



# **Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite**



Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054  
U.S.A.

Part No: 819-7196-10  
December 2006

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

# Contents

---

<b>Preface</b> .....	7
<b>1 Replicating Data With Sun StorageTek Availability Suite Software</b> .....	11
Task Summary of Replicating Data in a Sun StorageTek Availability Suite Protection Group .	11
Overview of Sun StorageTek Availability Suite Data Replication .....	12
Sun StorageTek Availability Suite Lightweight Resource Groups .....	13
Sun StorageTek Availability Suite Replication Resource Groups .....	13
Initial Configuration of Sun StorageTek Availability Suite Software .....	14
Sun StorageTek Availability Suite Volume Sets .....	14
▼ How to Set Up Raw-Disk Device Groups for Sun Cluster Geographic Edition Systems ....	16
▼ How to Configure a Sun StorageTek Availability Suite Volume in Sun Cluster .....	17
Enabling a Sun StorageTek Availability Suite Volume Set .....	18
▼ How to Configure the Sun Cluster Device Group That Is Controlled by Sun StorageTek Availability Suite .....	22
▼ How to Configure a Highly Available Cluster Global File System for Use With Sun StorageTek Availability Suite .....	23
<b>2 Administering Sun StorageTek Availability Suite Protection Groups</b> .....	25
Strategies for Creating Sun StorageTek Availability Suite Protection Groups .....	25
Creating a Protection Group While the Application Is Offline .....	26
Creating a Protection Group While the Application Is Online .....	26
Creating, Modifying, Validating, and Deleting a Sun StorageTek Availability Suite Protection Group .....	34
▼ How to Create and Configure a Sun StorageTek Availability Suite Protection Group .....	34
▼ How to Modify a Sun StorageTek Availability Suite Protection Group .....	36
▼ How to Validate a Sun StorageTek Availability Suite Protection Group .....	37
How the Data Replication Layer Validates the Application Resource Groups and Data Replication Entities .....	38

- ▼ How to Delete a Sun StorageTek Availability Suite Protection Group ..... 39
- Administering Sun StorageTek Availability Suite Application Resource Groups ..... 40
  - ▼ How to Add an Application Resource Group to a Sun StorageTek Availability Suite Protection Group ..... 41
  - ▼ How to Delete an Application Resource Group From a Sun StorageTek Availability Suite Protection Group ..... 43
- Administering Sun StorageTek Availability Suite Data Replication Device Groups ..... 44
  - ▼ How to Add a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group ..... 45
    - How the Data Replication Subsystem Verifies the Device Group ..... 47
  - ▼ How to Modify a Sun StorageTek Availability Suite Data Replication Device Group ..... 48
  - ▼ How to Delete a Data Replication Device Group From a Sun StorageTek Availability Suite Protection Group ..... 49
- Replicating the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster ..... 50
  - ▼ How to Replicate the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster ..... 50
- Activating and Deactivating a Protection Group ..... 52
  - ▼ How to Activate a Sun StorageTek Availability Suite Protection Group ..... 52
  - ▼ How to Deactivate a Sun StorageTek Availability Suite Protection Group ..... 54
- Resynchronizing a Sun StorageTek Availability Suite Protection Group ..... 57
  - ▼ How to Resynchronize a Sun StorageTek Availability Suite Protection Group ..... 57
- Checking the Runtime Status of Sun StorageTek Availability Suite Data Replication ..... 58
  - Displaying a Sun StorageTek Availability Suite Runtime Status Overview ..... 58
    - ▼ How to Check the Overall Runtime Status of Replication ..... 58
  - Displaying a Detailed Sun StorageTek Availability Suite Runtime Status ..... 59
- 3 Migrating Services That Use Sun StorageTek Availability Suite Data Replication ..... 61**
  - Detecting Cluster Failure on a System That Uses Sun StorageTek Availability Suite Data Replication ..... 61
    - Detecting Primary Cluster Failure ..... 61
    - Detecting Secondary Cluster Failure ..... 62
  - Migrating Services That Use Sun StorageTek Availability Suite With a Switchover ..... 62
    - ▼ How to Switch Over a Sun StorageTek Availability Suite Protection Group From Primary to Secondary ..... 63
      - Actions Performed by the Sun Cluster Geographic Edition Software During a Switchover ..... 64

Forcing a Takeover on Systems That Use Sun StorageTek Availability Suite .....	65
▼ How to Force Immediate Takeover of Sun StorageTek Availability Suite Services by a Secondary Cluster .....	66
Actions Performed by the Sun Cluster Geographic Edition Software During a Takeover .	67
Recovering Sun StorageTek Availability Suite Data After a Takeover .....	68
▼ How to Resynchronize and Revalidate the Protection Group Configuration .....	69
▼ How to Perform a Failback-Switchover on a System That Uses Sun StorageTek Availability Suite Replication .....	71
▼ How to Perform a Failback-Takeover on a System That Uses Sun StorageTek Availability Suite Replication .....	74
Recovering From a Sun StorageTek Availability Suite Data Replication Error .....	77
▼ How to Recover From a Data Replication Error .....	78
<b>A Sun Cluster Geographic Edition Properties for Sun StorageTek Availability Suite .....</b>	<b>79</b>
Sun StorageTek Availability Suite Properties .....	79
Sun Cluster Geographic Edition Resource Properties for Sun StorageTek Availability Suite That Must Not Be Changed .....	80
<b>Index .....</b>	<b>81</b>



# Preface

---

*Sun Cluster Geographic Edition Data Replication Guide for Sun StorEdge Availability Suite* provides procedures for administering Sun StorageTek Availability Suite data replication with Sun™ Cluster Geographic Edition software. This document is intended for experienced system administrators with extensive knowledge of Sun software and hardware. This document is not to be used as a planning or presales guide.

The instructions in this book assume knowledge of the Solaris™ Operating System (Solaris OS) and expertise with Sun Cluster software and with the volume manager software that is used with Sun Cluster software.

---

## Note –

This release of Sun Cluster Geographic Edition supports the following releases of availability suite software:

- Sun StorageTek Availability Suite 4
- Sun StorEdge Availability Suite 3.2.1

In this manual, references to Sun StorageTek Availability Suite software also apply to Sun StorEdge Availability Suite software unless specifically stated otherwise.

---

## Related Books

Information about related Sun Cluster Geographic Edition topics is available in the documentation that is listed in the following table. All Sun Cluster Geographic Edition documentation is available at <http://docs.sun.com>.

---

Topic	Documentation
Overview	<i>Sun Cluster Geographic Edition Overview</i>
Glossary	<i>Sun Java Enterprise System Glossary</i>
Hardware administration	Individual hardware administration guides

---

Topic	Documentation
Software installation	<i>Sun Cluster Geographic Edition Installation Guide</i>
System administration	<i>Sun Cluster Geographic Edition System Administration Guide</i> <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorageTek Availability Suite</i> <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>
Command and function references	<i>Sun Cluster Geographic Edition Reference Manual</i>

---

For a complete list of Sun Cluster documentation, see the release notes for your Sun Cluster software at <http://docs.sun.com>.

## Using UNIX Commands

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

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

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

## Related Third-Party Web Site References

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

---

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

---



## Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- Documentation (<http://www.sun.com/documentation/>)
- Support (<http://www.sun.com/support/>)
- Training (<http://www.sun.com/training/>)

## Typographic Conventions

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

TABLE P-1 Typographic Conventions

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

## Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

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

# Replicating Data With Sun StorageTek Availability Suite 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.

Sun Cluster Geographic Edition software supports the use of Sun StorageTek Availability Suite remote mirror software for data replication. Before you can replicate data with Sun StorageTek Availability Suite software, you must be familiar with the Sun StorageTek Availability Suite documentation, have the Sun StorageTek Availability Suite product, and have the latest Sun StorageTek Availability Suite patches installed on your system. For information about installing Sun StorageTek Availability Suite software and its latest patches, see *Sun StorageTek Availability Suite 4 Software Installation Guide*.

