroup1` device group are named `pair1` and `pair2`. The administrator chooses the device group and paired device names. In the third field of the file, each LUN is described by its serial number, followed by a colon, followed by the journal ID of its associated journal volume. In the logical device number (`ldev`) field, the controller unit (CU) is followed by a colon, which is followed by the logical device number. Both values are in hexadecimal format. All entries are supplied by the `raidscan` command, which is described in more detail in Hitachi's documentation. The `ldev` value that is supplied by the `raidscan` command is in decimal format, so you must convert the value to base 16 to obtain the correct format for the entry in the `ldev` field. You can only use the configuration shown in Table 1-2 with Hitachi Universal Replicator, as Hitachi TrueCopy does not support journaling.

**Note** – If you want to ensure the consistency of replicated data with Hitachi Universal Replicator on both the primary cluster and the secondary cluster, you must specify a journal volume ID in the third property configuration field of `HORCM_LDEV` for each device in a Hitachi Universal Replicator device group. Otherwise, journaling does not occur and Hitachi Universal Replicator's functionality in Geographic Edition configurations is no greater than the functionality of Hitachi TrueCopy.

**TABLE 1-2** Example `HORCM_LDEV` Section of the `/etc/horcm.conf` File on the Primary Cluster

# dev_group	dev_name	serial#:jid#	ldev
devgroup1	pair1	10136:0	00:12
devgroup1	pair2	10136:0	00:13

**Table 1-3** shows the configuration of one non-journaling Hitachi TrueCopy or Universal Replicator device group in the `/etc/horcm.conf` file as a property of the `HORCM_DEV` parameter. Each LUN in the device group is described on a single line consisting of five space-delimited entries. The table describes a device group named `devgroup2` that is composed of two LUNs in a single shared storage array that is attached to the nodes of the primary cluster. The LUNs have the device names `pair3` and `pair4` and are designated by their port, `CL1-A`, target `0`, and LU numbers, 3 and 4. The port number, target ID, and LU numbers are supplied by the `raidscan` command, which is described in more detail in Hitachi's documentation. For Hitachi TrueCopy and Universal Replicator, there is no entry in the MU number field.

**TABLE 1-3** Example `HORCM_DEV` Section of the `/etc/horcm.conf` File on the Primary Cluster

# dev_group	dev_name	port number	TargetID	LU number	MU number
devgroup2	pair3	CL1-A	0	3	-
devgroup2	pair4	CL1-A	0	4	-

## ▼ 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 containing 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**

The following commands illustrate configuring the device group on the primary cluster, configuring the same device group on the partner cluster, and adding the group to a Hitachi TrueCopy or Universal Replicator protection group.

*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 truecopy -p Nodelist=phys-paris-1,phys-paris-2 \
-o Primary -p cluster_dgs=rawdg -s paris-newyork-ps rawpg
```

**Next Steps** When configuring the partner cluster, create a raw-disk device group of 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.

Once 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 a Highly Available Local File System With ZFS for Hitachi Universal Replicator Replication

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

**Note** – Perform this procedure only if you are using Hitachi Universal Replicator. ZFS is not supported with Hitachi TrueCopy replication.

If you are not using ZFS, perform instead [“How to Configure a Highly Available Local File System for Hitachi TrueCopy or Universal Replicator Replication”](#) on page 21.

---

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

Observe the following requirements and restrictions for ZFS:

- Ensure that the zpool version on the cluster where you create the zpool is supported by the Oracle Solaris OS version of the partner cluster nodes. This is necessary so that the zpool can be imported by the partner cluster nodes, when that cluster becomes primary. You can do this by setting the zpool version to the default zpool version of the cluster that is running the earlier version of Oracle Solaris software.
- Mirrored and unmirrored ZFS 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 is not supported with Hitachi TrueCopy. Use ZFS only with Hitachi Universal Replicator.
- Ensure that Hitachi Universal Replicator is configured to preserve write ordering, even after a rolling failure.

Do not configure a storage-based replicated volume as a quorum device. The Geographic Edition software does not support Hitachi Universal Replicator S-VOL and Command Device as an Oracle Solaris Cluster quorum device. See [“Using Storage-Based Data Replication Within a Cluster”](#) in *Oracle Solaris Cluster System Administration Guide* for more information.

## 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 Hitachi Universal Replicator.

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

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

- g *app-rg*  
Specifies the application resource group.
  - p *zpool=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

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 c6t6006048000018790002353594D313137d0 c6t6006048000018790002353594D313143d0
```

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

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

## ▼ How to Configure a Highly Available Local File System for Hitachi TrueCopy or Universal Replicator 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 Hitachi Universal Replicator 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 mount points, have already been configured.

If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Geographic Edition software does not support Hitachi TrueCopy or Universal Replicator S-VOL and Command Device as a Oracle Solaris Cluster quorum device. See [“Using Storage-Based Data Replication Within a Cluster”](#) in *Oracle Solaris Cluster System Administration Guide* for more information.

- 1 **Create the required file system on the `vol1` volume at the command line.**

**2 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. Data must not be mounted on the secondary cluster or data on the primary will not be replicated to the secondary cluster. Otherwise, the data will not be replicated from the primary cluster to the secondary cluster.

---

**3 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 remounted 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*.

**Example 1–3 Configuring a Highly Available Local File System**

This example assumes that the `apprg1` resource group already exists.

**1. Create a UNIX file system (UFS).**

```
phys-paris-1# newfs dev/vx/dsk/oradg1/vol1
```

The following entry is created in the `/etc/vfstab` file:

```
# /dev/vs/dsk/oradg1/vol1 /dev/vx/rdisk/oradg1/vol1 /mounts/sample \
ufs 2 no logging
```

**2. Add the `HASStoragePlus` resource type.**

```
phys-paris-1# clresource create -g apprg1 -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=/mounts/sample -p Affinityon=TRUE \
-p GlobalDevicePaths=oradg1 rs-has
```

## Configuring Data Replication With Hitachi TrueCopy or Universal Replicator Software on the Secondary Cluster

This section describes the steps that you must complete on the secondary cluster before you can configure Hitachi TrueCopy or Universal Replicator data replication in the Geographic Edition software.

## Configuring the `/etc/horcm.conf` File on the Nodes of the Secondary Cluster

For more information about how to configure the `/etc/horcm.conf` file, see the documentation for Hitachi TrueCopy and Universal Replicator.

On each node of the secondary cluster, you must configure the `/etc/horcm.conf` file with the same Hitachi TrueCopy or Universal Replicator device group names and device names that are configured on the primary cluster, and assign them to LUNs and to journal volumes on the local shared storage array.

Table 1-4 and Table 1-5 show the entries in the `/etc/horcm.conf` file on the nodes of the secondary cluster for the device groups configured on the primary cluster in “Configuring the `/etc/horcm.conf` File on the Nodes of the Primary Cluster” on page 17. Table 1-4 shows the `HORCM_LDEV` parameter configured with two locally attached LUNs, designated by their serial numbers and logical device (`ldev`) numbers, and associated with a journal ID, as they were on the primary cluster.

**Note** – If you want to ensure the consistency of replicated data with Hitachi Universal Replicator on both the primary cluster and the secondary cluster, you must specify a journal volume ID in the third property configuration field of `HORCM_LDEV` for each device in a Hitachi Universal Replicator device group. Otherwise, journaling does not occur and Hitachi Universal Replicator’s functionality in Geographic Edition configurations is no greater than the functionality of Hitachi TrueCopy.

TABLE 1-4 Example `HORCM_LDEV` Section of the `/etc/horcm.conf` File on the Secondary Cluster

# dev_group	dev_name	serial#:jid#	ldev
devgroup1	pair1	10132:1	00:14
devgroup1	pair2	10132:1	00:15

The following table shows the `HORCM_DEV` parameter configured with two LUNs designated by their port, CL1-C, target 0, and LU numbers 22 and 23.

TABLE 1-5 Example `HORCM_DEV` Section of the `/etc/horcm.conf` File on the Secondary Cluster

# dev_group	dev_name	port number	TargetID	LU number	MU number
devgroup2	pair3	CL1-C	0	22	
devgroup2	pair4	CL1-C	0	23	

After you have configured the `/etc/horcm.conf` file on the secondary cluster, you can view the status of the pairs by using the `pairdisplay` command as follows:

```

phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 54321 1.. SMPL ---- ,-----,----- -
devgroup1 pair1(R) (CL1-C , 0, 20)12345 609..SMPL ---- ,-----,----- -
devgroup1 pair2(L) (CL1-A , 0, 2) 54321 2.. SMPL ---- ,-----,----- -
devgroup1 pair2(R) (CL1-C , 0, 21)12345 610..SMPL ---- ,-----,----- -

```

## 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.

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

**Before You Begin** Do not configure a storage-based replicated volume as a quorum device. The Geographic Edition software does not support Hitachi TrueCopy or Universal Replicator S-VOL and Command Device as a Oracle Solaris Cluster quorum device. See [“Using Storage-Based Data Replication Within a Cluster” in Oracle Solaris Cluster System Administration Guide](#) for more information.

#### 1 Start replication for the devgroup1 device group.

```
phys-paris-1# paircreate -g devgroup1 -vl -f async
```

```

phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 54321 1..P-VOL COPY ASYNC ,12345 609 -
devgroup1 pair1(R) (CL1-C , 0, 20)12345 609..S-VOL COPY ASYNC ,----- 1 -
devgroup1 pair2(L) (CL1-A , 0, 2) 54321 2..P-VOL COPY ASYNC ,12345 610 -
devgroup1 pair2(R) (CL1-C , 0, 21)12345 610..S-VOL COPY ASYNC ,----- 2 -

```

#### 2 Wait for the state of the pair to become PAIR on the secondary cluster.

```

phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL PAIR ASYNC ,-----, 1 -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321 1..P-VOL PAIR ASYNC,12345, 609 -
devgroup1 pair2(L) (CL1-C , 0, 21)12345 610..S-VOL PAIR ASYNC ,-----, 2 -
devgroup1 pair2(R) (CL1-A , 0, 2)54321 2..P-VOL PAIR ASYNC,12345, 610 -

```

#### 3 Split the pair by using the pairsplit command and confirm that the secondary volumes on cluster-newyork are writable by using the -rw option.

```

phys-newyork-1# pairsplit -g devgroup1 -rw
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL SSUS ASYNC, - - - - 1 -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321 1..P-VOL PSUS ASYNC,12345 609 W
devgroup1 pair2(L) (CL1-C , 0,21) 12345 610..S-VOL SSUS ASYNC, - - - - 2 -
devgroup1 pair2(R) (CL1-A , 0, 2) 54321 2..P-VOL PSUS ASYNC,12345 610 W

```

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

Use the same device group name that you used on the primary cluster.



