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

Part No: F20515
June 2019

ORACLE

Oracle Solaris Cluster Data Replication Guide for Hitachi TrueCopy and Universal Replicator

Part No: F20515

Access to Oracle Support

Oracle customers that have purchased support 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.

Référence: F20515

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 stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, 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 livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel 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 un risque de dommages corporels. Si vous utilisez ce logiciel ou ce 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 des applications dangereuses.

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 de 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle. 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` ou le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` si vous êtes malentendant.

Contents

Using This Documentation

- **Overview** – Provides procedures for administering Hitachi TrueCopy and Universal Replicator replication with Oracle Solaris Cluster Geographic Edition (Disaster Recovery) software
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Disaster Recovery software, and expertise with the volume manager software that is used with Disaster Recovery software

Product Documentation Library

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

Feedback

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

♦ ♦ ♦ CHAPTER 1

Replicating Data With Hitachi TrueCopy and Universal Replicator Software

This chapter contains the procedures for configuring and administering data replication 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.

The Disaster Recovery framework 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 or 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 or Universal Replicator product documentation.

The chapter contains the following sections:

For information about adding and removing data replication components, see "Adding a Data Replication Component to a Hitachi TrueCopy or Universal Replicator Protection Group" on page 45. 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 68.

Replicating Data in a Hitachi TrueCopy or Universal Replicator Protection Group (Task Map)

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

TABLE 1 Configuration 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 "Planning and Configuring Hitachi TrueCopy or Universal Replicator Data Replication" 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 32 or "How to Create a Protection Group for Oracle Real Application Clusters" on page 39.
Add a data replication component that is controlled by Hitachi TrueCopy or Universal Replicator.	See "How to Add a Data Replication Component to a Hitachi TrueCopy or Universal Replicator Protection Group" on page 46.
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 43.
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 57.
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 62.
Check the runtime status of replication.	See "Checking the Runtime Status of Hitachi TrueCopy and Universal Replicator Data Replication" on page 68.

TABLE 2 Administration Tasks for Hitachi TrueCopy and Universal Replicator Data Replication

Task	Description
Detect failure.	See "Detecting Cluster Failure" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.
Migrate services by using a switchover.	See "Migrating Replication Services by Switching Over Protection Groups" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

Task	Description
Migrate services by using a takeover.	See "Forcing a Takeover of a Protection Group" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.
Recover data after forcing a takeover.	See "Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy or Universal Replicator Replication" on page 71.
Detect and recover from a data replication error.	See "Recovering From a Hitachi TrueCopy or Universal Replicator Data Replication Error" on page 82.

Planning and Configuring Hitachi TrueCopy or Universal Replicator Data Replication

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

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 Disaster Recovery partnership, you must have Hitachi storage arrays that are supported for use with Hitachi Universal Replicator. See the Oracle Solaris Cluster Interoperability Matrix for Hitachi Data Systems Enterprise Storage (`http://www.oracle.com/technetwork/server-storage/solaris-cluster/hds-enterprise-matrix-4x-1726954.pdf`) 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 data replication component 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 data replication components to the protection group.
  3. Start the protection group.

  For details, see "Ensuring Data Consistency for Hitachi Universal Replicator in Asynchronous Mode" on page 34.

Overview of Initial Configuration Process

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

- Configuring a Hitachi TrueCopy or Universal Replicator data replication component, `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 `HAStoragePlus` 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 Disaster Recovery framework does not support using a Hitachi TrueCopy or Universal Replicator S-VOL and Command Device as an Oracle Solaris Cluster quorum device.
- If you use the Hitachi TrueCopy and Universal Replicator Command Control Interface (CCI) for data replication, you must use RAID Manager. For information about which version you should use, see *Installing and Configuring the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.
- All Hitachi TrueCopy or Universal Replicator data replication components with the same consistency group ID (CTGID) must be added to the same protection group.
- The Disaster Recovery framework uses the default CCI instance to manage the Hitachi TrueCopy or Universal Replicator devices. The Disaster Recovery framework starts the default CCI instance whenever a Hitachi TrueCopy data replication component is managed by the Disaster Recovery framework. Applications that are not under the control of the Disaster Recovery framework can also use the default CCI instance or any other instances without risk to Disaster Recovery or application processes or data.
- The Disaster Recovery framework supports the hardware configurations that are supported by the Oracle Solaris Cluster software. See the Oracle Solaris Cluster Interoperability Matrix for Hitachi Data Systems Enterprise Storage (`http://www.oracle.com/technetwork/server-storage/solaris-cluster/hds-enterprise-matrix-4x-1726954.pdf`) 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.

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 Disaster Recovery framework.

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 Disaster Recovery configuration, a Hitachi TrueCopy or Universal Replicator data replication component 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 Disaster Recovery 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 data replication component in the /etc/horcm.conf file. Additionally, in this file, each data replication component is assigned a name that is the same on all storage-attached nodes of all clusters that are participating in a Disaster Recovery partnership.

In the /etc/horcm.conf file, you configure each Hitachi TrueCopy or Universal Replicator data replication component as a property of either the HORCM_DEV parameter or the HORCM_LDEV parameter. Depending on their intended use, you might configure one data replication component in the /etc/horcm.conf file as a property of HORCM_DEV and another data replication component as a property of HORCM_LDEV. However, a single data replication component can only be configured as a property of HORCM_DEV or of HORCM_LDEV. For any one data replication component, the selected parameter, HORCM_DEV or HORCM_LDEV, must be consistent on all storage-attached nodes of all clusters that are participating in the Disaster Recovery 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 Disaster Recovery configuration. For information about configuring other parameters in the /etc/horcm.conf file, see the documentation for Hitachi TrueCopy or Universal Replicator.