This chapter describes the procedures for configuring data replication with Sun StorageTek Availability Suite software. This chapter contains the following sections:

- “Task Summary of Replicating Data in a Sun StorageTek Availability Suite Protection Group” on page 11
- “Overview of Sun StorageTek Availability Suite Data Replication” on page 12
- “Initial Configuration of Sun StorageTek Availability Suite Software” on page 14

## Task Summary of Replicating Data in a Sun StorageTek Availability Suite Protection Group

This section summarizes the steps for configuring Sun StorageTek Availability Suite data replication in a protection group.

TABLE 1-1 Administration Tasks for Sun StorageTek Availability Suite Data Replication

Task	Description
Perform an initial configuration of the Sun StorageTek Availability Suite software.	See “Initial Configuration of Sun StorageTek Availability Suite Software” on page 14.
Create a protection group that is configured for Sun StorageTek Availability Suite data replication.	See “How to Create and Configure a Sun StorageTek Availability Suite Protection Group” on page 34.
Add a device group that is controlled by Sun StorageTek Availability Suite.	See “How to Add a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group” on page 45.
Add an application resource groups to the protection group.	See “How to Add an Application Resource Group to a Sun StorageTek Availability Suite Protection Group” on page 41.
Replicate the protection group configuration to a secondary cluster.	See “How to Replicate the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster” on page 50.
Activate the protection group.	See “How to Activate a Sun StorageTek Availability Suite Protection Group” on page 52.
Verify the protection group configuration.	Perform a trial switchover or takeover and test some simple failure scenarios before bringing your system online. See Chapter 3.
Check the runtime status of replication.	See “Checking the Runtime Status of Sun StorageTek Availability Suite Data Replication” on page 58.
Detect failure.	See “Detecting Cluster Failure on a System That Uses Sun StorageTek Availability Suite Data Replication” on page 61.
Migrate services by using a switchover.	See “Migrating Services That Use Sun StorageTek Availability Suite With a Switchover” on page 62.
Migrate services by using a takeover.	See “Forcing a Takeover on Systems That Use Sun StorageTek Availability Suite” on page 65.
Recover data after forcing a takeover.	See “Recovering Sun StorageTek Availability Suite Data After a Takeover” on page 68.

## Overview of Sun StorageTek Availability Suite Data Replication

This section provides an overview of Sun StorageTek Availability Suite resource groups and outlines some limitations of Sun StorageTek Availability Suite replication on clusters of more than two nodes.

## Sun StorageTek Availability Suite Lightweight Resource Groups

To a protection group, you can add a device group that is controlled by the Sun StorageTek Availability Suite software. The Sun Cluster Geographic Edition software creates a lightweight resource group for each device group. The name of a lightweight resource group has the following format:

*AVSdevicegroupname-stor-rg*

For example, a device group named `avsdg` that is controlled by the Sun StorageTek Availability Suite software has a lightweight resource group named `avsdg-stor-rg`.

The lightweight resource group collocates the logical host and the device group, a requirement of data replication with the Sun StorageTek Availability Suite remote mirror software.

Each lightweight resource group contains two resources:

- A logical hostname resource for the local logical host that is used for replication of the device group. The name of this resource has the following format:

*SSEdevicegroup-lh*

- An HASoragePlus resource for controlling the collocation of the device group with the lightweight resource group. The name of this resource has the format *AVSdevicegroupname-stor*.

For more information about lightweight resource groups, see the Sun StorageTek Availability Suite documentation.

## Sun StorageTek Availability Suite Replication Resource Groups

When a device group that is controlled by the Sun StorageTek Availability Suite software is added to a protection group, the Sun Cluster Geographic Edition software creates a special replication resource for that device group in the replication resource group. By monitoring these replication resource groups, the Sun Cluster Geographic Edition software monitors the overall status of replication. One replication resource group with one replication resource is created for each protection group.

The name of the replication resource group has the following format:

*AVSprotectiongroupname-rep-rg*

The replication resource in the replication resource group monitors the replication status of the device group on the local cluster, which is reported by the Sun StorageTek Availability Suite remote mirror software.

The name of a replication resource has the following format:

*AVSdevicegroupname*-rep-rs

## Initial Configuration of Sun StorageTek Availability Suite Software

This section describes the initial steps you must perform before you can configure Sun StorageTek Availability Suite replication in the Sun Cluster Geographic Edition product.

The example protection group, `avspg`, in this section has been configured on a partnership that consists of two clusters, `cluster-paris` and `cluster-newyork`. An application, which is encapsulated in the `apprg1` resource group, is protected by the `avspg` protection group. The application data is contained in the `avsdg` device group. The volumes in the `avsdg` device group can be Solaris Volume Manager volumes, VERITAS Volume Manager volumes, or raw device volumes.

The resource group, `apprg1`, and the device group, `avsdg`, are present on both the `cluster-paris` cluster and the `cluster-newyork` cluster. The `avspg` protection group protects the application data by replicating data between the `cluster-paris` cluster and the `cluster-newyork` cluster.

---

**Note** – Replication of each device group requires a logical host on the local cluster and a logical host on the partner cluster.

---

You cannot use the slash character (/) in a cluster tag in the Sun Cluster Geographic Edition software. If you are using raw DID devices, you cannot use predefined DID device group names such as `dsk/s3`.

To use DIDs with raw device groups, see [“How to Set Up Raw-Disk Device Groups for Sun Cluster Geographic Edition Systems”](#) on page 16.

## Sun StorageTek Availability Suite Volume Sets

Before you can define a Sun StorageTek Availability Suite volume set, you must determine the following:

- **The data volumes to replicate** such as `vol-data-paris` in `avsdg` on `cluster-paris` and `vol-data-newyork` in `avsdg` on `cluster-newyork`.
- **The bitmap volume that is needed for replication**, such as `vol-bitmap-paris` in `avsdg` on `cluster-paris` and `vol-bitmap-newyork` in `avsdg` on `cluster-newyork`.

- **The logical host to use exclusively for replication of the device group** avsdg, such as the logical host logicalhost-paris-1 on cluster-paris and the logical host logicalhost-newyork-1 on cluster-newyork.

---

**Note** – The logical host that is used for Sun StorageTek Availability Suite replication must be different from the Sun Cluster Geographic Edition infrastructure logical host. For more information, see “Configuring Logical Hostnames” in *Sun Cluster Geographic Edition System Administration Guide* about configuring logical hostnames.

---

The volset file is located at `/var/cluster/geo/avs/devicegroupname-volset.ini` on all nodes of the primary and secondary clusters of the protection group. For example, the volset file for the device group avsdg is located at `/var/cluster/geo/avs/avsdg-volset.ini`.

The fields in the volume set file that are handled by the Sun Cluster Geographic Edition software are described in the following table. The Sun Cluster Geographic Edition software does not handle other parameters of the volume set, including disk queue, size of memory queue, and number of asynchronous threads. You must adjust these parameters manually by using Sun StorageTek Availability Suite commands.

Field	Meaning	Description
phost	Primary host	The logical host of the server on which the primary volume resides.
pdev	Primary device	Primary volume partition. Specify full path names only.
pbitmap	Primary bitmap	Volume partition in which the bitmap of the primary partition is stored. Specify full path names only.
shost	Secondary host	The logical host of the server on which the secondary volume resides.
sdev	Secondary device	Secondary volume partition. Specify full path names only.
sbitmap	Secondary bitmap	Volume partition in which the bitmap of the secondary partition is stored. Specify full path names only.
ip	Network transfer protocol	IP address.

Field	Meaning	Description
sync   async	Operating mode	sync is the mode in which the I/O operation is confirmed as complete only when the volume on the secondary cluster has been updated.  async is the mode in which the primary host I/O operation is confirmed as complete before updating the volumes on the secondary cluster.
g <i>iogroupname</i>	I/O group name	An I/O group name. The set must be configured in the same I/O group on both the primary and the secondary cluster. This parameter is optional and need only be configured if you have an I/O group.
C	C tag	The device group name or resource tag of the local data and bitmap volumes in cases where this information is not implied by the name of the volume. For example, /dev/md/avsset/rdisk/vol indicates a device group named avsset. As another example, /dev/vx/rdsk/avsdg/vol indicates a device group named avsdg.