You can use the same DIDs on each 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
```

**5 Verify that the device group `rawdg` was created.**

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

**6 Synchronize the volume manager information with the Oracle Solaris Cluster device group and verify the output.**

```
phys-newyork-1# cldevicegroup sync rawdg1
phys-newyork-1# cldevicegroup status
```

**7 Add an entry to the `/etc/vfstab` file on each node of the `newyork` cluster.**

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

**8 Create a mount directory on each node of the `newyork` cluster.**

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

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

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

**10 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=rawdg1 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.

**11 If necessary, confirm that the application resource group is correctly configured by bringing it online and taking it offline again.**

```
phys-newyork-1# clresourcegroup switch -emM -n phys-newyork-1 apprg1
phys-newyork-1# clresourcegroup offline apprg1
```

**12 Unmount the file system.**

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

**13 Take the Oracle Solaris Cluster device group offline.**

```
phys-newyork-1# cldevicegroup offline rawdg1
```

**14 Reestablish the Hitachi TrueCopy or Universal Replicator pair.**

```
phys-newyork-1# pairresync -g devgroup1
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL PAIR ASYNC,----- 1 -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321 1..P-VOL PAIR ASYNC,12345 609 W
devgroup1 pair2(L) (CL1-C , 0,21) 12345 610..S-VOL PAIR ASYNC,----- 2 -
devgroup1 pair2(R) (CL1-A , 0, 2) 54321 2..P-VOL PAIR ASYNC,12345 610 W
```

Initial configuration on the secondary cluster is now complete.

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

Perform this procedure only if you are using Hitachi Universal Replicator.

---

**Note** – Do not configure a replicated volume as a quorum device. Geographic Edition software does not support Hitachi TrueCopy or Universal Replicator S-VOL and Command Device as a Oracle Solaris Cluster quorum device. See [“Using Storage-Based Data Replication Within a Cluster” in Oracle Solaris Cluster System Administration Guide](#) for more information.

---

### 1 From one node of either cluster, start replication for the device group that replicates the LUNs in the zpool.

```
phys-paris-1# paircreate -g devgroup1 -vl -f async

phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 54321 1..P-VOL COPY ASYNC ,12345 609 -
devgroup1 pair1(R) (CL1-C , 0, 20)12345 609..S-VOL COPY ASYNC ,----- 1 -
devgroup1 pair2(L) (CL1-A , 0, 2) 54321 2..P-VOL COPY ASYNC ,12345 610 -
devgroup1 pair2(R) (CL1-C , 0, 21)12345 610..S-VOL COPY ASYNC ,----- 2 -
```

### 2 Wait for the state of the pair to become PAIR.

```
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL PAIR ASYNC,-----, 1 -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321 1..P-VOL PAIR ASYNC,12345, 609 -
devgroup1 pair2(L) (CL1-C , 0, 21)12345 610..S-VOL PAIR ASYNC,-----, 2 -
devgroup1 pair2(R) (CL1-A , 0, 2)54321 2..P-VOL PAIR ASYNC,12345, 610 -
```

### 3 From the primary cluster, export the zpool

#### a. From one node of the primary cluster, disable the HAStoragePlus 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
```

- 4 From one node of either cluster, split the pair by using the `pairsplit` command and confirm that the secondary volumes on `cluster-newyork` are writable by using the `-rw` option.**

```
phys-newyork-1# pairsplit -g devgroup1 -rw
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL SSUS ASYNC, ----- 1 -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321 1..P-VOL PSUS ASYNC,12345 609 W
devgroup1 pair2(L) (CL1-C , 0,21) 12345 610..S-VOL SSUS ASYNC, ----- 2 -
devgroup1 pair2(R) (CL1-A , 0, 2) 54321 2..P-VOL PSUS ASYNC,12345 610 W
```

- 5 From one node of the secondary cluster, import the zpool.**

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

- 6 From one node of the secondary cluster, add the `HASStoragePlus` resource for the zpool.**

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

- 7 From one node of the secondary cluster, unmanage the resource group where the `HASStoragePlus` 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
```

- 8 From one node of either cluster, reestablish the Hitachi Universal Replicator pair.**

```
phys-newyork-1# pairresync -g devgroup1
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL PAIR ASYNC, ----- 1 -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321 1..P-VOL PAIR ASYNC,12345 609 W
devgroup1 pair2(L) (CL1-C , 0,21) 12345 610..S-VOL PAIR ASYNC, ----- 2 -
devgroup1 pair2(R) (CL1-A , 0, 2) 54321 2..P-VOL PAIR ASYNC,12345 610 W
```

Initial configuration on the secondary cluster is now complete.



# Administering Hitachi TrueCopy and Universal Replicator Protection Groups

---

This chapter contains the procedures for configuring and administering data replication with Hitachi TrueCopy and Universal Replicator software. The chapter contains the following sections:

- “Strategies for Creating Hitachi TrueCopy and Universal Replicator Protection Groups” on page 29
- “Creating, Modifying, Validating, and Deleting a Hitachi TrueCopy or Universal Replicator Protection Group” on page 33
- “Administering Hitachi TrueCopy and Universal Replicator Application Resource Groups” on page 47
- “Administering Hitachi TrueCopy and Universal Replicator Data Replication Device Groups” on page 50
- “Replicating the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster” on page 59
- “Activating a Hitachi TrueCopy or Universal Replicator Protection Group” on page 60
- “Deactivating a Hitachi TrueCopy or Universal Replicator Protection Group” on page 65
- “Resynchronizing a Hitachi TrueCopy or Universal Replicator Protection Group” on page 69
- “Checking the Runtime Status of Hitachi TrueCopy and Universal Replicator Data Replication” on page 70

## Strategies for Creating Hitachi TrueCopy and Universal Replicator Protection Groups

Before you begin creating protection groups, consider the following strategies:

- Taking the application offline before creating the protection group.

This strategy is the most straightforward because you use a single command to create the protection group on one cluster, retrieve the information on the other cluster, and start the protection group. However, because the protection group is not brought online until the end of the process, you must take the application resource group offline to add it to the protection group.

- Creating the protection group while the application remains online.  
While this strategy allows you to create a protection group without any application outage, it requires issuing more commands.

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 31](#)

## Creating a Protection Group While the Application Is Offline

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

- Create the protection group from a cluster node.  
For more information, see [“How to Create and Configure a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters” on page 34](#) or [“How to Create a Protection Group for Oracle Real Application Clusters” on page 40](#).
- Add the data replication device group to the protection group.  
For more information, see [“How to Add a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator Protection Group” on page 51](#).
- Take the application resource group offline.
- Add the application resource group to the protection group.  
For more information, see [“How to Add an Application Resource Group to a Hitachi TrueCopy or Universal Replicator Protection Group” on page 48](#).
- On the other cluster, retrieve the protection group configuration.  
For more information, see [“How to Replicate the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster” on page 59](#).
- From either cluster, start the protection group globally.  
For more information, see [“How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group” on page 63](#).

## 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.

- Create the protection group from a cluster node.  
For more information, see [“How to Create and Configure a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters”](#) on page 34 or [“How to Create a Protection Group for Oracle Real Application Clusters”](#) on page 40.
- Add the data replication device group to the protection group.  
For more information, see [“How to Add a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 51.
- Start the protection group locally.  
For more information, see [“How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 63.
- Add the application resource group to the protection group.  
For more information, see [“How to Add an Application Resource Group to a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 48.

Complete the following steps on the other cluster.

- Retrieve the protection group configuration.  
For more information, see [“How to Replicate the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster”](#) on page 59.
- Activate the protection group locally.  
For more information, see [“How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 63.

**EXAMPLE 2-1** Creating a Hitachi TrueCopy or Universal Replicator 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 truecopy -p Nodelist=phys-paris-1,phys-paris-2 \
-o Primary -s paris-newyork-ps tcpg
Protection group "tcpg" has been successfully created
```

2. Add the device group, `tc dg`, to the protection group.

**EXAMPLE 2-1** Creating a Hitachi TrueCopy or Universal Replicator Protection Group While the Application Remains Online *(Continued)*

```
phys-paris-1# geopg add-device-group -p fence_level=async tcpg tcpg
```

3. Activate the protection group locally.

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

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

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

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

```
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 "tcpg" : Degraded
Partnership : paris-newyork-ps
Synchronization : OK

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

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

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

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

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

```
phys-newyork-1# geopg start -e local tcpg
Processing operation... this may take a while...
Protection group "tcpg" 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 a Hitachi TrueCopy or Universal Replicator 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 "tcpg"              : Degraded
Partnership                          : paris-newyork-ps
Synchronization                      : OK

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

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

```

## Creating, Modifying, Validating, and Deleting a Hitachi TrueCopy or Universal Replicator Protection Group

This section contains procedures for the following tasks:

- [“How to Create and Configure a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters” on page 34](#)
- [“Ensuring Data Consistency for Hitachi Universal Replicator in Asynchronous Mode” on page 36](#)
- [“Requirements to Support Oracle Real Application Clusters With Data Replication Software” on page 39](#)
- [“How to Create a Protection Group for Oracle Real Application Clusters” on page 40](#)
- [“How the Data Replication Subsystem Validates the Device Group” on page 43](#)
- [“How to Modify a Hitachi TrueCopy or Universal Replicator Protection Group” on page 44](#)
- [“Validating a Hitachi TrueCopy or Universal Replicator Protection Group” on page 45](#)
- [“How to Delete a Hitachi TrueCopy or Universal Replicator Protection Group” on page 46](#)

**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 a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters

Use the steps in this task to create and configure a Hitachi TrueCopy or Universal Replicator protection group. If you want to use Oracle Real Application Clusters (Oracle RAC), see “[How to Create a Protection Group for Oracle Real Application Clusters](#)” on page 40.

**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 Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster](#)” on page 59.

### 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 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 truecopy [-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 either primary or secondary.

-d *truecopy*

Specifies that the protection group data is replicated by the Hitachi TrueCopy or Universal Replicator software.

-p *propertysetting*

Specifies the properties of the protection group.

You can specify the following properties:

- **Description** – Describes the protection group.
- **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.
- **Ctgid** – Specifies the consistency group ID (CTGID) of the protection group.
- **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 a Hitachi TrueCopy or Universal Replicator Protection Group

This example creates a Hitachi TrueCopy or Universal Replicator protection group on `cluster-paris`, which is set as the primary cluster.