Journal Volumes

Entries in the /etc/horcm.conf file for Hitachi Universal Replicator data replication components 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 Disaster Recovery partnership temporarily fails. A journal volume can be used by more than one data replication component on the local cluster, but typically is assigned to just one data replication component. Hitachi TrueCopy does not support journalling.

If you want to implement journalling, you must configure Hitachi Universal Replicator data replication components as properties of the `HORCM_LDEV` parameter because only that parameter supports the association of data LUNs with journal volumes in the Disaster Recovery Hitachi Universal Replicator module. If you configure Hitachi Universal Replicator data replication components by using the `HORCM_DEV` parameter, no journalling 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 data replication components 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 Disaster Recovery framework is started. If you change the `/etc/horcm.conf` file on any node after the Disaster Recovery framework is started, and you do not anticipate rebooting, you must restart the `HORCM` daemon on that node by using the following commands:

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

Table 3, “Example `HORCM_LDEV` Section of the `/etc/horcm.conf` File on the Primary Cluster,” on page 18 shows the configuration of one journalling Hitachi Universal Replicator data replication component in the `/etc/horcm.conf` file as a property of the `HORCM_LDEV` parameter. Each LUN in the data replication component is described on a single line consisting of four space-delimited entries. The LUNs in the `devgroup1` data replication component are named `pair1` and `pair2`. The administrator chooses the data replication component 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 TrueCopy and Universal Replicator 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 3, "Example `HORCM_LDEV` Section of the `/etc/horcm.conf` File on the Primary Cluster," on page 18 with Hitachi Universal Replicator, as Hitachi TrueCopy does not support journalling.

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 data replication component. Otherwise, journalling does not occur and Hitachi Universal Replicator's functionality in Geographic Edition configurations is no greater than the functionality of Hitachi TrueCopy.

TABLE 3 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 4, "Example `HORCM_DEV` Section of the `/etc/horcm.conf` File on the Primary Cluster," on page 18 shows the configuration of one non-journalling Hitachi TrueCopy or Universal Replicator data replication component in the `/etc/horcm.conf` file as a property of the `HORCM_DEV` parameter. Each LUN in the data replication component is described on a single line consisting of five space-delimited entries. The table describes a data replication component 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 4 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 Disaster Recovery Systems

Disaster Recovery 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 Disaster Recovery.

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 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 Ctgid=5 -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 25 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 Disaster Recovery 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 storage pool (`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 22.

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

Observe the following requirements and restrictions for ZFS:

- ZFS is not supported with Hitachi TrueCopy. Use ZFS only with Hitachi Universal Replicator.
- 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 `zpool` are supported.
- ZFS `zpool` spares are not supported with storage-based replication in a Disaster Recovery 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.

- 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 Disaster Recovery software does not support Hitachi Universal Replicator S-VOL and Command Device as an Oracle Solaris Cluster quorum device.

1. Create a ZFS `zpool`.

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

`create` *appdataz*

Specifies the name of the `zpool` to create.

`mirror` c*N*t*X*d*Y* c*N*t*A*d*B*

Specifies the LUNs to replicate with Hitachi Universal Replicator.

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

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

`-g` *app-rg*

Specifies the application resource group.

`-p zpools=`*appdataz*

Specifies the `zpool`.

hasp4appdataz

Specifies the name of the HAStoragePlus resource to create.

Example 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 c6t600604800001879000235359 4D313137d0
 c6t600604800001879000235359 4D313143d0
```

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

```
# clresource create -g app-rg1 \
-t HAStoragePlus \
-p zpools=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 storage pool and you are using Hitachi Universal Replicator replication, 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 20.

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 Disaster Recovery framework does not support Hitachi TrueCopy or Universal Replicator S-VOL and Command Device as an Oracle Solaris Cluster quorum device.

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 Disaster Recovery framework handle mounting the file system by using the `HAStoragePlus` 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 `HAStoragePlus` 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 `HAStoragePlus` resource type, refer to the *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.

Example 3 Configuring a Highly Available Local File System

This example assumes that the `apprg1` resource group and the `oradg1` device group that uses DID `d4` already exist.

1. Create a UFS file system.

```
phys-paris-1# newfs /dev/global/rdsk/d4s0
```

2. Update the `/etc/vfstab` file on all nodes.

```
# echo "/dev/global/dsk/d4s0 /dev/global/rdsk/d4s0 /mounts/sample ufs 2 no logging"
 >> /etc/vfstab
```