The Sun Cluster Geographic Edition software does not modify the value of the Sun StorageTek Availability Suite parameters. The software controls only the role of the volume set during switchover and takeover operations.

For more information about the format of the volume set files, refer to the Sun StorageTek Availability Suite documentation.

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

Sun Cluster Geographic Edition supports the use of raw-disk device groups in addition to various volume managers. When you initially configure Sun 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 Sun Cluster Geographic Edition.

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

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

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

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

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



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

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

```
phys-paris-1# cldevicegroup create -n phys-paris-1,phys-paris-2 \
-t rawdisk -d d7,d8 rawdg
phys-paris-1# /usr/cluster/lib/dcs/dgconv -d d7 rawdg
phys-paris-1# /usr/cluster/lib/dcs/dgconv -d d8 rawdg
```

**3 On the partner cluster, unconfigure the predefined device groups for the devices that you want to use.**

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
```

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

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

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

**5 Use the new group name where a device group name is required.**

The following command adds `rawdg` to the AVS protection group `rawpg`.

```
phys-paris-1# geopg add-device-group -p local_logical_host=paris-1h \
-p remote_logical_host=newyork-1h rawdg rawpg
```

## ▼ How to Configure a Sun StorageTek Availability Suite Volume in Sun Cluster

This procedure configures Sun StorageTek Availability Suite volumes in a Sun Cluster environment. These volumes can be Solaris Volume Manager volumes, VERITAS Volume Manager volumes, or raw device volumes.

The volumes are encapsulated at the Sun Cluster device-group level. The Sun StorageTek Availability Suite software interacts with the Solaris Volume Manager disksets, or VERITAS Volume Manager disk group, or raw device through this device group interface. The path to the volumes depends on the volume type, as described in the following table.

Volume Type	Path
Solaris Volume Manager	/dev/md/ <i>disksetname</i> /rdsk/d#, where # represents a number
VERITAS Volume Manager	/dev/vx/rdsk/ <i>diskgroupname</i> / <i>volumename</i>
Raw device	/dev/did/rdsk/d#s#

- 1 **Create a diskset, *avset*, by using Solaris Volume Manager or a disk group, *avsdg*, by using VERITAS Volume Manager or a raw device on `cluster-paris` and `cluster-newyork`.**

For example, if you configure the volume by using a raw device, choose a raw device group, `dsk/d3`, on `cluster-paris` and `cluster-newyork`.

- 2 **Create two volumes in the diskset or disk group on `cluster-paris`.**

The Sun StorageTek Availability Suite software requires a dedicated bitmap volume for each data volume to track which modifications to the data volume when the system is in logging mode.

If you use a raw device to configure the volumes, create two partitions, `/dev/did/rdsk/d3s3` and `/dev/did/rdsk/d3s4`, on the `/dev/did/rdsk/d3` device on `cluster-paris`.

- 3 **Create two volumes in the diskset or disk group on `cluster-newyork`.**

If you use a raw device to configure the volumes, create two partitions, `/dev/did/rdsk/d3s5` and `/dev/did/rdsk/d3s6`, on the `/dev/did/rdsk/d3` device on `cluster-paris`.

## Enabling a Sun StorageTek Availability Suite Volume Set

You can enable the Sun StorageTek Availability Suite volume sets in one of two ways:

- Automatically, when the device group is added to the protection group, `avspg`  
Use the automatic procedures to prepare the `devicegroupname-volset.ini` file when you are setting up Sun StorageTek Availability Suite software for the first time. After you have prepared the file, when you add the device group to the protection group, set the `Enable_volume_set` property of a device group to `True`. The Sun StorageTek Availability Suite software reads the information in the `devicegroupname-volset.ini` file to automatically enable the device group.
- Manually, after the device group is added to the protection group, `avspg`  
Use the manual procedures to enable the volume sets when you are creating volumes on a system that has been configured.

## Automatically Enabling a Solaris Volume Manager Volume Set

In this example, the `cluster-paris` cluster is the primary and `avsset` is a device group that contains a Solaris Volume Manager diskset.

### EXAMPLE 1-1 Automatically Enabling a Solaris Volume Manager Volume Set

This example has the following entries in the `/var/cluster/geo/avs/avsset-volset.ini` file:

```
logicalhost-paris-1 /dev/md/avsset/rdisk/d100 /dev/md/avsset/rdisk/d101
logicalhost-newyork-1 /dev/md/avsset/rdisk/d100 /dev/md/avsset/rdisk/d101
ip async C avsset
```

The `avsset-volset.ini` file contains the following entries:

- `lh-paris-1` – Primary host
- `/dev/md/avsset/rdisk/d100` – Primary data
- `/dev/md/avsset/rdisk/d101` – Primary bitmap
- `lh-newyork-1` – Secondary host
- `/dev/md/avsset/rdisk/d100` – Secondary data
- `/dev/md/avsset/rdisk/d101` – Secondary bitmap
- `ip` – Protocol
- `async` – Mode
- `C` – C tag
- `avsset` – Diskset

The sample configuration file defines a volume set that replicates `d100` from `cluster-paris` to `d100` on `cluster-newyork` by using the bitmap volumes and logical hostnames that are specified in the file.

## Automatically Enabling a VERITAS Volume Manager Volume Set

In this example, the `cluster-paris` cluster is the primary and `avsdg` is a device group that contains a VERITAS Volume Manager disk group.

### EXAMPLE 1-2 Automatically Enabling a VERITAS Volume Manager Volume Set

This example has the following entries in the `/var/cluster/geo/avs/avsdg-volset.ini` file:

```
logicalhost-paris-1 /dev/vx/rdisk/avsdg/vol-data-paris \
/dev/vx/rdisk/avsdg/vol-bitmap-paris
logicalhost-newyork-1 /dev/vx/rdisk/avsdg/vol-data-newyork \
/dev/vx/rdisk/avsdg/vol-bitmap-ny
ip async C avsdg
```

The `avsdg-volset.ini` file contains the following entries:

- `lh-paris-1` – Primary host

**EXAMPLE 1-2** Automatically Enabling a VERITAS Volume Manager Volume Set *(Continued)*

- /dev/vx/rdisk/avsdg/vol-data-paris – Primary data
- /dev/vx/rdisk/avsdg/vol-bitmap-paris – Primary bitmap
- lh-newyork-1 is the secondary host.
- /dev/vx/rdisk/avsdg/vol-data-newyork – Secondary data
- /dev/vx/rdisk/avsdg/vol-bitmap-ny – Secondary bitmap
- ip – Protocol
- async – Mode
- C – C flag
- avsdg – Device group

The sample configuration file defines a volume set that replicates `vol-data-paris` from `cluster-paris` to `vol-data-newyork` on `cluster-newyork`. The volume set uses the bitmap volumes and logical hostnames that are specified in the file.

## Automatically Enabling a Raw Device Volume Set

In this example, the `cluster-paris` cluster is the primary and `rawdg` is the name of the device group that contains a raw device disk group, `/dev/did/rdisk/d3`.

**EXAMPLE 1-3** Automatically Enabling a Raw Device Volume Set

This example has the following entries in `/var/cluster/geo/avs/avsdg-volset.ini` file:

```
logicalhost-paris-1 /dev/did/rdisk/d3s3 /dev/did/rdisk/d3s4
logicalhost-newyork-1 /dev/did/rdisk/d3s5 /dev/did/rdisk/d3s6
ip async C rawdg
```

The `rawdg-volset.ini` file contains the following entries:

- `logicalhost-paris-1` – Primary host
- `/dev/did/rdisk/d3s3` – Primary data
- `/dev/did/rdisk/d3s4` – Primary bitmap
- `logicalhost-newyork-1` – Secondary host
- `/dev/did/rdisk/d3s5` – Secondary data
- `/dev/did/rdisk/d3s6` – Secondary bitmap
- ip – Protocol
- async – Mode
- C – C flag
- `rawdg` – Device group

The sample configuration file defines a volume set that replicates `d3s3` from `cluster-paris` to `d3s5` on `cluster-newyork`. The volume set uses the bitmap volumes and logical hostnames that are specified in the file.

## Manually Enabling Volume Sets

After you have added the device group to the protection group, `avspg`, you can manually enable the Sun StorageTek Availability Suite volume sets. Because the Sun Availability Suite commands are installed in different locations in the supported software versions, the following examples illustrate how to enable volume sets for each software version.

### EXAMPLE 1-4 Manually Enabling a Sun StorageTek Availability Suite 4 Volume Set

This example manually enables a Solaris Volume Manager volume set when using Sun StorageTek Availability Suite 4.