```
# geopg create -s paris-newyork-ps -o primary -d truecopy \  
-p NodeList=phys-paris-1,phys-paris-2 tcpg
```

**Example 2-3** Creating a Hitachi TrueCopy or Universal Replicator Protection Group for Application Resource Groups That Are Online

This example creates a Hitachi TrueCopy or Universal Replicator protection group, `tcpg`, for an application resource group, `resourcegroup1`, that is currently online on `cluster-newyork`.

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

```
# geopg create -s paris-newyork-ps -o primary -d truecopy \
-p nodeList=phys-paris-1,phys-paris-2 tcpg
```

2. Activate the protection group.

```
# geopg start -e local tcpg
```

3. Add the application resource group.

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

## Ensuring Data Consistency for Hitachi Universal Replicator in Asynchronous Mode

This section describes the protection group configuration that is required in Geographic Edition software to guarantee data consistency in asynchronous mode replication.

Asynchronous mode replication is implemented by using the `async` fence level of Hitachi Universal Replicator. The following discussion therefore applies only to the `async` fence level and to Hitachi Universal Replicator as implemented in the Geographic Edition module.

### Understanding Data Consistency in Geographic Edition

The Geographic Edition module supports Hitachi TrueCopy and Universal Replicator device groups in asynchronous mode replication. Routine operations for both Hitachi TrueCopy and Universal Replicator provide data consistency in asynchronous mode. However, in the event of a temporary loss of communications or of a “rolling disaster” where different parts of the system fail at different times, only Hitachi Universal Replicator software can prevent loss of consistency of replicated data for asynchronous mode. In addition, Hitachi Universal Replicator software can only ensure data consistency with the configuration described in this section and in [“Configuring the /etc/horcm.conf File on the Nodes of the Primary Cluster” on page 17](#) and [“Configuring the /etc/horcm.conf File on the Nodes of the Secondary Cluster” on page 23](#).

In Hitachi Universal Replicator software, the Hitachi storage arrays replicate data from primary storage to secondary storage. The application that produced the data is not involved. Even so, to guarantee data consistency, replication must preserve the application's I/O write ordering, regardless of how many disk devices the application writes.

During routine operations, Hitachi Universal Replicator software on the storage secondary array pulls data from cache on the primary storage array. If data is produced faster than it can be transferred, Hitachi Universal Replicator can commit backlogged I/O and a sequence number

for each write to a journal volume on the primary storage array. The secondary storage array pulls that data from primary storage and commits it to its own journal volumes, from where it is transferred to application storage. If communications fail and are later restored, the secondary storage array begins to resynchronize the two sites by continuing to pull backlogged data and sequence numbers from the journal volume. Sequence numbers control the order in which data blocks are committed to disk so that write ordering is maintained at the secondary site despite the interruption. As long as journal volumes have enough disk space to record all data that is generated by the application that is running on the primary cluster during the period of failure, consistency is guaranteed.

In the event of a rolling disaster, where only some of the backlogged data and sequence numbers reach the secondary storage array after failures begin, sequence numbers determine which data should be committed to data LUNs to preserve consistency.

---

**Note** – In the Geographic Edition module with Hitachi Universal Replicator, journal volumes are associated with application storage in the `/etc/horc.m.conf` file. That configuration is described in [“Journal Volumes” on page 16](#) and [“Configuring the /etc/horc.m.conf File on the Nodes of the Primary Cluster” on page 17](#). For information about how to configure journal volumes on a storage array, see the Hitachi documentation for that array.

---

## Using Consistency Group IDs to Ensure Data Consistency

Along with journal volumes, consistency group IDs (CTGIDs) ensure data consistency even if the storage for an application data service includes devices in multiple Hitachi device groups. A CTGID is an integer that is assigned to one or more Hitachi device groups. It designates those devices that must be maintained in a state of replication consistent with each other. Consistency is maintained among all devices with the same CTGID whether the devices are members of a single Hitachi device group or several Hitachi device groups. For example, if Hitachi Universal Replicator stops replication on the devices of one device group that is assigned the CTGID of 5, it stops replication on all other devices in device groups with the CTGID of 5.

To ensure data consistency, an exact correspondence must therefore exist between the device groups that are used by a single application data service and a CTGID. All device groups that are used by a single data service must have the same unique CTGID. No device group can have that CTGID unless it is used by the data service.

To ensure this correspondence, the Geographic Edition software allows the administrator to set a CTGID property on each protection group. The device groups that are added to the protection group must all have the same CTGID as the protection group. If other device groups are assigned the same CTGID as the device groups in the protection group, the Geographic Edition software generates an error. For example, if the protection group `app1-pg` has been assigned the CTGID of 5, all device groups included in `app1-pg` must have the CTGID of 5. Moreover, all CTGIDs of device groups that are included in `app1-pg` must have the CTGID of 5.

You are not required to set a CTGID on a protection group. The Hitachi storage software will automatically assign a unique CTGID to an asynchronously replicated device group when it is initialized. Thereafter, the pairs in that device group will be maintained in a state of consistency with each other. Thus, if an application data service in a protection group uses storage in just one asynchronously replicated Hitachi device group, you can let the Hitachi storage array assign the device group's CTGID. You do not have to also set the CTGID of the protection group.

Similarly, if you do not need data consistency, or if your application does not write asynchronously to your Hitachi device groups, then setting the CTGID on the protection group has little use. However, if you do not assign a CTGID to a protection group, any later configuration changes to the device group or to the protection group might lead to conflicts. Assignment of a CTGID to a protection group provides the most flexibility for later changes and the most assurance of device group consistency.

## ▼ **Configuring Consistency Group IDs for Hitachi Universal Replicator Device Groups in Asynchronous Mode**

You can assign a consistency group ID (CTGID) to a protection group by setting the property `ctgid=consistency-group-ID` as an option to the `geopg create` command. You can assign CTGID values to device groups in one of two ways:

- You can add uninitialized device groups to the protection group. They are initialized and acquire the CTGID of the protection group when the protection group is started with the `geopg start` command.
- You can initialize a device group with the CTGID that you plan to use for the protection group that will hold that device group. After you create the protection group with that CTGID, you must assign the device group to it.

The following procedure demonstrates these two methods of setting the CTGID for the devices that are used by an application data service. The procedure configures a protection group named `app1-pg` with a CTGID of 5. This protection group contains the `app1-rg` resource group and the Hitachi Universal Replicator `devgroup1` device group, which uses the `async` fence level.

- Before You Begin**
- Configure a Hitachi Universal Replicator device group with journal volumes in the `/etc/horcm.conf` file as described in [“Configuring the /etc/horcm.conf File on the Nodes of the Primary Cluster”](#) on page 17 and [“Configuring the /etc/horcm.conf File on the Nodes of the Secondary Cluster”](#) on page 23.
  - Configure the devices in each device group as raw-disk devices.
  - Configure a Oracle Solaris Cluster resource group that includes a resource of type `HASStoragePlus` in addition to any other resources that are required for its application data service. This `HASStoragePlus` resource must use the disk devices of a previously configured Hitachi Universal Replicator device group as described in [“How to Configure a Highly Available Local File System for Hitachi TrueCopy or Universal Replicator Replication”](#) on page 21.

- 1 On the primary cluster, create the Geographic Edition protection group with a specified CTGID, and add the resource group.

```
phys-paris-1# geopg create -s paris-newyork-ps -o primary -d truecopy -p ctgid=5 \
-p nodelist=phys-paris-1,phys-paris-2 app1-pg
```

```
phys-paris-1# geopg add-resource-group app1-rg app1-pg
```

- 2 Add device groups to the protection group by using one of the following methods:

- Add device groups that have been configured in the `/etc/horcm.conf` file but have not been initialized by using the `paircreate` command.

```
phys-paris-1# geopg add-device-group -p fence_level=async devgroup1 app1-pg
```

- Assign CTGIDs to device groups when they are initialized by using the Hitachi `paircreate` command, and add the device groups to the protection group that has the same value for the CTGID property.

In the following example, a device group is initialized with the CTGID of 5 and then added to the `app1-pg` protection group:

```
phys-paris-1# paircreate -g devgroup1 -vl -f async 5
```

```
phys-paris-1# geopg add-device-group -p fence_level=async devgroup1 app1-pg
```

- 3 Start the protection group.

```
phys-paris-1# geopg start -e local app1-pg
```

Uninitialized device groups, if any, are initialized and assigned the CTGID of 5.

## Requirements to Support Oracle Real Application Clusters With Data Replication Software

Geographic Edition software supports Oracle Real Application Clusters (Oracle RAC) with Hitachi TrueCopy and Universal Replicator 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 39.
- The node list of the protection group must be the same as the node list of Oracle RAC framework resource group.
- If one cluster is running Oracle RAC on a different number of nodes than another cluster, ensure that all nodes on both clusters have the same resource groups defined.

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

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see “[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 truecopy \
-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 truecopy

Specifies that the protection group data is replicated by the Hitachi TrueCopy or Universal Replicator 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, set this property to true.
- **Timeout** – Specifies the timeout period for the protection group in seconds.
- **NodeList** – Lists the host names of the machines that can be primary for the replication subsystem.
- **Ctgid** – Specifies the consistency group ID (CTGID) 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.

### 3 Add a Hitachi TrueCopy or Universal Replicator device group to the protection group.

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

-p *propertysetting*

Specifies the properties of the protection group.

You can specify the `Fence_level` properties which defines the fence level that is used by the disk device group. The fence level determines the level of consistency among the primary and secondary volumes for that disk device group. You must set this to never.



**Caution** – To avoid application failure on the primary cluster, specify a `Fence_level` of `never` or `async`. If the `Fence_level` parameter is not set to `never` or `async`, data replication might not function properly when the secondary site goes down.

If you specify a `Fence_level` of `never`, the data replication roles do not change after you perform a takeover.

Do not use programs that would prevent the `Fence_level` parameter from being set to `data` or `status` because these values might be required in special circumstances.

If you have special requirements to use a `Fence_level` of `data` or `status`, consult your Oracle representative.

---

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 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 2-4 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 `racdbdg` controls the data which is replicated by the Hitachi TrueCopy or Universal Replicator device group `VG01`. 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 Truecopy \  
-p cluster_dgs=racdbdg -p external_dependency_allowed=true pg1  
Protection group "pg1" successfully created.
```

2. Add the Hitachi TrueCopy or Universal Replicator device group `VG01` to protection group `pg1`.

```
# geopg add-device-group --property fence_level=never VG01 pg1  
Device group "VG01" 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 `scaldbdg-rg`, to the protection group.

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

## How the Data Replication Subsystem Validates the Device Group

Before creating the protection group, the data replication layer validates that the `horcmd` daemon is running.

The data replication layer validates that the `horcmd` daemon is running on at least one node that is specified in the `NodeList` property.

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.

---

A Oracle Solaris Cluster resource group is automatically created when the protection group is created.

This resource in this resource group monitors data replication. The name of the Hitachi TrueCopy or Universal Replicator data replication resource group is `rg-tc-protectiongroupname`.



---

**Caution** – These automatically created replication resource groups are for Geographic Edition internal implementation purposes only. Use caution when you modify these resource groups by using Oracle Solaris Cluster commands.

---

## ▼ How to Modify a Hitachi TrueCopy or Universal Replicator 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 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 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 2-5** 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=400 tcpg
```

## Validating a Hitachi TrueCopy or Universal Replicator Protection Group

During protection group validation, the Hitachi TrueCopy or Universal Replicator data replication subsystem validates the following:

- The `horcmd` daemon is running on at least one node that is specified in the `NodeList` property of the protection group. The data replication layer also confirms that a path to a Hitachi TrueCopy or Universal Replicator storage device exists from the node on which the `horcmd` daemon is running.
- 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.
- The properties are validated for each Hitachi TrueCopy or Universal Replicator 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 `geopg validate` command. This command checks the current state of the protection group and its entities.