3. Add the `HAStoragePlus` resource type.

```
phys-paris-1# clresource create -g apprg1 -t SUNW.HAStoragePlus \
-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 Disaster Recovery framework.

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 data replication component 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 5, "Example `HORCM_LDEV` Section of the `/etc/horcm.conf` File on the Secondary Cluster," on page 24 and Table 6, "Example `HORCM_DEV` Section of the `/etc/horcm.conf` File on the Secondary Cluster," on page 24 show the entries in the `/etc/horcm.conf` file on the nodes of the secondary cluster for the data replication components configured on the primary cluster in "Configuring the `/etc/horcm.conf` File on the Nodes of the Primary Cluster" on page 17. Table 5, "Example `HORCM_LDEV` Section of the `/etc/horcm.conf` File on the Secondary Cluster," on page 24 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 data replication component. Otherwise, journalling does not occur and Hitachi Universal Replicator's functionality in Geographic Edition configurations is no greater than the functionality of Hitachi TrueCopy.

TABLE 5 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 6 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 If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Disaster Recovery framework does not support Hitachi TrueCopy or Universal Replicator S-VOL and Command Device as an Oracle Solaris Cluster quorum device.

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/rdsk/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 `HAStoragePlus` resource in `apprg1`.**

    ```
    phys-newyork-1# clresource create -g apprg1 -t SUNW.HAStoragePlus \
    -p FilesystemMountPoints=/mounts/sample -p Affinityon=TRUE \
    -p GlobalDevicePaths=rawdg1 rs-hasp
    ```

    This `HAStoragePlus` resource is required for Disaster Recovery configurations, because the framework 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
phs-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.

Creating and Validating a Hitachi TrueCopy or Universal Replicator Protection Group

This section contains the following topics:

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 28
- "Creating a Protection Group While the Application Is Online" on page 29

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 32 or "How to Create a Protection Group for Oracle Real Application Clusters" on page 39.
- Add the data replication component to the protection group.

  For more information, see "How to Add a Data Replication Component to a Hitachi TrueCopy or Universal Replicator Protection Group" on page 46.
- 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 43.
- 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 57.
- 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 62.

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 32 or "How to Create a Protection Group for Oracle Real Application Clusters" on page 39.
- Add the data replication component to the protection group.

  For more information, see "How to Add a Data Replication Component to a Hitachi TrueCopy or Universal Replicator Protection Group" on page 46.
- Start the protection group locally.

  For more information, see "How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group" on page 62.
- 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 43.

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 57.
- Activate the protection group locally.

  For more information, see "How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group" on page 62.

EXAMPLE 4 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 \
```

```
-p Ctgid=5 -o Primary -s paris-newyork-ps hdspg
Protection group "hdspg" has been successfully created
```

2. Add the data replication component, `hdsdg`, to the protection group.

```
phys-paris-1# geopg add-replication-comonent -p Fence_level=async hdsdg hdspg
```

3. Activate the protection group locally.

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

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

```
phys-paris-1# geopg add-resource-group apprg1 hdspg
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 "hdspg"           : 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 hdspg
Protection group "hdspg" has been successfully created.
```

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

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

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

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

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

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

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

   Protection group "hdspg"               : 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 and Validating a Hitachi TrueCopy or Universal Replicator Protection Group

This section contains procedures for the following tasks:

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 4.4 Geographic Edition Installation and Configuration 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, see "How to Create a Protection Group for Oracle Real Application Clusters" on page 39.

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 Disaster Recovery 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 57.

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

   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

   `#` `geopg create -s` *partnership* **`-o`** *local-role* **`-d truecopy [-p`** *property* **`[-p…]]`** `\`
   protection-group

   `-s` *partnership* — Specifies the name of the partnership.

   `-o` *local-role* — Specifies the role of this protection group on the local cluster as either `primary` or `secondary`.

   `-d truecopy` — Specifies that the protection group data is replicated by the Hitachi TrueCopy or Universal Replicator software.

   `-p` *property* — 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` – Optional. Specifies the Oracle Solaris Cluster device group where the data is written. The LUNs in this device group must correspond to the LUNs that are replicated in the Hitachi TrueCopy or Universal Replicator data replication component that you added to the protection group. Setting this property ensures that the Disaster Recovery framework takes offline the specified device group during switchovers and takeovers.

For more information about the properties you can set, see Appendix A, "Standard Disaster Recovery Framework Properties," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

protection-group Specifies the name of the protection group.

For information about the names and values that are supported by the Disaster Recovery framework, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

For more information about the `geopg` command, refer to the `geopg(8)` man page.

Example 5 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 Ctgid=5 -p Nodelist=phys-paris-1,phys-paris-2 hdspg
```

Example 6 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, `hdspg`, 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 Ctgid=5 -p Nodelist=phys-paris-1,phys-paris-2 hdspg
   ```

2. Activate the protection group.

   ```
   # geopg start -e local hdspg
   ```

3. Add the application resource group.

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

Ensuring Data Consistency for Hitachi Universal Replicator in Asynchronous Mode

This section describes the protection group configuration that is required in Disaster Recovery 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 Disaster Recovery

The Disaster Recovery module supports Hitachi TrueCopy and Universal Replicator data replication components 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 Disaster Recovery with Hitachi Universal Replicator, journal volumes are associated with application storage in the `/etc/horcm.conf` file. That configuration is described in "Journal Volumes" on page 16 and "Configuring the `/etc/horcm.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 data replication components. A CTGID is an integer that is assigned to one or more Hitachi Universal Replicator data replication components. 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 Universal Replicator data replication component or several Hitachi Universal Replicator data replication components. For example, if Hitachi Universal Replicator stops replication on the devices of one data replication component that is assigned the CTGID of `5`, it stops replication on all other devices in data replication components with the CTGID of `5`.

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

To ensure this correspondence, the Disaster Recovery software allows the administrator to set a CTGID property on each protection group. The data replication components that are added to the protection group must all have the same CTGID as the protection group. If other data replication components are assigned the same CTGID as the data replication components in the protection group, the Disaster Recovery framework generates an error. For example, if the protection group `app1-pg` has been assigned the CTGID of `5`, all data replication components included in `app1-pg` must have the CTGID of `5`. Moreover, all CTGIDs of data replication components 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 Universal Replicator storage software will automatically assign a unique CTGID to an asynchronously replicated data replication component when it is initialized. Thereafter, the pairs in that data replication component 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 Universal Replicator data replication component, you can let the Hitachi Universal Replicator storage array assign the data replication component'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 Universal Replicator data replication components, 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 data replication component 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 data replication component consistency.