```
phys-paris-1# /usr/sbin/sndradm -e localhost-paris-1 \
/dev/md/avssset/rdisk/d100 /dev/md/avssset/rdisk/d101 \
localhost-newyork-1 /dev/md/avssset/rdisk/d100 \
/dev/md/avssset/rdisk/d101 ip async C avssset
```

### EXAMPLE 1-5 Manually Enabling a Sun StorEdge Availability Suite 3.2.1 Volume Set

This example manually enables a Solaris Volume Manager volume set when using Sun StorEdge Availability Suite 3.2.1.

```
phys-paris-1# /usr/opt/SUNWesm/sbin/sndradm -e localhost-paris-1 \
/dev/md/avssset/rdisk/d100 /dev/md/avssset/rdisk/d101 \
localhost-newyork-1 /dev/md/avssset/rdisk/d100 \
/dev/md/avssset/rdisk/d101 ip async C avssset
```

### EXAMPLE 1-6 Manually Enabling a VERITAS Volume Manager Volume Set

This example manually enables a VERITAS Volume Manager volume set when using Sun StorageTek Availability Suite 4.

```
phys-paris-1# /usr/sbin/sndradm -e localhost-paris-1
/dev/vx/rdisk/avsdg/vol-data-paris /dev/vx/rdisk/avsdg/vol-bitmap-paris
localhost-newyork-1 /dev/vx/rdisk/avsdg/vol-data-newyork
/dev/vx/rdisk/avsdg/vol-bitmap-newyork ip async C avsdg
```

This example manually enables a VERITAS Volume Manager volume set when using Sun StorEdge Availability Suite 3.2.1.

```
phys-paris-1# /usr/opt/SUNWesm/sbin/sndradm -e localhost-paris-1
/dev/vx/rdisk/avsdg/vol-data-paris /dev/vx/rdisk/avsdg/vol-bitmap-paris
localhost-newyork-1 /dev/vx/rdisk/avsdg/vol-data-newyork
/dev/vx/rdisk/avsdg/vol-bitmap-newyork ip async C avsdg
```

**EXAMPLE 1-7** Manually Enabling a Raw Device Volume Set

This example manually enables a raw device volume set when using Sun StorageTek Availability Suite 4.

```
phys-paris-1# /usr/sbin/sndradm -e logicalhost-paris-1
/dev/did/rdisk/d3s3 /dev/did/rdisk/d3s4 logicalhost-newyork-1 /dev/did/rdisk/d3s5
/dev/did/rdisk/d3s6 ip async C dsk/d3
```

This example manually enables a raw device volume set when using Sun StorEdge Availability Suite 3.2.1.

```
phys-paris-1# /usr/opt/SUNWesm/sbin/sndradm -e logicalhost-paris-1
/dev/did/rdisk/d3s3 /dev/did/rdisk/d3s4 logicalhost-newyork-1 /dev/did/rdisk/d3s5
/dev/did/rdisk/d3s6 ip async C dsk/d3
```

Information about the `sndradm` command execution is written to the Sun StorageTek Availability Suite log file at the following locations:

- When using Sun StorageTek Availability Suite 4, `/var/adm/ds.log`
- When using Sun StorEdge Availability Suite 3.2.1, `/var/opt/SUNWesm/ds.log`

Refer to this file if errors occur while manually enabling the volume set.

## ▼ How to Configure the Sun Cluster Device Group That Is Controlled by Sun StorageTek Availability Suite

Sun StorageTek Availability Suite software supports Solaris Volume Manager, VERITAS Volume Manager, and raw device volumes.

- 1 Ensure that the device group that contains the volume set that you want to replicate is registered with Sun Cluster software.**

```
# cldevicegroup show -v dg1
```

For more information about this command, refer to the `cldevicegroup(1CL)` man page.

- 2 If you are using a VERITAS Volume Manager device group, synchronize the VERITAS Volume Manager configuration by using the Sun Cluster command `clsetup` or `cldevicegroup`.**

- 3 Ensure that the device group is displayed in the output of the `cldevicegroup show` command.**

```
# cldevicegroup show -v dg1
```

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

- 4 Repeat steps 1–3 on both clusters, `cluster-paris` and `cluster-newyork`.**

## ▼ How to Configure a Highly Available Cluster Global File System for Use With Sun StorageTek Availability Suite

- 1 **Create the required file system on the volume set that you created in the previous step, `vol-data-paris`.**  
The application writes to this file system.
- 2 **Add an entry to the `/etc/vfstab` file that contains information such as the mount location.**

---

**Note** – You must specify 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 Sun Cluster software and the Sun Cluster Geographic Edition framework handle mounting the file system by using the `HASStoragePlus` resource when the application is brought online on the primary cluster. You must not mount data on the secondary cluster because data on the primary will not be replicated to the secondary cluster.

---

- 3 **To handle the new file system, add the `HASStoragePlus` resource to the application resource group, `apprg1`.**  
Adding this resource ensures that the necessary file systems are remounted before the application is started.  
  
For more information about the `HASStoragePlus` resource type, refer to the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.
- 4 **Repeat steps 1–3 on both `cluster-paris` and `cluster-newyork`.**

### Example 1–8 Configuring a Highly Available Cluster Global File System for Solaris Volume Manager Volumes

This example configures a highly available cluster global file system for Solaris Volume Manager volumes. This example assumes that the resource group `apprg1` already exists.

1. Create a UNIX file system (UFS).

```
# newfs /dev/md/avsset/rdisk/d100
```

This command creates the following entry in the `/etc/vfstab` file:

```
/dev/md/avsset/dsk/d100 /dev/md/avsset/rdisk/d100
/global/sample ufs 2 no logging
```

2. Add the `HASStoragePlus` resource.

```
# clresource create -g apprg1 -t SUNWHASStoragePlus \  
-p FilesystemMountPoints=/global/sample -p Affinityon=TRUE rs-hasp
```

### Example 1-9 Configuring a Highly Available Cluster Global File System for VERITAS Volume Manager Volumes

This example assumes that the apprg1 resource group already exists.

1. Create a UNIX file system (UFS).

```
# newfs /dev/vx/rdisk/avsdg/vol-data-paris
```

This command creates the following entry in the `/etc/vfstab` file:

```
/dev/vx/dsk/avsdg/vol-data-paris /dev/vx/rdisk/avsdg/vol-data-paris  
/global/sample ufs 2 no logging
```

2. Add the HASStoragePlus resource.

```
# clresource create -g apprg1 -t SUNWHASStoragePlus \  
-p FilesystemMountPoints=/global/sample -p Affinityon=TRUE rs-hasp
```

### Example 1-10 Configuring a Highly Available Cluster Global File System for Raw Device Volumes

This example assumes that the apprg1 resource group already exists.

1. Create a UNIX file system (UFS).

```
# newfs /dev/did/rdsk/d3s3
```

This command creates the following entry in the `/etc/vfstab` file:

```
/dev/did/dsk/d3s3 /dev/did/rdsk/d3s3  
/global/sample ufs 2 no logging
```

2. Add the HASStoragePlus resource.