If the protection group and its entities are valid, then the `Configuration` status of the protection groups is set to `OK`. If the `geopg 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 `geopg validate` command again.

### ▼ How to Validate a Hitachi TrueCopy or Universal Replicator 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 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 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 2-6 Validating the Configuration of a Protection Group

This example validates a protection group.

```
# geogg validate tcpg
```

## ▼ How to Delete a Hitachi TrueCopy or Universal Replicator 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 the local cluster.

---

**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 2-8](#) and [Example 2-10](#) for examples of this procedure.

---

### 1 Log in to a 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 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 Hitachi TrueCopy or Universal Replicator device group in the protection group. This command does not alter the pair state of the Hitachi TrueCopy or Universal Replicator device group.

```
# geogg delete protectiongroupname
```

*protectiongroupname*

Specifies the name of the protection group.

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

### Example 2-7 Deleting a Protection Group

This example deletes a protection group from both partner clusters.

`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# geopg delete tcpg
# rlogin phys-newyork-1 -l root
phys-newyork-1# geopg delete tcpg
```

### Example 2-8 Deleting a Hitachi TrueCopy or Universal Replicator Protection Group While Keeping Application Resource Groups Online

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

```
# geopg remove-resource-group apprg1,apprg2 tcpg
# geopg stop -e global tcpg
# geopg delete tcpg
```

## Administering Hitachi TrueCopy and Universal Replicator 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 application resources, installation, application configuration files, and resource groups, 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 a Hitachi TrueCopy or Universal Replicator Protection Group” on page 48
- “How to Delete an Application Resource Group From a Hitachi TrueCopy or Universal Replicator Protection Group” on page 49

## ▼ How to Add an Application Resource Group to a Hitachi TrueCopy or Universal Replicator 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, you should bring the application resources groups participating in that protection group online only on the same primary cluster. Setting the `Auto_start_on_new_cluster` property to `False` prevents the Oracle Solaris Cluster resource group manager from automatically starting the application resource groups. In this case, the start up of resource groups is reserved to 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 apprg
```

- 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 protection group can be activated or deactivated and the resource group can be either `Online` or `Unmanaged`.

If the resource group is `Unmanaged` and the protection group is `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 active 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.

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

Specifies the name 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 Configuration status is OK on the local cluster, but the add operation is unsuccessful on the partner cluster, the Configuration status is set to Error on the partner cluster.

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 2–9 Adding an Application Resource Group to a Protection Group**

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

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

## ▼ How to Delete an Application Resource Group From a Hitachi TrueCopy or Universal Replicator 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. For example, you cannot remove a resource group that belongs to the data replication management entity.

**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.

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

Specifies the name 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.

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 2-10 Deleting an Application Resource Group From a Protection Group**

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

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

## Administering Hitachi TrueCopy and Universal Replicator Data Replication Device Groups

This section provides the following information about administering Hitachi TrueCopy and Universal Replicator data replication device groups:

- [“How to Add a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator Protection Group” on page 51](#)

- “Validations Made by the Data Replication Subsystem” on page 52
- “How the State of the Hitachi TrueCopy or Universal Replicator Device Group Is Validated” on page 53
- “How to Modify a Hitachi TrueCopy or Universal Replicator Data Replication Device Group” on page 57
- “How to Delete a Data Replication Device Group From a Hitachi TrueCopy or Universal Replicator Protection Group” on page 57

For details about configuring a Hitachi TrueCopy or Universal Replicator data replication protection group, see “How to Create and Configure a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters” on page 34.

## ▼ How to Add a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator 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.

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

Specifies the properties of the data replication device group.

You can specify the `Fence_level` property which defines the fence level that is used by the device group. The fence level determines the level of consistency among the primary and secondary volumes for that device group.

You can set this property to `data`, `status`, `never`, or `async`. When you use a `Fence_level` of `never` or `async`, the application can continue to write to the primary cluster even after failure on the secondary cluster. However, when you set the `Fence_level` property to `data` or `status`, the application on the primary cluster might fail because the secondary cluster is not available for the following reasons:

- Data replication link failure
- Secondary cluster and storage is down
- Storage on the secondary cluster is down




---

**Caution** – To avoid application failure on the primary cluster, specify a `Fence_level` of `never` or `async`.

If you specify a `Fence_level` of `never`, the data replication roles do not change after you perform a takeover.

If you have special requirements to use a `Fence_level` of `data` or `status`, consult your Oracle representative.

---

The other properties you can specify depend on the type of data replication you are using. For details about these properties, see [Appendix A, “Standard Geographic Edition Properties,”](#) in *Oracle Solaris Cluster Geographic Edition System Administration Guide*.

*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 2–11 Adding a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator Protection Group

This example creates a Hitachi TrueCopy or Universal Replicator data replication device group in the `tcpg` protection group.

```
# geopg add-device-group -p Fence_level=data devgroup1 tcpg
```

## Validations Made by the Data Replication Subsystem

When the Hitachi TrueCopy or Universal Replicator device group, configured as `dev_group` in the `/etc/horcm.conf` file, is added to a protection group, the data replication layer makes the following validations.

- Validates that the `horcmd` daemon is running on at least one node in the `NodeList` property of the protection group.
- Checks that the path to the storage device exists from all the nodes that are specified in the `NodeList` property. The storage device controls the new Hitachi TrueCopy or Universal Replicator device group.

- The Hitachi TrueCopy and Universal Replicator device group properties that are specified in the `geopg add -device -group` command are validated as described in the following table.

Hitachi TrueCopy or Universal Replicator Device Group Property	Validation
<i>devicegroupname</i>	Checks that the specified Hitachi TrueCopy or Universal Replicator device group is configured on all of the cluster nodes that are specified in the <code>NodeList</code> property.
<code>Fence_level</code>	If a pair is already established for this Hitachi TrueCopy or Universal Replicator device group, the data replication layer checks that the specified <code>Fence_level</code> matches the already established fence level.  If a pair is not yet established, for example, if a pair is in the SMPL state, any <code>Fence_level</code> is accepted.

When a Hitachi TrueCopy or Universal Replicator device group is added to a protection group, a Oracle Solaris Cluster resource is automatically created by this command. This resource monitors data replication. The name of the resource is `r - tc - protectiongroupname - devicegroupname`. This resource is placed in the corresponding Oracle Solaris Cluster resource group, which is named `rg - tc - protectiongroupname`.



**Caution** – You must use caution before you modify these replication resources with Oracle Solaris Cluster commands. These resources are for internal implementation purposes only.

## How the State of the Hitachi TrueCopy or Universal Replicator Device Group Is Validated

For validation purposes, Geographic Edition gives each Hitachi TrueCopy or Universal Replicator device group a state according to the current state of its pair. This state is returned by the `pairvolchk -g devicegroup -ss` command.

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 Hitachi TrueCopy or Universal Replicator Device Group

An individual Hitachi TrueCopy or Universal Replicator device group can be in one of the following states:

- SMPL
- Regular Primary

- Regular Secondary
- Takeover Primary
- Takeover Secondary

The state of a particular device group is determined by using the value that is returned by the `pairvolchk -g devicegroup -ss` command. The following table describes the device group state associated with the values returned by the `pairvolchk` command.

TABLE 2-1 Individual Hitachi TrueCopy and Universal Replicator Device Group States

Output of <code>pairvolchk</code>	Individual Device Group State
11 = SMPL	SMPL
22 / 42 = PVOL_COPY 23 / 42 = PVOL_PAIR 26 / 46 = PVOL_PDUB 47 = PVOL_PFUL 48 = PVOL_PFUS	Regular Primary
24 / 44 = PVOL_PSUS 25 / 45 = PVOL_PSUE  For these return codes, determining the individual device group category requires that the <code>horcmd</code> process be active on the remote cluster so that the <code>remote-pair-state</code> for this device group can be obtained.	Regular Primary, if <code>remote-cluster-state != SSWS</code> or Takeover Secondary, if <code>remote-cluster-state == SSWS</code>  SSWS, when you use the <code>pairdisplay -g devicegroup -fc</code> command.
32 / 52 = SVOL_COPY 33 / 53 = SVOL_PAIR 35 / 55 = SVOL_PSUE 36 / 56 = SVOL_PDUB 57 = SVOL_PFUL 58 = SVOL_PFUS	Regular Secondary

TABLE 2-1 Individual Hitachi TrueCopy and Universal Replicator Device Group States (Continued)

Output of <code>pairvolchk</code>	Individual Device Group State
34 / 54 = SVOL_PSUS	Regular Secondary, if <code>local-cluster-state !=SSWS</code>  or Takeover Primary, if <code>local-cluster-state == SSWS</code>  SSWS, when you use the <code>pairdisplay -g devicegroup -fc</code> command.

## Determining the Aggregate Hitachi TrueCopy or Universal Replicator Device Group State

If a protection group contains only one Hitachi TrueCopy or Universal Replicator device group, then the aggregate device group state is the same as the individual device group state.