▼ Configuring Consistency Group IDs for Hitachi Universal Replicator Data Replication Components 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 data replication components in one of two ways:

- You can add uninitialized data replication components 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 data replication component with the CTGID that you plan to use for the protection group that will hold that data replication component. After you create the protection group with that CTGID, you must assign the data replication component 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` data replication component, which uses the `async` fence level.

Before You Begin

- Configure a Hitachi Universal Replicator data replication component 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 as described in "How to Set Up Raw-Disk Device Groups for Disaster Recovery Systems" on page 19.
- Configure an Oracle Solaris Cluster resource group that includes a resource of type HAStoragePlus in addition to any other resources that are required for its application data service. This HAStoragePlus resource must use the disk devices of a previously configured Hitachi Universal Replicator data replication component as described in "How to Configure a Highly Available Local File System for Hitachi TrueCopy or Universal Replicator Replication" on page 22.

1. **On the primary cluster, create the Disaster Recovery 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 data replication components to the protection group by using one of the following methods:**

   - Add data replication components 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-replication-component -p Fence_level=async devgroup1 app1-pg
     ```

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

     In the following example, a data replication component 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-replication-component -p Fence_level=async devgroup1 app1-pg
     ```

3. **Start the protection group.**

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

   Uninitialized data replication components, if any, are initialized and assigned the CTGID of 5.

Requirements to Support Oracle Real Application Clusters With Data Replication Software

The Disaster Recovery framework supports Oracle Real Application Clusters with Hitachi TrueCopy and Universal Replicator software. Observe the following requirements when you configure Oracle Real Application Clusters:

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

- You must create storage resources for dynamic database files that are replicated in their own resource groups. These storage resources must be separate from the resource group that holds the storage resource for Oracle Clusterware.
- To be able to leave Oracle RAC infrastructure resource groups outside of Disaster Recovery control, you must run Disaster Recovery binaries on both cluster partners and set the Oracle RAC protection group `External_Dependency_Allowed` property to `true`.
- Do not add the Oracle Clusterware OCR and Voting Disk device group to the protection group's `Cluster_dgs` property.
- Do not add Oracle RAC infrastructure resource groups to the protection group. Only add the `rac_server_proxy` resource group and resource groups for device groups that are replicated to the protection group. Also, you must set to `False` the `Auto_start_on_new_cluster` resource group property for the `rac_server_proxy` resource group and resource groups and for device groups that are replicated.

▼ 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 38.
- 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 Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

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

-s *partnership* Specifies the name of the partnership.

-o *local-role* Specifies the role of this protection group on the local cluster as `primary`.

`-d truecopy` — Specifies that the protection group data is replicated by the Hitachi TrueCopy or Universal Replicator software.

`-p` *property* — 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.
- `Cluster_dgs` – Optional. Specifies the Oracle Solaris Cluster device group where the replicated data is written. Specify this property if you want Disaster Recovery to unmount the file systems on this device group and take offline the device group. Do not specify OCR or Voting Disk device groups to this property.

For more information about the properties you can set, see Appendix A, "Standard Disaster Recovery Framework Properties," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

protection-group — Specifies the name of the protection group.

For information about the names and values that are supported by the Disaster Recovery framework, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

For more information about the `geopg` command, refer to the `geopg`(8) man page.

3. Add a Hitachi TrueCopy or Universal Replicatordata replication component to the protection group.

```
# geopg add-replication-component [-p property [-p…]] protection-group
```

`-p` *property* — 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 Disaster Recovery Framework Properties," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

protection-group 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 resource-group protection-group
```

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

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

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

protection-group Specifies the name of the protection group.

Example 7 Creating a Protection Group for Oracle RAC

This example creates the protection group `pg1` with the `External_dependency_allowed` property set to `true`. The example adds a replication component to the protection group, adds resource groups that contain a RAC server proxy, and that contain resources for Oracle ASM device groups for Oracle Database files that use replicated LUNs. 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.

```
# geopg create -s pts1 -o PRIMARY -d truecopy \
-p External_Dependency_Allowed=true pg1
Protection group "pg1" successfully created.
```

2. Add the Hitachi TrueCopy or Universal Replicator data replication component `VG01` to protection group `pg1`.

```
# geopg add-replication-component --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 `asm-dg-dbfiles-rg`, `hasp4rac-rg`, and `scaldbdg-rg`, to the protection group.

```
# geopg add-resource-group rac_server_proxy-rg,asm-dg-dbfiles-rg,hasp4rac-rg,
scaldbdg-rg pg1
```

How the Data Replication Subsystem Validates the Data Replication Component

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.

An 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-`*protection-group*.

Caution - These automatically created replication resource groups are for Disaster Recovery internal implementation purposes only. Use caution when you modify these resource groups by using Oracle Solaris Cluster commands.

Adding an Application Resource Group to a Hitachi TrueCopy or Universal Replicator Protection Group

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

All the entities you configure for the application resource group on the primary cluster, such as 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.