```
# clresource create -g apprg1 -t SUNWHASStoragePlus \  
-p FilesystemMountPoints=/global/sample -p Affinityon=TRUE rs-hasp
```



# Administering Sun StorageTek Availability Suite Protection Groups

---

This chapter describes the procedures for administering data replication with Sun StorageTek Availability Suite software. This chapter contains the following sections:

- “Strategies for Creating Sun StorageTek Availability Suite Protection Groups” on page 25
- “Creating, Modifying, Validating, and Deleting a Sun StorageTek Availability Suite Protection Group” on page 34
- “Administering Sun StorageTek Availability Suite Application Resource Groups” on page 40
- “Administering Sun StorageTek Availability Suite Data Replication Device Groups” on page 44
- “Replicating the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster” on page 50
- “Activating and Deactivating a Protection Group” on page 52
- “Resynchronizing a Sun StorageTek Availability Suite Protection Group” on page 57
- “Checking the Runtime Status of Sun StorageTek Availability Suite Data Replication” on page 58

## Strategies for Creating Sun StorageTek Availability Suite Protection Groups

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

- Stopping the application before creating the protection group  
This strategy is the most straightforward. However, because the protection group is not brought online until the end of the process, you must unmanage the application resource group to add it to the protection group.
- Creating the protection group while the application remains online  
While this strategy enables you to create a protection group without any application outage, it requires issuing more commands.

Before you create a protection group by using the steps in the following sections, ensure that the following prerequisites are met.

- The application has been configured by Sun Cluster software on both clusters.
- Corresponding device groups are configured for data replication.

## 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 node on one cluster.  
For more information, see [“How to Create and Configure a Sun StorageTek Availability Suite Protection Group”](#) on page 34.
- Add the data replication device group to the protection group.  
For more information, see [“How to Add a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group”](#) on page 45.
- Take the application resource group to the unmanaged state.
- Add the application resource group to the protection group on one cluster.  
For more information, see [“How to Add an Application Resource Group to a Sun StorageTek Availability Suite Protection Group”](#) on page 41.
- On the other cluster, retrieve the protection group configuration.  
For more information, see [“How to Replicate the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster”](#) on page 50.
- From either cluster, activate the protection group globally.  
For more information, see [“How to Activate a Sun StorageTek Availability Suite Protection Group”](#) on page 52.

## 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 node on one cluster.  
For more information, see [“How to Create and Configure a Sun StorageTek Availability Suite Protection Group”](#) on page 34.

- Add the data replication device group to the protection group.  
For more information, see [“How to Add a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group”](#) on page 45.
- Activate the protection group locally.  
For more information, see [“How to Activate a Sun StorageTek Availability Suite Protection Group”](#) on page 52.
- Add the application resource group to the protection group.  
For more information, see [“How to Add an Application Resource Group to a Sun StorageTek Availability Suite Protection Group”](#) on page 41.

Complete the following steps on the other cluster.

- Retrieve the protection group configuration.  
For more information, see [“How to Replicate the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster”](#) on page 50.
- Activate the protection group locally.  
For more information, see [“How to Activate a Sun StorageTek Availability Suite Protection Group”](#) on page 52.

**EXAMPLE 2-1** Creating a Sun StorageTek Availability Suite 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 the `cluster-paris` cluster.
  - a. On a node of the cluster where an application resource group is already running, check and fix any `NodeList` inconsistencies that might exist between the resource group and the device group with which the resource group has affinities.
  - b. Create the protection group with a matching `NodeList`.

```
phys-paris-1# clresourcegroup show -v apprg1 | grep NodeList
NodeList:                phys-paris-1 phys-paris-2
phys-paris-1# cldevicegroup show -v avsdg1 | grep "Node List:"
Node List:                phys-paris-2, phys-paris-1
```

The node list of the device group is in a different order from the node list of the resource group. The order of the node list of the device group is changed as follows:

```
phys-paris-1# cldevicegroup set -p preferenced=true -n phys-paris-1,phys-paris-2 avsdg
```

You can also change the node list of a resource group to meet this requirement.

**EXAMPLE 2-1** Creating a Sun StorageTek Availability Suite Protection Group While the Application Remains Online *(Continued)*

The protection group is created with a `Nodelist` identical to the `Nodelist` of the resource group and device group:

```
phys-paris-1# geopg create -d avs -p Nodelist=phys-paris-1,phys-paris-2 \
-o Primary -s paris-newyork-ps avspg
phys-paris-1# Protection group "avspg" has been successfully created
```

2. Add the Sun StorageTek Availability Suite device group, `avsdg`, to the protection group.

```
phys-paris-1# geopg add-device-group -p Local_logical_host=lh-paris-1 \
-p Remote_logical_host=lh-newyork-1 -p Enable_volume_set=True avsdg avspg
```

3. Verify that the data replication resource groups and the lightweight resource groups have been created and are online. Also, verify that the Sun StorageTek Availability Suite volume set has been enabled by setting the `Enable-volume-set` property to `True`.

```
phys-paris-1# dsstat
name          t      s      pct      role      ckps      dkps      tps      svf
/avsdg/rdsk/d100 P      L      100.00    net      -         0         0         0
/avsdg/rdsk/d101                bmp      0         0         0         0
```

```
phys-paris-1# clresource list -v
Resource Name      Resource Type      Resource Group
-----
geo-clustername    SUNW.LogicalHostname:2  geo-infrastructure
geo-hbmonitor       SUNW.HBmonitor      geo-infrastructure
geo-failovercontrol SUNW.scmasa          geo-infrastructure
lh-paris-1         SUNW.LogicalHostname:2  avsdg-stor-rg
avsdg-stor         SUNW.HASStoragePlus:4  avsdg-stor-rg
avsdg-rep-rs       SUNW.GeoCtlAVS        avspg-rep-rg
avs-lh             SUNW.LogicalHostname:2  apprg1
avs-stor           SUNW.HASStoragePlus:4  apprg1
avs-server-res     SUNW.oracle_server:6   apprg1
avs-listener-res   SUNW.oracle_listener:5 apprg1
```

```
phys-paris-1# clresourcegroup status
```

```
=== Cluster Resource Groups ===
```

Group Name	Node Name	Suspended	Status
geo-clusterstate	phys-paris-1	No	Online
	phys-paris-2	No	Online
geo-infrastructure	phys-paris-1	No	Online

**EXAMPLE 2-1** Creating a Sun StorageTek Availability Suite Protection Group While the Application Remains Online (Continued)

```

                                phys-paris-2  No      Offline
avsdg-stor-rg                 phys-paris-1  No      Online
                                phys-paris-2  No      Offline
avspg-rep-rg                  phys-paris-1  No      Online
                                phys-paris-2  No      Offline
apprg1                         phys-paris-1  No      Online
                                phys-paris-2  No      Offline

```

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

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

Resource Name	Node Name	State	Status Message
-----	-----	-----	-----
geo-clustername	phys-paris-1	Online	Online - LogicalHostname online.
	phys-paris-2	Offline	Offline
geo-hbmonitor	phys-paris-1	Online	Online
	phys-paris-2	Offline	Offline
geo-failovercontrol	phys-paris-1	Online	Online
	phys-paris-2	Offline	Offline
lh-paris-1	phys-paris-1	Online	Online - LogicalHostname online.
	phys-paris-2	Offline	Offline
avsdg-stor	phys-paris-1	Online	Online
	phys-paris-2	Offline	Offline
avsdg-rep-rs	phys-paris-1	Online	Degraded - Logging
	phys-paris-2	Offline	Offline
avs-lh	phys-paris-1	Online	Online - LogicalHostname online.
	phys-paris-2	Offline	Offline
avs-stor	phys-paris-1	Online	Online
	phys-paris-2	Offline	Offline
avs-server-res	phys-paris-1	Online	Online
	phys-paris-2	Offline	Offline

**EXAMPLE 2-1** Creating a Sun StorageTek Availability Suite Protection Group While the Application Remains Online *(Continued)*

```

avs-listener-res      phys-paris-1  Online      Online
                    phys-paris-2  Offline     Offline

```

For the Sun StorageTek Availability Suite 4 software, this command is located at `/usr/sbin/dsstat`. For the Sun StorEdge Availability Suite 3.2.1 software, this command is located at `/usr/opt/SUNWesm/sbin/dsstat`.

4. Activate the protection group locally.