When a protection group contains multiple Hitachi TrueCopy or Universal Replicator device groups, the aggregate device group state is obtained as described in the following table.

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

Condition	Aggregate Device Group State
All individual device group states are SMPL	SMPL
All individual device group states are either Regular Primary or SMPL	Regular Primary
All individual device group states are either Regular Secondary or SMPL	Regular Secondary
All individual device group states are either Takeover Primary or SMPL	Takeover Primary
All individual device group states are either Takeover Secondary or SMPL	Takeover Secondary

The aggregate device group state cannot be obtained for any other combination of individual device group states. This is considered a pair-state validation failure.

## Validating the Local Role of the Protection Group Against the Aggregate Device Group State

The local role of a Hitachi TrueCopy or Universal Replicator protection group is validated against the aggregate device group state as described in the following table.

TABLE 2-3 Validating the Aggregate Device Group State Against the Local Role of a Protection Group

Aggregate Device Group State	Valid Local Protection Group Role
SMPL	primary or secondary
Regular Primary	primary
Regular Secondary	secondary
Takeover Primary	primary
Takeover Secondary	secondary

## EXAMPLE 2-12 Validating the Aggregate Device Group State

This example validates the state of a Hitachi TrueCopy or Universal Replicator device group against the role of the Hitachi TrueCopy or Universal Replicator protection group to which it belongs.

First, the protection group is created as follows:

```
phys-paris-1# geogg create -s paris-newyork-ps -o primary -d truecopy tcpg
```

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

```
phys-paris-1# geogg add-device-group -p fence_level=async devgroup1 tcpg
```

The current state of a Hitachi TrueCopy or Universal Replicator device group, devgroup1, is provided in the output of the pairdisplay command as follows:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345 1..P-VOL PAIR ASYNC,54321 609 -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL PAIR ASYNC,----- 1 -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345 2..P-VOL PAIR ASYNC,54321 610 -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL PAIR ASYNC,----- 2 -
```

The pairvolchk -g <DG> -ss command is run and returns a value of 23.

```
phys-paris-1# pairvolchk -g devgroup1 -ss
pairvolchk : Volstat is P-VOL.[status = PAIR fence = ASYNC]
phys-paris-1# echo $?
23
```

The output of the pairvolchk command is 23, which corresponds in [Table 2-1](#) to an individual device group state of Regular Primary. Because the protection group contains only one device group, the aggregate device group state is the same as the individual device group state. The device group state is valid because the local role of the protection group, specified by the -o option, is primary, as specified in [Table 2-3](#).



## ▼ How to Modify a Hitachi TrueCopy or Universal Replicator 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.

```
# geogg modify-device-group -p property [-p...] TCdevicegroupname 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*.

*TCdevicegroupname*

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 2–13 Modifying the Properties of a Hitachi TrueCopy or Universal Replicator Data Replication Device Group

This example modifies the properties of a data replication device group that is part of a Hitachi TrueCopy or Universal Replicator protection group.

```
# geogg modify-device-group -p fence_level=async tcgd tcpg
```

## ▼ How to Delete a Data Replication Device Group From a Hitachi TrueCopy or Universal Replicator 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 a Hitachi TrueCopy or Universal Replicator Protection Group](#)” on page 46. For information about deleting application resource groups from a protection group, refer to “[How to Delete an Application Resource Group From a Hitachi TrueCopy or Universal Replicator Protection Group](#)” on page 49.

### 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.

```
# geopg 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 a Hitachi TrueCopy or Universal Replicator protection group, the corresponding Oracle Solaris Cluster resource, *r-tc-protectiongroupname-devicegroupname*, is removed from the replication resource group. As a result, the deleted device group is no longer monitored. The resource group is removed when the protection group is deleted.

## Example 2–14 Deleting a Replication Device Group From a Hitachi TrueCopy or Universal Replicator Protection Group

This example removes a Hitachi TrueCopy or Universal Replicator data replication device group.

```
# geopg remove-device-group tcdg tcpg
```

# Replicating the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster

After you have configured data replication, resource groups, and resources on your primary and secondary clusters, you can replicate the configuration of the protection group to the secondary cluster.

## ▼ How to Replicate the Hitachi TrueCopy or Universal Replicator Protection Group Configuration to a Secondary Cluster

**Before You Begin** Before you replicate the configuration of a Hitachi TrueCopy or Universal Replicator protection group to a secondary 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 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 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
```

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
```

### 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 the only 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 2–15 Replicating the Hitachi TrueCopy or Universal Replicator Protection Group to a Partner Cluster

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

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

## Activating a Hitachi TrueCopy or Universal Replicator Protection Group

When you activate a protection group, the protection group assumes the role that you assigned to it during configuration. For more information about configuring protection groups, see [“How to Create and Configure a Hitachi TrueCopy or Universal Replicator Protection Group That Does Not Use Oracle Real Application Clusters”](#) on page 34.

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 a Hitachi TrueCopy or Universal Replicator 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 aggregate device group state as described in [Table 2–3](#). If validation is successful, data replication is started.
- 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.
- When a protection group is activated on the secondary cluster, the application resource groups are *not* started.

The Hitachi TrueCopy or Universal Replicator command that is used to start data replication depends on the following factors:

- Aggregate device group state
- Local role of the protection group
- Current pair state

If a protection group has a consistency group defined, the fence level is `async` and the device group is in `SMPL` state, then you create the device group with the `pair create` command when the `geopg start` command is run with the `-f` flag. If a protection group has a consistency group defined, the fence level is not `async` and the device group is in `SMPL` state then you create the device group with the `pair create` command when you run the `geopg start` command with the `-fgflags`. On arrays that only support the Hitachi TrueCopy software, the `-fg fence` level option to the `geopg` command is not supported. Thus, on such arrays, the user should only define the `ctgid` on the protection group, if that protection group only has device groups of fence level `async`.

The following table describes the Hitachi TrueCopy or Universal Replicator command that is used to start data replication for each of the possible combinations of factors. In the commands, `dg` is the device group name and `fl` is the fence level that is configured for the device group.

**TABLE 2-4** Commands Used to Start Hitachi TrueCopy or Universal Replicator Data Replication

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopy or Universal Replicator Start Command
SMPL	primary or secondary	paircreate -vl -g dg -f fl paircreate -vl -g dg -f fl ctgid paircreate -vr -g dg -f fl paircreate -vr -g dg -f fl ctgid  All commands require that the horcmd process is running on the remote cluster. Device pairs can be started with or without a specified CTGID.
Regular Primary	primary	If the local state code is 22, 23, 25, 26, 29, 42, 43, 45, 46, or 47, no command is run because data is already being replicated.  If the local state code is 24, 44, or 48, then the following command is run: pairresync -g dg [-l].  If the local state code is 11, then the following command is run: paircreate -vl -g dg -f fl.  Both commands require that the horcmd process is running on the remote cluster.
Regular Secondary	secondary	If the local state code is 32, 33, 35, 36, 39, 52, 53, 55, 56, or 57, no command is run because data is already being replicated.  If the local state code is 34, 54, or 58, then the following command is run: pairresync -g dg  If the local state code is 11, the following command is run: paircreate -vr -g dg -f fl  Both commands require that the horcmd process is up on the remote cluster.
Takeover Primary	primary	If the local state code is 34 or 54, the following command is run: pairresync -swaps -g.  If the local state code is 11, then the following command is run: paircreate -vl -g dg -f fl.  The paircreate command requires that the horcmd process is running on the remote cluster.

TABLE 2-4 Commands Used to Start Hitachi TrueCopy or Universal Replicator Data Replication  
(Continued)

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopy or Universal Replicator Start Command
Takeover Secondary	secondary	<p>If the local state code is 24, 44, 25, or 45, the following command is run: <code>pairresync -swapp -g dg</code>.</p> <p>If the local state code is 11, the following command is run: <code>paircreate -vr -g dg -f fl</code>.</p> <p>Both commands require that the <code>horcmd</code> process is running on the remote cluster.</p>

## ▼ How to Activate a Hitachi TrueCopy or Universal Replicator 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 Activate the protection group.

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

```
# geogg 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, such as `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 Oracle Solaris Cluster commands to bring resource groups and resources online.

### Example 2-16 How the Geographic Edition Software Issues the Command to Start Replication

This example illustrates how the Geographic Edition determines the Hitachi TrueCopy or Universal Replicator command that is used to start data replication.

First, the Hitachi TrueCopy or Universal Replicator protection group is created.

```
phys-paris-1# geopg create -s paris-newyork-ps -o primary -d truecopy tcpg
```

A device group, `devgroup1`, is added to the protection group.

```
phys-paris-1# geopg add-device-group -p fence_level=async devgroup1 tcpg
```

The current state of a Hitachi TrueCopy or Universal Replicator device group, `devgroup1`, is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345 1..SMPL ---- -, ----- -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..SMPL ---- -, ----- -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345 2..SMPL ---- -, ----- -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..SMPL ---- -, ----- -
```

The aggregate device group state is `SMPL`.

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

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

The Geographic Edition software runs the `paircreate -g devgroup1 -vl -f async` command at the data replication level. If the command is successful, the state of `devgroup1` is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345 1..P-VOL COPY ASYNC,54321 609 -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL COPY ASYNC,----- 1 -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345 2..P-VOL COPY ASYNC,54321 610 -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL COPY ASYNC,----- 2 -
```

### Example 2-17 Activating a Hitachi TrueCopy or Universal Replicator Protection Group Globally

This example activates a protection group globally.

```
# geopg start -e global tcpg
```

The protection group, `tcpg`, is activated on both clusters where the protection group is configured.



**Example 2–18** Activating a Hitachi TrueCopy or Universal Replicator 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 tcpg
```

## Deactivating a Hitachi TrueCopy or Universal Replicator 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 a Hitachi TrueCopy or Universal Replicator 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 as described in [Table 2–3](#). 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 that are 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 Hitachi TrueCopy or Universal Replicator command that is used to stop data replication depends on the following factors:

- Aggregate device group state
- Local role of the protection group
- Current pair state

The following table describes the Hitachi TrueCopy or Universal Replicator command used to stop data replication for each of the possible combinations of factors. In the commands, *dg* is the device group name.

TABLE 2-5 Commands Used to Stop Hitachi TrueCopy or Universal Replicator Data Replication

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopy or Universal Replicator Stop Command
SMPL	primary or secondary	No command is run because no data is being replicated.
Regular Primary	primary	If the local state code is 22, 23, 26, 29, 42, 43, 46, or 47, then the following command is run: <code>pairsplit -g dg [-l]</code> .  If the local state code is 11, 24, 25, 44, 45, or 48, then no command is run because no data is being replicated.
Regular Secondary	secondary	If the local state code is 32, 33, 35, 36, 39, 52, 53, 55, 56, or 57, the following command is run: <code>pairsplit -g dg</code> .  If the local state code is 33 or 53 and the remote state is PSUE, no command is run to stop replication.  If the local state code is 11, 34, 54, or 58, then no command is run because no data is being replicated.
Takeover Primary	primary	No command is run because no data is being replicated.
Takeover Secondary	secondary	No command is run because no data is being replicated.

## ▼ How to Deactivate a Hitachi TrueCopy or Universal Replicator 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, its application resource groups are also unmanaged.