▼ 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 Disaster Recovery 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 rights profile to complete this procedure. For more information, see Chapter 4, “Administering Rights Profiles” in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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 resource-group  protection-group
```

resource-group

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

protection-group

Specifies the name of the protection group.

For information about the names and values that are supported by Disaster Recovery software, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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 8 Adding an Application Resource Group to a Protection Group

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

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

Adding a Data Replication Component to a Hitachi TrueCopy or Universal Replicator Protection Group

This section provides the following information about adding Hitachi TrueCopy or Universal Replicator data replication components:

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 32.

▼ How to Add a Data Replication Component to a Hitachi TrueCopy or Universal Replicator Protection Group

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

   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

   ```
   # geopg add-replication-component -p property [-p…] replication-component protection-group
   ```

   -p *property*

   Specifies the properties of the data replication component.

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

   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 Disaster Recovery Framework Properties," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

replication-component

Specifies the name of the new data replication component.

protection-group

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

For information about the names and values that are supported by Disaster Recovery software, see Appendix B, "Legal Names and Values of Disaster Recovery Framework Entities," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

For more information about the `geopg` command, refer to the `geopg`(8) man page.

Example 9 Adding a Data Replication Component to a Hitachi Universal Replicator Protection Group

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

```
# geopg add-replication-component -p Fence_level=async devgroup1 hurpg
```

Validations Made by the Data Replication Subsystem

When the Hitachi TrueCopy or Universal Replicator data replication component, 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 data replication component.
- The Hitachi TrueCopy or Universal Replicator; data replication component properties that are specified in the `geopg add-replication-component` command are validated as described in the following table.

Hitachi TrueCopy or Universal Replicator Data Replication Component Property	Validation
data-replication-component	Checks that the specified Hitachi TrueCopy or Universal Replicator data replication component is configured on all of the cluster nodes that are specified in the `Nodelist` property.
`Fence_level`	If a pair is already established for this Hitachi TrueCopy or Universal Replicator data replication component, the data replication layer checks that the specified `Fence_level` matches the already established fence level. If a pair is not yet established, for example, if a pair is in the `SMPL` state, any `Fence_level` is accepted.

When a Hitachi TrueCopy or Universal Replicator data replication component 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-`*protection-group*`-`*replication-component*. This resource is placed in the corresponding Oracle Solaris Cluster resource group, which is named `rg-tc-`*protection-group*.

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 Data Replication Component Is Validated

For validation purposes, Disaster Recovery gives each Hitachi TrueCopy or Universal Replicator data replication component 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 data replication component 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 Data Replication Component

An individual Hitachi TrueCopy or Universal Replicator data replication component can be in one of the following states:

- `SMPL`
- `Regular Primary`
- `Regular Secondary`
- `Takeover Primary`
- `Takeover Secondary`

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

TABLE 7 Individual Hitachi TrueCopy and Universal Replicator Data Replication Component States

Output of `pairvolchk`	Individual Data Replication Component 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 data replication component category requires that the `horcmd` process be active on the remote cluster so that the `remote-pair-state` for this data replication component can be obtained.
	`Regular Primary`, if `remote-cluster-state !=SSWS`
	or
	`Takeover Secondary`, if `remote-cluster-state == SSWS`
	`SSWS`, when you use the `pairdisplay -g` *devicegroup* `-fc` command.

Output of `pairvolchk`	Individual Data Replication Component State

`32 / 52 = SVOL_COPY`
`33 / 53 = SVOL_PAIR`
`35 / 55 = SVOL_PSUE`
`36 / 56 = SVOL_PDUB`
`57 = SVOL_PFUL`
`58 = SVOL_PFUS`
	`Regular Secondary`

`34 / 54 = SVOL_PSUS`	`Regular Secondary`, if `local-cluster-state !=SSWS`
	or
	`Takeover Primary`, if `local-cluster-state == SSWS`
	`SSWS`, when you use the `pairdisplay -g` *devicegroup* `-fc` command.

Determining the Aggregate Hitachi TrueCopy or Universal Replicator Data Replication Component State

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

When a protection group contains multiple Hitachi TrueCopy or Universal Replicator data replication components, the aggregate data replication component state is obtained as described in the following table.

TABLE 8 Conditions That Determine the Aggregate Data Replication Component State

Condition	Aggregate Data Replication Component State
All individual data replication component states are `SMPL`	`SMPL`
All individual data replication component states are either `Regular Primary` or `SMPL`	`Regular Primary`
All individual data replication component states are either `Regular Secondary` or `SMPL`	`Regular Secondary`
All individual data replication component states are either `Takeover Primary` or `SMPL`	`Takeover Primary`
All individual data replication component states are either `Takeover Secondary` or `SMPL`	`Takeover Secondary`

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

Validating the Local Role of the Protection Group Against the Aggregate Data Replication Component State

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

TABLE 9 Validating the Aggregate Data Replication Component State Against the Local Role of a Protection Group

Aggregate Data Replication Component State	Valid Local Protection Group Role
SMPL	primary or secondary
Regular Primary	primary
Regular Secondary	secondary
Takeover Primary	primary
Takeover Secondary	secondary

EXAMPLE 10 Validating the Aggregate Data Replication Component State

This example validates the state of a Hitachi TrueCopy or Universal Replicator data replication component 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# geopg create -s paris-newyork-ps -o primary -d truecopy -p Ctgid=5 hdspg
```

A data replication component, `devgroup1`, is added to the protection group, `hdspg`, as follows:

```
phys-paris-1# geopg add-replication-component -p Fence_level=async devgroup1 hdspg
```

The current state of a Hitachi TrueCopy or Universal Replicator data replication component, `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
parivolchk : 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 7, “Individual Hitachi TrueCopy and Universal Replicator Data Replication Component States,” on page 49 to an individual data replication component state of `Regular Primary`. Because the protection group contains only one data replication component, the aggregate data replication component state is the same as the individual data replication component state. The data replication component state is valid because the local role of the protection group, specified by the `-o` option, is `primary`, as specified in Table 9, “Validating the Aggregate Data Replication Component State Against the Local Role of a Protection Group,” on page 51.