```

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

```

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

```

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

```

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 "avspg"             : Unknown
  Partnership                         : paris-newyork-ps
  Synchronization                    : Error

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

Cluster cluster-newyork              : Unknown

```

**EXAMPLE 2-1** Creating a Sun StorageTek Availability Suite Protection Group While the Application Remains Online *(Continued)*

```

Role                               : Unknown
Activation State                    : Unknown
Configuration                       : Unknown
Data Replication                    : Unknown
Resource Groups                     : Unknown

```

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

```

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

```

7. Verify that the data replication resource groups and the lightweight resource groups have been created and brought online.

```

phys-newyork-1# dsstat
name          t  s  pct  role  ckps  dkps  tps  svt
/avsdg/rdisk/d100  S  L  100.00  net   -    0    0    0
/avsdg/rdisk/d101                bmp   0    0    0    0

```

```

phys-newyork-1# clresource list -v
Resource Name      Resource Type      Resource Group
-----
geo-clustername   SUNW.LogicalHostname:2  geo-infrastructure
geo-hbmonitor     SUNW.HBmonitor      geo-infrastructure
geo-failovercontrol SUNW.scmasa         geo-infrastructure
lh-newyork-1      SUNW.LogicalHostname:2  avsdg-stor-rg
avsdg-stor        SUNW.HASStoragePlus:4  avsdg-stor-rg
avsdg-rep-rs     SUNW.GeoCtlAVS       avspg-rep-rg
avs-lh           SUNW.LogicalHostname:2  apprg1
avs-stor         SUNW.HASStoragePlus:4  apprg1
avs-server-res   SUNW.oracle_server:6   apprg1
avs-listener-res SUNW.oracle_listener:5  apprg1

```

```

phys-newyork-1# clresourcegroup status

```

```

=== Cluster Resource Groups ===

```

Group Name	Node Name	Suspended	Status
geo-clusterstate	phys-newyork-1	No	Online
	phys-newyork-2	No	Online
geo-infrastructure	phys-newyork-1	No	Online
	phys-newyork-2	No	Offline

**EXAMPLE 2-1** Creating a Sun StorageTek Availability Suite Protection Group While the Application Remains Online *(Continued)*

```

avsdg-stor-rg      phys-newyork-1  No    Online
                   phys-newyork-2  No    Offline

avspg-rep-rg      phys-newyork-1  No    Online
                   phys-newyork-2  No    Offline

apprg1            phys-newyork-1  No    Unmanaged
                   phys-newyork-2  No    Unmanaged
    
```

```
phys-newyork-1# clresource status
```

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

Resource Name	Node Name	State	Status Message
-----	-----	----	-----
geo-clustername	phys-newyork-1	Online	Online - LogicalHostname online.
	phys-newyork-2	Offline	Offline
geo-hbmonitor	phys-newyork-1	Online	Online
	phys-newyork-2	Offline	Offline
geo-failovercontrol	phys-newyork-1	Online	Online
	phys-newyork-2	Offline	Offline
lh-newyork-1	phys-newyork-1	Online	Online - LogicalHostname online.
	phys-newyork-2	Offline	Offline
avsdg-stor	phys-newyork-1	Online	Online
	phys-newyork-2	Offline	Offline
avsdg-rep-rs	phys-newyork-1	Online	Degraded - Logging
	phys-newyork-2	Offline	Offline
avs-lh	phys-newyork-1	Offline	Offline
	phys-newyork-2	Offline	Offline
avs-stor	phys-newyork-1	Offline	Offline
	phys-newyork-2	Offline	Offline
avs-server-res	phys-newyork-1	Offline	Offline
	phys-newyork-2	Offline	Offline
avs-listener-res	phys-newyork-1	Offline	Offline
	phys-newyork-2	Offline	Offline



**EXAMPLE 2-1** Creating a Sun StorageTek Availability Suite Protection Group While the Application Remains Online *(Continued)*

For the Sun StorageTek Availability Suite 4 software, this command is located at `/usr/sbin/dsstat`. For the Sun StorEdge Availability Suite 3.2.1 software, this command is located at `/usr/opt/SUNWesm/sbin/dsstat`.

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

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

9. 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      : cluster-newyork
  Synchronization      : OK
  ICRM Connection      : OK

Heartbeat "paris-to-newyork" monitoring "cluster-newyork": OK
  Heartbeat plug-in "ping_plugin"      : Inactive
  Heartbeat plug-in "tcp_udp_plugin": OK

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

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

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

# Creating, Modifying, Validating, and Deleting a Sun StorageTek Availability Suite Protection Group

This section contains procedures for the following tasks:

- [“How to Create and Configure a Sun StorageTek Availability Suite Protection Group” on page 34](#)
- [“How to Modify a Sun StorageTek Availability Suite Protection Group” on page 36](#)
- [“How to Validate a Sun StorageTek Availability Suite Protection Group” on page 37](#)
- [“How to Delete a Sun StorageTek Availability Suite Protection Group” on page 39](#)

---

**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 data replication of `NONE`.

---

## ▼ How to Create and Configure a Sun StorageTek Availability Suite Protection Group

**Before You Begin** 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 Sun Cluster Geographic Edition namespace. You cannot use the same protection group name in two partnerships on the same system.

---

You can also replicate the existing configuration of a protection group from a remote cluster to the local cluster. For more information, see [“Replicating the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster” on page 50](#).

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

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

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

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

<code>-s partnershipname</code>	Specifies the name of the partnership.
<code>-d avs</code>	Specifies that the protection group data is replicated by Sun StorageTek Availability Suite software.
<code>-o localrole</code>	Specifies the role of this protection group on the local cluster as either primary or secondary.
<code>-p propertysetting</code>	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.
- `Enable_volume_set` – Enables the volume set automatically at protection group creation time when set to `true`.

For more information about automatically enabling a volume set, see [“Enabling a Sun StorageTek Availability Suite Volume Set” on page 18](#).

- `NodeList` – Lists the host names of the machines that can be primary for the device group in the protection group.

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

`protectiongroupname` Specifies the name of the protection group.

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

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

Before creating the protection group, the data replication layer validates that the configuration is correct.

If the validation is successful, the local `Configuration` status is set to `OK` and the `Synchronization` status is set to `Error`.

If the validation is unsuccessful, the protection group is not created.

**Example 2-2** Creating and Configuring a Protection Group

This example creates a Sun StorageTek Availability Suite protection group on the cluster-`paris` cluster, which is set as the primary cluster.

```
phys-paris-1# geopg create -s paris-newyork-ps -d avs -o primary \
-p Nodelist=phys-paris-1,phys-paris-2 avspg
```

## ▼ How to Modify a Sun StorageTek Availability Suite Protection Group

**Before You Begin** Ensure that the protection group you want to modify exists locally.

### 1 Log in to a cluster node.

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

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

If the partner cluster contains a protection group of the same name, the `geopg set-prop` command also propagates the new configuration information to the partner cluster.

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

`-p propertysetting` Specifies the properties of the protection group.

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

`protectiongroupname` Specifies the name of the protection group.

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

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

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

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

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

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

```
# geopg set-prop -p Timeout=300 avspg
```

## ▼ How to Validate a Sun StorageTek Availability Suite Protection Group

**Before You Begin** When the Configuration status of a protection group is displayed as Error in the `geoadm` status output, you can validate the configuration by using the `geopg validate` command. This command checks the current state of the protection group and its entities.

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

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

Before validating the configuration of a protection group, ensure that the protection group you want to validate exists locally and that the common agent container is online on all nodes of both clusters in the partnership.