```
# geopg 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 2–19 How the Geographic Edition Software Issues the Command to Stop Replication

This example illustrates how the Geographic Edition software determines the Hitachi TrueCopy or Universal Replicator command that is used to stop data replication.

The current state of the Hitachi TrueCopy or Universal Replicator device group, `devgroup1`, is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345 1..P-VOL PAIR ASYNC,54321 609 -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL PAIR ASYNC,----- 1 -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345 2..P-VOL PAIR ASYNC,54321 610 -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL PAIR ASYNC,----- 2 -
```

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

```
phys-paris-1# geogg add-device-group -p fence_level=async devgroup1 tcpg
```

The Geographic Edition software runs the `pairvolchk -g <DG> -ss` command at the data replication level, which returns a value of 43.

```
# pairvolchk -g devgroup1 -ss
Volstat is P-VOL.[status = PAIR fence = ASYNC]
phys-paris-1# echo $?
43
```

Next, the protection group, `tcpg`, is deactivated by using the `geogg stop` command.

```
phys-paris-1# geogg stop -s local tcpg
```

The Geographic Edition software runs the `pairsplit -g devgroup1` command at the data replication level.

If the command is successful, the state of `devgroup1` is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345 1..P-VOL PSUS ASYNC,54321 609 -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL SSUS ASYNC,----- 1 -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345 2..P-VOL PSUS ASYNC,54321 610 -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL SSUS ASYNC,----- 2 -
```

### Example 2–20 Deactivating a Protection Group on All Clusters

This example deactivates a protection group on all clusters.

```
# geopg stop -e global tcpg
```

### Example 2–21 Deactivating a Protection Group on a Local Cluster

This example deactivates a protection group on the local cluster.

```
# geopg stop -e local tcpg
```

### Example 2–22 Stopping Data Replication While Leaving the Protection Group Online

This example stops only data replication on a local cluster.

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

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 tcpg
```

### Example 2–23 Deactivating a Hitachi TrueCopy or Universal Replicator Protection Group While Keeping Application Resource Groups Online

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

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

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

2. Deactivate the protection group.

```
# geopg stop -e global tcpg
```

# Resynchronizing a Hitachi TrueCopy or Universal Replicator 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 a Hitachi TrueCopy or Universal Replicator Protection Group” on page 65](#).

### 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.

```
# geopg update protectiongroupname
```

```
protectiongroupname
```

Specifies the name of the protection group.

### Example 2–24 Resynchronizing a Protection Group

This example resynchronizes a protection group.

```
# geopg update tcpg
```

# Checking the Runtime Status of Hitachi TrueCopy and Universal Replicator Data Replication

You can obtain an overall view of the status of replication, as well as a more detailed runtime status of the Hitachi TrueCopy or Universal Replicator replication resource groups. The following sections describe the procedures for checking each status.

- “Displaying a Hitachi TrueCopy or Universal Replicator Runtime Status Overview” on page 70
- “Displaying a Detailed Hitachi TrueCopy or Universal Replicator Runtime Status” on page 71

## Displaying a Hitachi TrueCopy or Universal Replicator Runtime Status Overview

The status of each Hitachi TrueCopy or Universal Replicator 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 Hitachi TrueCopy or Universal Replicator device group.

# `clresource status`

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 2–6](#).

## Displaying a Detailed Hitachi TrueCopy or Universal Replicator 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:

*rg-tc\_truecopyprotectiongroupname*

If you add a Hitachi TrueCopy or Universal Replicator 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:

*r-tc\_truecopyprotectiongroupname-truecopydevicegroupname*

You can monitor the status of replication of this device group by checking the `Status` and `Status Message` of this resource. Use the `clresource status` command to display the 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 Hitachi TrueCopy or Universal Replicator replication resource group is `Online`.

**TABLE 2–6** Status and Status Messages of an Online Hitachi TrueCopy or Universal Replicator Replication Resource Group

Status	Status Message
Online	P-Vol/S-Vol:PAIR
Online	P-Vol/S-Vol:PAIR:Remote horcmd not reachable
Online	P-Vol/S-Vol:PFUL
Online	P-Vol/S-Vol:PFUL:Remote horcmd not reachable
Degraded	SMPL:SMPL

**TABLE 2-6** Status and Status Messages of an Online Hitachi TrueCopy or Universal Replicator Replication Resource Group *(Continued)*

Status	Status Message
Degraded	SMPL:SMPL:Remote horcmd not reachable
Degraded	P-Vol/S-Vol:COPY
Degraded	P-Vol/S-Vol:COPY:Remote horcmd not reachable
Degraded	P-Vol/S-Vol:PSUS
Degraded	P-Vol/S-Vol:PSUS:Remote horcmd not reachable
Degraded	P-Vol/S-Vol:PFUS
Degraded	P-Vol/S-Vol:PFUS:Remote horcmd not reachable
Faulted	P-Vol/S-Vol:PDFUB
Faulted	P-Vol/S-Vol:PDUB:Remote horcmd not reachable
Faulted	P-Vol/S-Vol:PSUE
Faulted	P-Vol/S-Vol:PSUE:Remote horcmd not reachable
Degraded	S-Vol:SSWS:Takeover Volumes
Faulted	P-Vol/S-Vol:Suspicious role configuration. Actual Role=x, Config Role=y

For more information about these values, refer to the Hitachi TrueCopy or Universal Replicator documentation.

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



# Migrating Services That Use Hitachi TrueCopy and Universal Replicator 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 Hitachi TrueCopy or Universal Replicator Data Replication” on page 73
- “Migrating Services That Use Hitachi TrueCopy or Universal Replicator Data Replication With a Switchover” on page 75
- “Forcing a Takeover on a System That Uses Hitachi TrueCopy or Universal Replicator Data Replication” on page 77
- “Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy or Universal Replicator Replication” on page 81
- “Recovering From a Switchover Failure on a System That Uses Hitachi TrueCopy or Universal Replicator Replication” on page 91
- “Recovering From a Hitachi TrueCopy or Universal Replicator Data Replication Error” on page 94

## Detecting Cluster Failure on a System That Uses Hitachi TrueCopy or Universal Replicator 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 74
- “Detecting Secondary Cluster Failure” on page 74

## Detecting Primary Cluster Failure

When the primary cluster for a given 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.

- 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 timeout 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](#).

## Detecting Secondary Cluster Failure

When a secondary cluster for a given 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 occur:

- Heartbeat failure is detected by a partner cluster.
- The heartbeat is activated in emergency mode to verify that the secondary cluster is dead.
- 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 Hitachi TrueCopy or Universal Replicator Data Replication With a Switchover

Perform a switchover of a Hitachi TrueCopy or Universal Replicator protection group when you want to migrate services to the partner cluster in an orderly fashion. A switchover consists of the following actions:

- Application services are 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`.

This section provides the following information:

- [“Validations That Occur Before a Switchover” on page 75](#)
- [“Results of a Switchover From a Replication Perspective” on page 76](#)
- [“How to Switch Over a Hitachi TrueCopy or Universal Replicator Protection Group From Primary to Secondary” on page 76](#)

## 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 Hitachi TrueCopy or Universal Replicator device group is in a valid aggregate device group state. Then, it checks that the local device group states on the target primary cluster, `cluster-newyork`, are 23, 33, 43, or 53. The local device group state is returned by the `pairvolchk -g device-group-name -ss` command. These values correspond to a `PVOL_PAIR` or `SVOL_PAIR` state. The Hitachi TrueCopy or Universal Replicator commands that are run on the new primary cluster, `cluster-newyork`, are described in the following table.

TABLE 3-1 Hitachi TrueCopy and Universal Replicator Switchover Validations on the New Primary Cluster

Aggregate Device Group State	Valid Device Group State on Local Cluster	Hitachi TrueCopy or Universal Replicator Switchover Commands That Are Run on <code>cluster-newyork</code>
SMPL	None	None

**TABLE 3-1** Hitachi TrueCopy and Universal Replicator Switchover Validations on the New Primary Cluster *(Continued)*

Aggregate Device Group State	Valid Device Group State on Local Cluster	Hitachi TrueCopy or Universal Replicator Switchover Commands That Are Run on <code>cluster-newyork</code>
Regular primary	23, 43	No command is run, because the Hitachi TrueCopy or Universal Replicator device group is already in the PVOL_PAIR state.
Regular secondary	33, 53	<code>horctakeover -g dg [-t]</code>  The <code>-t</code> option is specified when the <code>fence_level</code> of the Hitachi TrueCopy or Universal Replicator device group is <code>async</code> . The value is calculated as 80% of the <code>Timeout</code> property of the protection group. For example, if the protection group has a <code>Timeout</code> of 200 seconds, the value of <code>-t</code> used in this command is 80% of 200 seconds, or 160 seconds.
Takeover primary	None	None
Takeover secondary	None	None

## 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 PVOL\_PAIR volumes that were in place before the switchover become the SVOL\_PAIR volumes. The SVOL\_PAIR volumes in place before the switchover become the PVOL\_PAIR volumes. Data replication will continue from the new PVOL\_PAIR volumes to the new SVOL\_PAIR 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 a Hitachi TrueCopy or Universal Replicator 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 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 tcpg
```

## Forcing a Takeover on a System That Uses Hitachi TrueCopy or Universal Replicator 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 not be consistent 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 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 78
- “[Results of a Takeover From a Replication Perspective](#)” on page 79
- “[How to Force Immediate Takeover of Hitachi TrueCopy or Universal Replicator Services by a Secondary Cluster](#)” on page 80

## 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 Hitachi TrueCopy or Universal Replicator device group is in a valid aggregate device group state. Then, the replication subsystem checks that the local device group states on the target primary cluster, `cluster-newyork`, are not 32 or 52. These values correspond to a `SVOL_COPY` state, for which the `horctakeover` command fails. The Hitachi TrueCopy or Universal Replicator commands that are used for the takeover are described in the following table.