♦ ♦ ♦ CHAPTER 2

Administering Hitachi TrueCopy and Universal Replicator Protection Groups

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

Administering Hitachi TrueCopy and Universal Replicator Application Resource Groups

This section describes how to remove 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.

How to Remove an Application Resource Group From a Hitachi TrueCopy or Universal Replicator Protection 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 rights profile to complete this procedure. For more information, see Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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 resource-group protection-group
   ```

   resource-group

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

   protection-group

   Specifies the name of the protection group.

   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 11 Removing an Application Resource Group From a Protection Group

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

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

Administering Hitachi TrueCopy and Universal Replicator Data Replication Components

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

- "How to Modify a Hitachi TrueCopy or Universal Replicator Data Replication Component" on page 55
- "How to Remove a Data Replication Component From a Hitachi TrueCopy or Universal Replicator Protection Group" on page 56

For details about configuring a Hitachi TrueCopy and 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 32.

▼ How to Modify a Hitachi TrueCopy or Universal Replicator Data Replication Component

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

   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

2. **Modify the data replication component.**

   This command modifies the properties of a data replication component in a protection group on the local cluster. Then the command propagates the new configuration to the partner cluster if the partner cluster contains a protection group of the same name.

   `# geopg modify-replication-component -p` *property* **`[-p…]`** *replication-component protection-group*

   -p *property*

   Specifies the properties of the data replication component.

   For more information about the properties you can set, see Appendix A, "Standard Disaster Recovery Framework Properties," in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

   replication-component

   Specifies the name of the new data replication component.

protection-group

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

Example 12 Modifying the Properties of a Hitachi TrueCopy or Universal Replicator Data Replication Component

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

```
# geopg modify-replication-component -p Fence_level=async hdsdg hdspg
```

How to Remove a Data Replication Component From a Hitachi TrueCopy or Universal Replicator Protection Group

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

Removing a data replication component does not stop replication or change the replication status of the data replication component.

For information about deleting protection groups, refer to "How to Delete a Protection Group" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*. For information about removing application resource groups from a protection group, refer to "How to Remove an Application Resource Group From a Hitachi TrueCopy or Universal Replicator Protection Group" on page 54.

1. **Log in to a cluster node.**
   You must be assigned the Geo Management rights profile to complete this procedure. For more information, see Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

2. **Remove the data replication component.**
   This command removes a data replication component 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-replication-component replication-component protection-group
   ```

replication-component

Specifies the name of the data replication component.

protection-group

Specifies the name of the protection group.

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

Example 13 Removing a Data Replication Component From a Hitachi TrueCopy or Universal Replicator Protection Group

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

```
# geopg remove-replication-component hdsdg hdspg
```

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 data replication components 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 Disaster Recovery framework 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 Disaster Recovery framework 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 rights profile to complete this procedure. For more information, see Chapter 4, “Administering Rights Profiles” in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

   `phys-newyork-1` is the only node on the secondary cluster. For a reminder of which node is `phys-newyork-1`, see “Example Disaster Recovery Framework Cluster Configuration” in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

   ```
   phys-newyork-1# geopg get -s partnership [protectiogroup]
   ```

   `-s` *partnership*

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

   protection-group

   Specifies the name of the protection group.

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

Note - The `geopg get` command replicates Disaster Recovery related entities. For information about how to replicate Oracle Solaris Cluster entities, see "How to Back Up the Cluster Configuration" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

Example 14 Replicating a Hitachi TrueCopy or Universal Replicator Protection Group to a Partner Cluster

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

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

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 32.

You can activate a protection group in the following ways:

- Globally – Activates a protection group on both clusters where the protection group is configured.
- Primary cluster only – Secondary cluster remains inactive.
- 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 data replication component state as described in Table 9, "Validating the Aggregate Data Replication Component State Against the Local Role of a Protection Group," on page 51. If validation is successful, data replication is started.
- Data replication is started on the data replication components 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 data replication component 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 data replication component is in `SMPL` state, then you create the data replication component with the `paircreate` 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 data replication component is in `SMPL` state then you create the data replication component with the `paircreate` command when you run the `geopg start` command with the `-fg`flags.

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 data replication components 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 data replication component name and `fl` is the fence level that is configured for the data replication component.

TABLE 10 Commands Used to Start Hitachi TrueCopy or Universal Replicator Data Replication

Aggregate Data Replication Component 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`

Aggregate Data Replication Component State	Valid Local Protection Group Role	Hitachi TrueCopy or Universal Replicator Start Command
		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.
`Takeover Secondary`	`secondary`	If the local state code is 24, 44, 25, or 45, the following command is run: `pairresync -swapp -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 running on the remote cluster.

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 Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

2. **Activate the protection group.**

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

   ```
   # geopg start -e scope [-n] protection-group
   ```

   -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.

   protection-group

   Specifies the name of the protection group.

   The `geopg start` command uses Oracle Solaris Cluster commands to bring resource groups and resources online.