### 1 Log in to a cluster node.

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

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

This command validates the configuration of the protection group on the local cluster only.

```
# geopg validate protectiongroupname
```

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

**Example 2-4** Validating the Configuration of a Protection Group

This example validates a protection group.

```
# geogg validate avspg
```

## How the Data Replication Layer Validates the Application Resource Groups and Data Replication Entities

During protection group validation, the Sun StorageTek Availability Suite data replication layer validates the application resource groups and the data replication entities as follows:

- Verifies that an application resource group in the protection group has its `Auto_start_on_new_cluster` property set to `False`.  
When you bring a protection group online on the primary cluster, bring the application resources groups participating in that protection group online only on the same primary cluster. Setting the `Auto_start_on_new_cluster` property to `False` prevents the Sun Cluster resource group manager from automatically starting the application resource groups. In this case, the startup of resource groups is reserved for the Sun Cluster Geographic Edition software.  
Application resource groups should be online only on primary cluster when the protection group is activated.
- Verifies that the `NodeList` property of an application resource group that has affinities with a device group defined by the `HASStoragePlus` resource contains the same entries in identical order to the `NodeList` property of the protection group.
- Verifies that the `NodeList` property of a device group in the protection group contains the same entries in identical order to the `NodeList` property of the protection group.
- Verifies that a lightweight resource group is created for each device group in the protection group. Each lightweight resource group contains two resources, a logical hostname resource and a `HASStoragePlus` resource. For more information about lightweight resource groups and their resources, see [“Sun StorageTek Availability Suite Lightweight Resource Groups” on page 13](#).
- Verifies that a replication resource of the type `GeoCtlAVS` is created in the replication resource group of each device group in the protection group. For information about the format of the replication resource group, see [“Sun StorageTek Availability Suite Replication Resource Groups” on page 13](#).
- Verifies that the `NodeList` property of the lightweight resource group and replication resource group contains the same entries in identical order to the `NodeList` property of the protection group.

If the `Enable_volume_set` property of a successfully validated device group is set to `True`, then volume sets defined in the `/var/cluster/geo/avs/avsdg-volset.ini` file are enabled. Other volume sets for the device group are disabled. If you want to enable the other volume sets, you can add the volume sets to the `/var/cluster/geo/avs/avsdg-volset.ini` file or set the `Enable_volume_set` property to `False`.

When validation is complete, the Sun Cluster Geographic Edition software creates the lightweight resource group, the replication resource group, and the resources for this replication resource group, if nonexistent, and brings them online. If a resource group or resource of the same name already exists, the Sun Cluster Geographic Edition operations might modify their properties. Sun Cluster Geographic Edition software cannot create a new resource group or resource of the same name if one already exists.

The `Configuration` status is set to `OK` after successful validation. If validation is not successful, the `Configuration` status is set to `Error`.

## ▼ How to Delete a Sun StorageTek Availability Suite Protection Group

**Before You Begin** To delete a protection group on all clusters, run the `geogg delete` command on each cluster where the protection group exists.

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

- The protection group exists locally.
- The protection group is offline on the local cluster.

---

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

---

- 1 Log in to a node on the cluster where you want to delete the protection group, for example, `cluster-paris`.**

The `cluster-paris` is the primary cluster. See “Example Sun Cluster Geographic Edition Cluster Configuration” in *Sun Cluster Geographic Edition System Administration Guide* for a sample cluster configuration.

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

## 2 Delete the protection group.

This command deletes the configuration of the protection group from the local cluster. The command also removes the lightweight resource group and the replication resource group for each device group in the protection group.

```
# geopg delete protectiongroupname
```

*protectiongroupname* Specifies the name of the protection group.

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

### Example 2-5 Deleting a Protection Group

This example deletes a protection group from both partner clusters.

```
# rlogin cluster-paris -l root
cluster-paris# geopg delete avspg
# rlogin cluster-newyork -l root
cluster-newyork# geopg delete avspg
```

### Example 2-6 Deleting a Protection Group While Keeping Application Resource Groups Online

This example keeps online two application resource groups (`apprg1` and `apprg2`) while deleting the protection group that they share, `avspg`.

Remove the application resource groups from the protection group, then delete the protection group.

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

# Administering Sun StorageTek Availability Suite Application Resource Groups

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

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



This section contains information about the following tasks:

- “How to Add an Application Resource Group to a Sun StorageTek Availability Suite Protection Group” on page 41
- “How to Delete an Application Resource Group From a Sun StorageTek Availability Suite Protection Group” on page 43

## ▼ How to Add an Application Resource Group to a Sun StorageTek Availability Suite 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 to add already 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 show` command.

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

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

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

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

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

- The `NodeList` property of the failover application resource group that has affinities with a device group defined by the resource must contain the same entries in identical order to the `NodeList` property of the protection group.
- The application resource group must not have dependencies on resource groups and resources outside of this protection group. To add several application resource groups that share dependencies, you must 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 will fail.

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

If the resource group is `Unmanaged` and the protection group is activated after the configuration of the protection group has changed, then the local state of the protection group becomes `Error`.

If the resource group to add is `Online` and the protection group is deactivated, the request is rejected. You must activate the protection group before adding an online resource group.

### 1 Log in to a cluster node.

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

### 2 Add an application resource group to the protection group.

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

*resourcegroup* Specifies the name of the application resource group.

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

*protectiongroup* Specifies the name of the protection group.

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.

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

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

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

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

If the application resource group is a failover type resource group that shares affinities with a device group in the same protection group, the Sun Cluster Geographic Edition software alters

its `RG_affinities` property to include a strong, positive affinity to an internal resource group, called a *lightweight resource group*. This affinity includes failover delegation.

The application resource group must not have strong, positive affinities that have failover delegation with other resource groups. Otherwise, trying to include a strong positive affinity with failover delegation on the lightweight resource group fails.

The Sun Cluster Geographic Edition software also creates strong dependencies between the `HASStoragePlus` resource in the application resource group and the `HASStoragePlus` resource in the lightweight resource group for this device group. This redirection occurs when the protection group is brought online or when an online application resource group is added to an online protection group.

Do not modify dependencies and resource group affinities between application resource groups and lightweight resource groups.

### Example 2-7 Adding an Application Resource Group to a Sun StorageTek Availability Suite Protection Group

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

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

## ▼ How to Delete an Application Resource Group From a Sun StorageTek Availability Suite Protection Group

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

**Before You Begin** Ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The resource group to remove is part of the application resource groups of the protection group. For example, you cannot remove a resource group that belongs to the data replication management entity.

### 1 Log in to a cluster node.

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

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

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

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

*resourcegrouplist* Specifies the name of the application resource group.

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

*protectiongroup* Specifies the name of the protection group.

If the resource group that is being removed shares dependencies with other resource groups in the protection group, then you must also remove all other resource groups that share dependencies with the resource group that is being removed.

If the remove operation failed 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.

Sun Cluster Geographic Edition software removes the affinity and resource dependencies between the application resource group and the lightweight resource group.

**Example 2-8 Deleting an Application Resource Group From a Protection Group**

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

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

## Administering Sun StorageTek Availability Suite Data Replication Device Groups

This section describes the following tasks for administering data replication device groups in a Sun StorageTek Availability Suite protection group:

- [“How to Add a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group” on page 45](#)
- [“How to Modify a Sun StorageTek Availability Suite Data Replication Device Group” on page 48](#)
- [“How to Delete a Data Replication Device Group From a Sun StorageTek Availability Suite Protection Group” on page 49](#)

For details about configuring a Sun StorageTek Availability Suite protection group, see [“How to Create and Configure a Sun StorageTek Availability Suite Protection Group”](#) on page 34.

## ▼ How to Add a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group

### Before You Begin

A protection group is the container for the application resource groups, which contain data for services that are protected from disaster. Sun Cluster Geographic Edition software protects the data by replicating it from the primary cluster to the secondary cluster. By adding a Sun Cluster device group to a protection group, Sun Cluster Geographic Edition software monitors the replication status of all volumes in the device group that belong to a Sun StorageTek Availability Suite volume set. Sun Cluster Geographic Edition software also controls the role and state of the volume set during protection group operations like start, stop, switchover, and takeover.