TABLE 3-2 Hitachi TrueCopy or Universal Replicator Takeover Validations on the New Primary Cluster

Aggregate Device Group State	Valid Local State Device Group State	Hitachi TrueCopy or Universal Replicator Takeover Commands That Are Run on <code>cluster-newyork</code>
SMPL	All	No command is run.
Regular primary	All	No command is run.
Regular secondary	All Regular secondary states except 32 or 52  For a list of Regular secondary states, refer to <a href="#">Table 2-1</a> and <a href="#">Table 2-2</a> .	<code>horctakeover -S -g dg [-t]</code>  The <code>-t</code> option is given when the <code>fence_level</code> of the Hitachi TrueCopy or Universal Replicator device group is <code>async</code> . The value is calculated as 80% of the <code>Timeout</code> property of the protection group. For example, if the protection group has a <code>Timeout</code> of 200 seconds, the value of <code>-t</code> used in this command will be 80% of 200 seconds, or 160 seconds.
Takeover primary	All	No command is run.
Takeover secondary	All	<code>pairsplit -R-g dg pairsplit -S-g dg</code>

## Results of a Takeover From a Replication Perspective

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

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

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



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

## ▼ How to Force Immediate Takeover of Hitachi TrueCopy or Universal Replicator 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 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 `tcpg` 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 tcpg
```

**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 Hitachi TrueCopy or Universal Replicator 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 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.

To continue using the new primary, `cluster-newyork`, as the primary cluster and the original primary cluster, `cluster-paris`, as the secondary after the original primary is running again, resynchronize and revalidate the protection group configuration without performing a switchover or takeover.

This section provides the following information:

- “[How to Resynchronize and Revalidate the Protection Group Configuration](#)” on page 81
- “[How to Perform a Failback-Switchover on a System That Uses Hitachi TrueCopy or Universal Replicator Replication](#)” on page 83
- “[How to Perform a Failback-Takeover on a System That Uses Hitachi TrueCopy or Universal Replicator Replication](#)” on page 86

## ▼ 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 `cluster-paris` could be reached during the takeover from `cluster-newyork`.

**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.**

```
# 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.

```
# geopg update protectiongroupname
```

*protectiongroupname*

Specifies the name of the protection group.

For more information about synchronizing protection groups, see “[Resynchronizing a Hitachi TrueCopy or Universal Replicator Protection Group](#)” on page 69.

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

```
# geopg validate protectiongroupname
```

*protectiongroupname*

Specifies a unique name that identifies a single protection group.

For more information, see “[How to Validate a Hitachi TrueCopy or Universal Replicator Protection Group](#)” on page 45.

**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`.

```
# 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.




---

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

---

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

For more information about the `geopg start` command, see [“How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group” on page 63](#).

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

The state of the protection group on `cluster-newyork` must be OK.

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

Refer to the Protection Group section of the output.

The protection group has a local state of OK when the Hitachi TrueCopy or Universal Replicator device groups on `cluster-newyork` have a state of `PVOL_PAIR` and the Hitachi TrueCopy or Universal Replicator device groups on `cluster-paris` have a state of `SVOL_PAIR`.

## ▼ How to Perform a Failback-Switchover on a System That Uses Hitachi TrueCopy or Universal Replicator 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` could be reached during the takeover from `cluster-newyork`.

**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 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 a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 69.

## 2 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# geogg validate protectiongroupname
```

*protectiongroupname*

Specifies a unique name that identifies a single protection group.

For more information, see [“How to Validate a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 45.

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

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

```
phys-paris-1# geogg 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.



**Caution** – Do not use the `-n` option 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 `geogg start` command, see [“How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 63.

## 4 Confirm that the data is completely synchronized.

The state of the protection group on `cluster-newyork` must be OK.

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

Refer to the Protection Group section of the output.

The protection group has a local state of OK when the Hitachi TrueCopy or Universal Replicator device groups on `cluster-newyork` have a state of PVOL\_PAIR and the Hitachi TrueCopy or Universal Replicator device groups on `cluster-paris` have a state of SVOL\_PAIR.

- 5 On both partner clusters, ensure that the protection group is activated.

```
# geoadm status
```

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

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

For more information, see [“How to Switch Over a Hitachi TrueCopy or Universal Replicator Protection Group From Primary to Secondary”](#) on page 76.

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

- 7 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 Hitachi TrueCopy or Universal Replicator protection group.

```
# clresourcegroup status -v
# clresource status -v
```

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 Hitachi TrueCopy and Universal Replicator Data Replication”](#) on page 70.

## ▼ How to Perform a Failback-Takeover on a System That Uses Hitachi TrueCopy or Universal Replicator 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** – Conditionally, you can 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 use the `-n` option when you run the `geogg start` command.

---

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

- 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 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# geogg stop -e local protectiongroupname
```

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

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

**e. Place the Hitachi TrueCopy or Universal Replicator device group, `devgroup1`, in the SMPL state.**

Use the `pairsplit` commands to place the Hitachi TrueCopy or Universal Replicator device groups that are in the protection group on both `cluster-paris` and

`cluster-newyork` in the SMPL state. The `pairsplit` command you use depends on the pair state of the Hitachi TrueCopy or Universal Replicator device group. The following table gives some examples of the command you need to use on `cluster-paris` for some typical pair states.

Pair State on <code>cluster-paris</code>	Pair State on <code>cluster-newyork</code>	<code>pairsplit</code> Command Used on <code>cluster-paris</code>
PSUS or PSUE	SSWS	<code>pairsplit -R -g dgname</code> <code>pairsplit -S -g dgname</code>
SSUS	PSUS	<code>pairsplit -S -g dgname</code>

If the command is successful, the state of `devgroup1` is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345 1..SMPL ---- -,----- ---- -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..SMPL ---- -,----- ---- -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345 2..SMPL ---- -,----- ---- -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..SMPL ---- -,----- ---- -
.
```

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

```
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 69](#).

**2 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 a Hitachi TrueCopy or Universal Replicator Protection Group” on page 45](#).



### 3 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`.

```
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
```

Prevents the start of data replication at protection group startup.

---

**Note** – You must use the `-n` option.

---

*protectiongroupname*

Specifies the name of the protection group.

For more information, see [“How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 63.

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

### 4 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 Hitachi TrueCopy or Universal Replicator Services by a Secondary Cluster”](#) on page 80.

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`.

## 5 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 a Hitachi TrueCopy or Universal Replicator Protection Group”](#) on page 63.

## 6 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
```

Check the runtime status of the application resource group and data replication for each Hitachi TrueCopy or Universal Replicator protection group.

```
# clresourcegroup status -v
```

```
# clresource status -v
```

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 Hitachi TrueCopy and Universal Replicator Data Replication”](#) on page 70.

# Recovering From a Switchover Failure on a System That Uses Hitachi TrueCopy or Universal Replicator Replication

When you run the `geopg switchover` command, the `horctakeover` command runs at the Hitachi TrueCopy or Universal Replicator data replication level. If the `horctakeover` command returns a value of 1, the switchover is successful.

In Hitachi TrueCopy and Universal Replicator terminology, a switchover is called a *swap-takeover*. In some cases, the `horctakeover` command might not be able to perform a swap-takeover. In these cases, a return value other than 1 is returned, which is considered a switchover failure.

---

**Note** – In a failure, the `horctakeover` command usually returns a value of 5, which indicates a SVOL-SSUS-takeover.

---

One reason the `horctakeover` command might fail to perform a swap-takeover is because the data replication link, ESCON/FC, is down.

Any result other than a swap-takeover implies that 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.

- [“Switchover Failure Conditions” on page 91](#)
- [“Recovering From Switchover Failure” on page 92](#)
- [“How to Make the Original Primary Cluster Primary for a Hitachi TrueCopy or Universal Replicator Protection Group” on page 93](#)
- [“How to Make the Original Secondary Cluster Primary for a Hitachi TrueCopy or Universal Replicator Protection Group” on page 93](#)

## Switchover Failure Conditions

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

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

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

While processing the `geopg switchover` command, the `horctakeover` command performs an SVOL-SSUS-takeover and returns a value of 5 for the Hitachi TrueCopy or Universal Replicator device group, `devgroup1`. As a result, the `geopg switchover` command returns with the following failure message:

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

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

```
cluster-paris:
    tcpg role: Secondary
cluster-newyork:
    tcpg role: Secondary
```

```
phys-newyork-1# pairdisplay -g devgroup1 -fc
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S, Status,Fence,%, P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL SSWS ASYNC,100 1 -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321 1..P-VOL PSUS ASYNC,100 609 -
```

## 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 Hitachi TrueCopy or Universal Replicator device group, `devgroup1`, in the SMPL state.

Use the `pairsplit` commands to place the device groups that are in the protection group on both `cluster-paris` and `cluster-newyork` in the SMPL state. For the pair states that are shown in the previous section, run the following `pairsplit` commands:

```
phys-newyork-1# pairsplit -R -g devgroup1
phys-newyork-1# pairsplit -S -g devgroup1
```

2. Designate one of the clusters Primary for the protection group.

Designate 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.

Designate 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 `horctakeover` command did not perform a swap-takeover, 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 that appears on the original primary cluster, you must not make the original secondary cluster Primary.

---

## ▼ How to Make the Original Primary Cluster Primary for a Hitachi TrueCopy or Universal Replicator Protection Group

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

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

- 2 Resynchronize the configuration of the protection group.

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

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

After the `geopg update` command completes successfully, `tcpg` has the following role on each cluster:

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

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

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

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 a Hitachi TrueCopy or Universal Replicator Protection Group

- 1 Resynchronize the configuration of the protection group.

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

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

After the `geopg update` command completes successfully, `tcpg` has the following role on each cluster:

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

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

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

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 a Hitachi TrueCopy or Universal Replicator 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 section provides the following information:

- [“How to Detect Data Replication Errors” on page 94](#)
- [“How to Recover From a Hitachi TrueCopy or Universal Replicator Data Replication Error” on page 95](#)

### How to Detect Data Replication Errors

For information about how different Resource status values map to actual replication pair states, see [Table 2-6](#).

You can check the status of the replication resources by using the `clresource status` command as follows:

```
phys-paris-1# clresource status -v
```

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

```
=== Cluster Resources ===
```

Resource Name	de Name	State	Status Message
r-tc-tcpg1-devgroup1	phys-paris-2	Offline	Offline
	phys-paris-1	Online	Faulted - P-VOL:PSUE
hasp4nfs	phys-paris-2	Offline	Offline
	phys-paris-1	Offline	Offline

The aggregate resource status for all device groups in the protection group is provided by using the `geoadm status` command. For example, the output of the `clresource status` command in the preceding example indicates that the Hitachi TrueCopy or Universal Replicator device

group, devgroup1, is in the PSUE state on cluster-paris. Table 2–6 indicates that the PSUE 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 "tcpg" : 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

Pending Operations
  Protection Group : "tcpg"
  Operations       : start
```

## ▼ How to Recover From a Hitachi TrueCopy or Universal Replicator 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 Hitachi TrueCopy or Universal Replicator documentation to determine the causes of the FAULTED state. This state is indicated as PSUE.