Example 15 How the Disaster Recovery Software Issues the Command to Start Replication

This example illustrates how the Disaster Recovery 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 -p Ctgid=5 hdspg
```

A data replication component, `devgroup1`, is added to the protection group.

```
phys-paris-1# geopg add-replication-component -p Fence_level=async devgroup1 hdspg
```

The current state of a Hitachi TrueCopy or Universal Replicator data replication component, `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 data replication component state is `SMPL`.

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

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

The Disaster Recovery framework 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 16 Activating a Hitachi TrueCopy or Universal Replicator Protection Group Globally

This example activates a protection group globally.

```
# geopg start -e global hdspg
```

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

Example 17 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 hdspg
```

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 data replication component state as described in Table 9, “Validating the Aggregate Data Replication Component State Against the Local Role of a Protection Group,” on page 51. If validation is successful, data replication is stopped.
- Data replication is stopped on the data replication components 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 data replication component 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 data replication component name.

TABLE 11 Commands Used to Stop Hitachi TrueCopy or Universal Replicator Data Replication

Aggregate Data Replication Component 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: `pairsplit -g dg [-l]`. 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: `pairsplit -g dg`. 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 rights profile to complete this procedure. For more information, see Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

2. **Deactivate the protection group.**
   When you deactivate a protection group, its application resource groups are also unmanaged.

   # geopg stop -e *scope* **[-D]** *protection-group*

   -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.

protection-group

Specifies the name of the protection group.

Example 18 How the Disaster Recovery Software Issues the Command to Stop Replication

This example illustrates how the Disaster Recovery framework 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 data replication component, `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 data replication component, `devgroup1`, is added to the protection group as follows:

```
phys-paris-1# geopg add-replication-component -p Fence_level=async devgroup1 hdspg
```

The Disaster Recovery framework 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, `hdspg`, is deactivated by using the `geopg stop` command.

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

The Disaster Recovery framework 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 19 Deactivating a Protection Group on All Clusters

This example deactivates a protection group on all clusters.

```
# geopg stop -e global hdspg
```

Example 20 Deactivating a Protection Group on a Local Cluster

This example deactivates a protection group on the local cluster.

```
# geopg stop -e local hdspg
```

Example 21 Stopping Data Replication While Leaving the Protection Group Online

This example stops only data replication on a local cluster.

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

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 hdspg
```

Example 22 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, `hdspg`, on both clusters.

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

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

2. Deactivate the protection group.

   ```
   # geopg stop -e global hdspg
   ```

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.

Overview of 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 data replication component. The status of all the resources under a protection group are aggregated in the replication status. This replication status is the second component of the protection group state. For more information about the states of protection groups, refer to "Monitoring the Runtime Status of the Disaster Recovery Framework" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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 Chapter 4, "Administering Rights Profiles" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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 data replication component.**

```
# clresource status
```

Refer to the `Status` and `Status Message` fields for the data replication component you want to check.

See Also For more information about these fields, see Table 12, “Status and Status Messages of an Online Hitachi TrueCopy or Universal Replicator Replication Resource Group,” on page 70.

Hitachi TrueCopy or Universal Replicator Runtime Status and Status Messages

The Disaster Recovery framework 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_`*protection-group*

If you add a Hitachi TrueCopy or Universal Replicator data replication component to a protection group, Disaster Recovery framework creates a resource for each data replication component. This resource monitors the status of replication for its data replication component. The name of each resource has the following format:

`r-tc-`*protection-group*`-`*replication-component*

You can monitor the status of replication of this data replication component 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 12 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
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`(8CL) man page.

♦ ♦ ♦ CHAPTER 3

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:

Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy or Universal Replicator Replication

This section provides the following information:

Overview of Recovering Services After a Takeover

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*.

The Disaster Recovery framework 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 4.4 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.

▼ 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 Disaster Recovery infrastructure is enabled on the cluster. For more information about booting a cluster, see "How to Boot a Cluster" in *Administering an Oracle Solaris Cluster 4.4 Configuration*.
- 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 *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

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

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

d. Verify that the protection group is stopped.

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

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

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

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

protection-group

Specifies the name of the protection group.

For more information, see the `geopg`(8) man page.

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# geopg validate protection-group
   ```

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

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

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

   `-e local`

   Specifies the scope of the command.

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

   

   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 62.

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 data replication components on `cluster-newyork` have a state of `PVOL_PAIR` and the Hitachi TrueCopy or Universal Replicator data replication components 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.**

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

`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 component you want to check. For more information about these fields, see Table 7, “Individual Hitachi TrueCopy and Universal Replicator Data Replication Component States,” on page 49.

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 68.

▼ 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 `geopg 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 partnership
   ```

   partnership

   Specifies the name of the partnership.

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

   For more information about synchronizing partnerships, see "Resynchronizing a Partnership" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

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

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

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

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

protection-group

Specifies the name of the protection group.

d. Verify that the protection group is stopped.

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

e. Place the Hitachi TrueCopy or Universal Replicator data replication component, `devgroup1`, in the SMPL state.

Use the `pairsplit` commands to place the Hitachi TrueCopy or Universal Replicator data replication components 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 data replication component. The following table gives some examples of the command you need to use on `cluster-paris` for some typical pair states.

Pair State on `cluster-paris`	Pair State on `cluster-newyork`	`pairsplit` Command Used on `cluster-paris`

PSUS or PSUE	SSWS	`pairsplit -R -g dgname`
		
		`pairsplit -S -g dgname`

SSUS	PSUS	`pairsplit -S -g dgname`

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 protection-group
```

For more information, see the `geopg`(8) man page.