**2 Recover from the faulted state by using the Hitachi TrueCopy or Universal Replicator 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 Hitachi TrueCopy or Universal Replicator protection group.

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

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

Specifies the name of the Hitachi TrueCopy or Universal Replicator protection group.

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

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



# Geographic Edition Properties for Hitachi TrueCopy and Universal Replicator

---

This appendix provides the properties of Geographic Edition data replication device groups. The following sections are in this appendix:

- [“Hitachi TrueCopy and Universal Replicator Properties” on page 97](#)
- [“Hitachi TrueCopy and Universal Replicator Properties That Must Not Be Changed” on page 99](#)

---

**Note** – The property values, such as True and False, are *not* case sensitive.

---

## Hitachi TrueCopy and Universal Replicator Properties

The following table describes the Hitachi TrueCopy and Universal Replicator properties that the Geographic Edition software defines.

TABLE A-1 Hitachi TrueCopy and Universal Replicator Properties

Property	Description
Data Replication Property: Cluster_dgs (string array)	<p>Lists the 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 listed in the cluster_dgs protection group property 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 Hitachi TrueCopy and Universal Replicator Properties (Continued)

Property	Description
Data Replication Property: NodeList (string array)	<p>Lists the host names of the machines that can be primary for the replication mechanism. This list is comma delimited.</p> <p>Tuning recommendations: This property can be tuned at any time.</p> <p>Category: Optional</p> <p>Default: All nodes in the cluster</p>
Device Group Property: Fence_level (enum)	<p>Defines the fence level that is used by the device group. The fence level determines the level of consistency among the primary and secondary volumes for that device group. Possible values are Never and Async. To use the data or status fence levels, contact your Oracle representative.</p> <p><b>Note</b> – If you specify a Fence_level of never, the data replication roles do not change after you perform a takeover.</p> <p>For more information about setting this property, see <a href="#">“How to Add a Data Replication Device Group to a Hitachi TrueCopy or Universal Replicator Protection Group”</a> on page 51.</p> <p>Tuning recommendations: This property can only be tuned when the protection group is offline.</p> <p>Category: Required</p> <p>Default: None</p>
Data Replication Property: Ctgid (integer)	<p>Specifies the consistency group ID (CTGID) of the protection group. Once the CTGID of a protection group has been set, all Hitachi TrueCopy or Universal Replicator device groups thereafter added to the protection group either must be uninitialized or must already have the same CTGID as the protection group.</p> <p>Attempting to add an initialized device group to a protection group results in an error if the CTGID of the device group differs from the CTGID of the protection group. A device group with the same CTGID as a protection group must be added to that protection group.</p> <p>Tuning recommendations: This property can only be tuned at creation.</p> <p>Category: Optional</p> <p>Default: None</p>

# Hitachi TrueCopy and Universal Replicator Properties That Must Not Be Changed

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

For Hitachi TrueCopy and Universal Replicator, do not edit the following properties:

- `Dev_group` – Specifies the Hitachi TrueCopy or Universal Replicator device group that contains the volumes that are being replicated.
- `Replication_role` – Defines the local data replication role.



# Index

---

## A

- activating protection groups, 60–65
- administering
  - data replication, 11–27, 29–72
  - device groups, 50–58
- administration tasks, 12–13
- aggregate state, of device groups, 55
- application resource groups
  - administering, 47–50
  - creating, 48–49
  - removing, 49–50
- asynchronous mode, requirements, 13
- asynchronous mode replication
  - Hitachi Universal Replicator
    - data consistency, 13–14
- asynchronous replication
  - data consistency
    - Hitachi Universal Replicator, 36–39

## C

- commands
  - to start replication, 62–63
  - to stop replication, 66
- configuring
  - /etc/horcm.conf file
    - on primary cluster, 16
    - on secondary cluster, 23–24
  - Hitachi TrueCopy or Universal Replicator software, 13–27
    - on primary cluster, 15–22

- configuring, Hitachi TrueCopy or Universal Replicator software (*Continued*)
  - on secondary cluster, 22–27
  - local file system, 21–22
  - protection groups, 34–36
  - ZFS zpools, 19–21
- consistency group IDs
  - setting
    - on Hitachi device groups, 38–39
    - on protection groups, 38–39
- creating
  - application resource group, 48–49
  - protection groups, 34–36
    - while application offline, 30
    - while application online, 31–33
  - replication device group, 51–52

## D

- data consistency
  - Hitachi Universal Replicator
    - asynchronous replication, 36–39
    - guaranteeing, 13–14
- data recovery, 81–90
  - failback-switchover, 83–86
  - failback-takeover, 86–90
- deactivating protection groups, 65–68
- deleting
  - application resource groups, 49–50
  - protection groups, 46–47
  - replication device group, 57–58

detecting failure, 73–74

device groups

adding to protection group, 51–52

administering, 50–58

modifying, 57

property validations, 52–53

removing, 57–58

state validations, 53–56

aggregate state, 55

individual state, 53–55

DID, with raw-disk device groups, 18–19

disaster recovery, data consistency, 13–14

## E

error

detection, 94–95

recovery, 95–96

/etc/horcm.conf file, 17–18

on primary cluster, 16, 17–18

on secondary cluster, 23–24

## F

failback-switchover, 83–86

failback-takeover, 86–90

failure

detecting, 73–74

primary cluster, 74

secondary cluster, 74

failure conditions, switchover, 91–92

## H

HAStoragePlus resource

configuring, 21–22

configuring with ZFS, 19–21

help, 10

Hitachi TrueCopy and Universal Replicator

administering data replication with, 11–27

administration tasks, 12–13

configuring primary cluster, 15–22

Hitachi TrueCopy and Universal Replicator

(*Continued*)

initial software configuration, 13–27

requirements, 14–15

Hitachi TrueCopy or Universal Replicator

activating protection groups, 60–65

administering data replication with, 29–72

data recovery, 81–90

failback-switchover, 83–86

failback-takeover, 86–90

deactivating protection groups, 65–68

detecting failure, 73–74

device groups

properties, 52–53

subsystem validations, 52–53

migrating services that use, 73–96

properties of, 97–98

recovering from errors, 94–96

recovering from switchover failure, 91–94

runtime status

detailed, 71–72

overall, 70–71

state and status messages, 71–72

starting replication, 62–63

stopping replication, 66

Hitachi Universal Replicator

asynchronous mode replication

data consistency, 13–14

consistency group ID, 36–39

HORCM\_DEV

/etc/horcm.conf, 16, 17–18

HORCM\_LDEV

/etc/horcm.conf, 16, 17–18

horctakeover command, switchover failure, 91–94

## I

individual state, of device groups, 53–55

## L

local file-system configuration, 21–22

**M**

- migrating services, 73–96
- modifying
  - protection groups, 44–45
  - replication device group, 57

**O**

- Oracle Real Application Clusters, requirements, 39–40

**P**

- primary cluster
  - configuration of, 15–22
  - data recovery, 81–90
  - failure detection, 74
  - restoring as primary, 93
  - switchover, 75–77
- properties, Hitachi TrueCopy or Universal Replicator, 97–98
- protection groups
  - activating, 60–65
  - adding application resource group to, 48–49
  - adding device group to, 51–52
  - configuring, 34–36
  - creating, 34–36
    - while application offline, 30
    - while application online, 31–33
    - while application resource group online, 36
  - creation strategies, 29–33
  - deactivating, 65–68
  - deleting, 46–47
  - local role
    - validated against aggregate state, 55–56
  - modifying, 44–45
  - modifying device group from, 57
  - removing
    - application resource groups, 49–50
    - device group from, 57–58
  - replicating configuration of, 59–60
  - resynchronizing, 69
  - validating, 45–46

**R**

- raw-disk device groups, 18–19
- recovery
  - See* data recovery
  - from replication error, 94–96
  - from switchover failure, 91–94
- replication
  - adding device group, 51–52
  - configuration, 24–26, 26–27
  - detecting errors in, 94–95
  - error recovery, 94–96
  - forcing takeover, 77–80
  - Hitachi TrueCopy and Universal Replicator software, 11–27
  - Hitachi TrueCopy or Universal Replicator stop command, 66
  - Hitachi TrueCopy or Universal Replicator start command, 62–63
  - initial configuration of, 13–27
  - migrating services that use, 73–96
  - modifying device group, 57
  - protection group configuration, 59–60
  - removing device groups, 57–58
  - runtime status details, 71–72
  - runtime status of, 70–72
  - runtime status overview, 70–71
  - switchover failure recovery, 91–94
  - task summary, 12–13
- requirements
  - asynchronous mode, 13
  - Hitachi TrueCopy and Universal Replicator, 14–15
  - Oracle Real Application Clusters, 39–40
  - ZFS zpools, 20
- resource groups
  - application, 47–50
  - Hitachi TrueCopy or Universal Replicator replication status, 71–72
- resynchronizing, protection groups, 69
- runtime status
  - detailed, 71–72
  - overview, 70–71
  - replication, 70–72
  - state and status messages, 71–72

## **S**

- secondary cluster
  - configuring, 22–27
  - failure detection, 74
  - making primary, 93–94
  - switchover, 75–77
- state, device group, 53–56
- switchover, 75–77
  - failure
    - conditions, 91–92
    - recovering from, 92–93
  - Hitachi TrueCopy or Universal Replicator, 76–77
  - results of, 76
  - validations, 75–76
- switchover failure, recovering from, 91–94

## **T**

- takeover, 77–80
  - failback-switchover, 83–86
  - failback-takeover, 86–90
  - forcing, 80
  - results of, 79–80
  - validations, 78–79
- technical support, 10

## **V**

- validating
  - device group properties, 52–53
  - protection groups, 45–46

## **Z**

- ZFS zpools
  - configuration, 19–21
  - requirements, 20