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 a error state.

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

   For more information, see the `geopg`(8) man page.

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 protection-group
   ```

   `-e local`

   Specifies the scope of the command.

   .

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

   `-n`

   Prevents the start of data replication at protection group startup.

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

   For more information, see "How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group" on page 62.

   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] protection-group
   ```

   `-f`

   Forces the command to perform the operation without your confirmation.

   The local state of the protection group on `cluster-newyork` is `Offline`.

   For more information about the `geopg takeover` command, see the `geopg`(8) man page.

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.

To start monitoring the local state of the protection group, you must activate the protection group on `cluster-newyork`.

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

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

`-e local`

Specifies the scope of the command.

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

`-n`

Prevents the start of data replication at protection group startup.

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

For more information about the `geopg start` command, see "How to Activate a Hitachi TrueCopy or Universal Replicator Protection Group" on page 62.

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 component you want to check. For more information about these fields, see Table 7, "Individual Hitachi TrueCopy and Universal Replicator Data Replication Component States," on page 49.

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 68.

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

This section describes the initial conditions that lead to a switchover failure and how to recover from a switchover failure.

Overview of Recovering Services After a Switchover

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. The Disaster Recovery framework does not start the applications on the new intended primary cluster in a switchover failure scenario.

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 hdspg
```

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 data replication component, `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:
  hdspg role: Secondary
cluster-newyork:
  hdspg 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 data replication component, `devgroup1`, in the SMPL state.

   Use the `pairsplit` commands to place the data replication components 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. Follow procedures in "Migrating Replication Services by Switching Over Protection Groups" in *Administering the Disaster Recovery Framework for Oracle Solaris Cluster 4.4*.

   - 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`.

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 data replication component.

This section provides the following information:

- "How to Detect Data Replication Errors" on page 82
- "How to Recover From a Hitachi TrueCopy or Universal Replicator Data Replication Error" on page 84

▼ How to Detect Data Replication Errors

Before You Begin For information about how different `Resource status` values map to actual replication pair states, see Table 12, "Status and Status Messages of an Online Hitachi TrueCopy or Universal Replicator Replication Resource Group," on page 70.

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

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

Output is similar to the following:

```
=== Cluster Resources ===

Resource Name          de Name        State      Status Message
```

```
-------------           ---------    -----      --------------
r-tc-hdspg1-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
```

2. Display the aggregate resource status for all data replication components in the protection group.

For example, the output of the `clresource status` command in the preceding example indicates that the Hitachi TrueCopy or Universal Replicator data replication component, `devgroup1`, is in the `PSUE` state on `cluster-paris`. Table 12, “Status and Status Messages of an Online Hitachi TrueCopy or Universal Replicator Replication Resource Group,” on page 70 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 "hdspg"   : 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         : "hdspg"
      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 data replication component, this state is automatically detected by the resource and is reported as a new protection group state.

3. **Revalidate the protection group configuration.**

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

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

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

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

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

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

♦ ♦ ♦ APPENDIX A

Geographic Edition Properties for Hitachi TrueCopy and Universal Replicator

This appendix provides the properties of Disaster Recovery data replication components.

This appendix contains the following sections:

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 Disaster Recovery framework defines.

TABLE 13 Hitachi TrueCopy and Universal Replicator Properties

Property	Description
Data Replication Property: `Cluster_dgs` (string array)	Lists the Oracle Solaris Cluster device groups where the data is written. The list is comma delimited. Only applications that belong to the protection group should write to these device groups. The Oracle Solaris Cluster device groups listed in the `Cluster_dgs` protection group property must exist and have the same name on both the primary cluster and the secondary cluster. The LUNs in this device group must correspond to LUNs getting replicated in the Hitachi data replication component added to the protection group. Specify this property if you want geographic edition to unmount file systems that use this Oracle Solaris Cluster device group. Tuning recommendations: This property can only be tuned when the protection group is offline. Category: Optional

Property	Description
	Default: Empty
Data Replication Property: `Nodelist` (string array)	Lists the host names of the machines that can be primary for the replication mechanism. This list is comma delimited. Tuning recommendations: This property can be tuned at any time. Category: Optional Default: All nodes in the cluster
Data Replication Property: `Fence_level` (enum)	Defines the fence level that is used by the data replication component. The fence level determines the level of consistency among the primary and secondary volumes for that data replication component. Possible values are `Never` and `Async`. To use the `data` or `status` fence levels, contact your Oracle representative. **Note -** If you specify a `Fence_level` of `never`, the data replication roles do not change after you perform a takeover. For more information about setting this property, see "How to Create a Protection Group for Oracle Real Application Clusters" on page 39. Tuning recommendations: This property can only be tuned when the protection group is offline. Category: Required Default: None
Data Replication Property: `Ctgid` (integer)	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 data replication components thereafter added to the protection group either must be uninitialized or must already have the same CTGID as the protection group. Attempting to add an initialized data replication component to a protection group results in an error if the CTGID of the data replication component differs from the CTGID of the protection group. A data replication component with the same CTGID as a protection group must be added to that protection group. Tuning recommendations: This property can only be tuned at creation. Category: Optional Default: None

Hitachi TrueCopy and Universal Replicator Properties That Must Not Be Changed

The Disaster Recovery framework 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 data replication component that contains the volumes that are being replicated.
- `Replication_role` – Defines the local data replication role.

Index

R

S

T

V

Z