Before you add a device group to a protection group, ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The protection group is offline on the local cluster and the partner cluster, if the partner cluster can be reached.
- The device group exists on both the local cluster and the partner cluster.
- The `NodeList` property of the device group contains the same entries in identical order to the `NodeList` property of the protection group.
- The `Local_logical_host` property specifies a valid hostname that can be hosted by the local cluster and that is reserved for this device group.
- The `Remote_logical_host` property specifies a valid hostname that can be hosted by the remote cluster and that has been reserved for this device group.
- If the `Enable_volume_set` property is set to `true`, then the `/var/cluster/geo/avs/avsdg-volset.ini` file must exist and contain valid entries on all nodes of both partner clusters. For information about configuring this file, see [“Enabling a Sun StorageTek Availability Suite Volume Set”](#) on page 18.

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

## 2 Add a data replication device group to the protection group.

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

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

*-p property* Specifies the properties of the data replication device group.

You can specify the following Sun StorageTek Availability Suite properties:

- *Local\_logical\_host* – Specifies the name of the local logical host that is used to replicate the device group.
- *Remote\_logical\_host* – Specifies the name of the remote logical host that is used to replicate the device group.
- *Enable\_volume\_set* – Specifies whether the volume sets in the file should be enabled automatically. Set to either `True` or `False`.

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

*AVSdevicegroupname* Specifies the name of the new data replication device group.

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

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

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

### Example 2-9 Adding a Data Replication Device Group to a Sun StorageTek Availability Suite Protection Group

This example creates a Sun StorageTek Availability Suite data replication device group in the `avspg` protection group.

```
# geopg add-device-group -p Local_logical_host=lh-paris-1 \  
-p Remote_logical_host=lh-newyork-1 avsdg avspg
```

## How the Data Replication Subsystem Verifies the Device Group

The device group that Sun StorageTek Availability Suite controls is added to a protection group. The data replication layer verifies that the device group exists and that the value of its `NodeList` property contains the same entries in identical order to the `NodeList` property of the protection group.

When the `geopg add-device-group` command runs, a lightweight resource group for the device group is created and brought online. The lightweight resource group contains the following resources:

- A logical hostname resource that is used for data replication, as specified in the `LocalLogicalHost` property
- An `HASStoragePlus` resource that controls the collocation of the device group with the lightweight resource group

For more information about lightweight resource groups and their resources, see [“Sun StorageTek Availability Suite Lightweight Resource Groups” on page 13](#).

When you run the `geopg add-device-group` command, a replication resource of the type `GeoCtlAVS` is created in the replication resource group of each device group in the protection group and brought online. For information about the format of the replication resource group, see [“Sun StorageTek Availability Suite Replication Resource Groups” on page 13](#).

The `NodeList` property of the lightweight resource group and replication resource group contains the same entries in identical order to the `NodeList` property of the protection group.

If a resource or resource group of the same name is already configured on the local cluster, Sun Cluster Geographic Edition verifies the configuration and sets the `Configuration` to `Error` if the configuration is not correct.

If the `EnableVolumeSet` property of this device group is set to `True`, then volume sets that are defined in the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file are enabled. Otherwise, Sun Cluster Geographic Edition software controls and monitors all volume sets that you enable manually by using the Sun StorageTek Availability Suite commands.

If the `geopg add-device-group` command is unsuccessful, the configuration of the protection group is not modified.

If the `geopg add-device-group` command is successful and the `Configuration` status on the local cluster is set to `OK`, then the new configuration is propagated to the partner cluster. This propagation causes the whole protection group configuration to revalidate on the partner cluster. During revalidation, the same entities are created on the partner cluster, including the lightweight resource group and the replication resource group. Volume sets are also enabled on

the partner clusters if the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file exists on the partner cluster and contains correctly defined volume sets. If the validation is unsuccessful, then the Configuration status on the partner cluster is set to Error on the partner cluster.



**Caution** – Do not use Sun Cluster commands to change, remove, or bring offline these resources or resource groups. Use only Sun Cluster Geographic Edition commands to administer lightweight resource groups, replication resource groups, and resources that are internal entities managed by Sun Cluster Geographic Edition software. Altering the configuration or state of these entities directly with Sun Cluster commands could result in unrecoverable failure.

If the device group on the partner cluster is validated successfully and the `Enable_volume_set` property of this device group is set to `true`, then the volume sets that are defined in the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file are enabled on the partner cluster. Other volume sets of the device group are disabled.

After the device group has been added to the protection group, you can enable or disable the volume sets of the device group directly by using the Sun StorageTek Availability Suite commands. The `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file is used only when the protection group that contains the device group is successfully validated for the first time.

## ▼ How to Modify a Sun StorageTek Availability Suite Data Replication Device Group

### 1 Log in to a cluster node.

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

### 2 Modify the device group.

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

```
# geopg modify-device-group -p property [-p...] \  
AVSdevicegroupname protectiongroupname
```

`-p property` Specifies the properties of the data replication device group.

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



<i>AVSdevicegroupname</i>	Specifies the name of the new data replication device group.
<i>protectiongroupname</i>	Specifies the name of the protection group that will contain the new data replication device group.

## ▼ How to Delete a Data Replication Device Group From a Sun StorageTek Availability Suite Protection Group

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

Before you remove a data replication device group, ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The protection group is offline on the local cluster and the partner cluster, if the partner cluster can be reached.
- The device group is managed by the protection group.

For information about deleting protection groups, refer to [“How to Delete a Sun StorageTek Availability Suite Protection Group” on page 39](#).

### 1 Log in to a cluster node.

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

### 2 Remove the device group.

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

This command removes the device group from the protection group. This command also disables all volume sets associated with the device group and deletes the lightweight resource group and replication resource group for this device group.

```
# geopg remove-device-group AVSdevicegroupname protectiongroupname
```

*AVSdevicegroupname* Specifies the name of the data replication device group

*protectiongroupname* Specifies the name of the protection group

**Example 2–10** Deleting a Replication Device Group From a Sun StorageTek Availability Suite Protection Group

This example deletes a data replication device group from a Sun StorageTek Availability Suite protection group.

```
# geopg remove avsdg avspg
```

## Replicating the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster

You can replicate the configuration of a protection group to the partner cluster either before or after you configure data replication, resource groups, and resources on both clusters.

### ▼ How to Replicate the Sun StorageTek Availability Suite Protection Group Configuration to a Partner Cluster

**Before You Begin** Before you replicate the configuration of a Sun StorageTek Availability Suite protection group to a partner cluster, ensure that the following conditions are met:

- The protection group is defined on the remote cluster, not on the local cluster.
- The device groups in the protection group on the remote cluster exist on the local cluster.
- The 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 groups is set to `False`. You can view this property by using the `clresourcegroup show` command.

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

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

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

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

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

**1 Log in to phys-newyork-1.**

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

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

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

```
# geopg get -s partnershipname \
AVSprotectiongroup
```

`-s partnershipname` Specifies the name of the partnership from which the protection group configuration information is gathered

`AVSprotectiongroup` Specifies the name of the protection group

---

**Note** – The `geopg get` command replicates Sun Cluster Geographic Edition related entities. For information about how to replicate Sun Cluster entities, see “Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

---

**Example 2–11 Replicating the Sun StorageTek Availability Suite Protection Group to a Partner Cluster**

This example replicates the configuration of `avspg` to `cluster-newyork`.

The configuration of the protection group is retrieved from the remote cluster, in this example `cluster-paris`, and then validated by the data replication subsystem on the local cluster, `cluster-newyork`.

If the validation is successful, the `Configuration` status is set to `OK` and the protection group is created on the local cluster. This protection group contains a device group and an application group that are configured almost identically to the device group and application group on the remote cluster.

If the validation fails, the protection group is not created on the local cluster. Fix the cause of the error and replicate it again.

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

# Activating and Deactivating a Protection Group

When you activate a protection group, it assumes the role that you assigned to it during configuration.

For more information about configuring protection groups, see [“How to Create and Configure a Sun StorageTek Availability Suite Protection Group”](#) on page 34.

## ▼ How to Activate a Sun StorageTek Availability Suite Protection Group

**Before You Begin** You can activate a protection group in the following ways:

- Globally, which activates a protection group on both clusters where the protection group has been configured
- On the primary cluster only
- On a secondary cluster only

When you activate a protection group, the data replication product you are using determines the clusters on which data replication can start. For example, the Sun StorageTek Availability Suite software allows data replication to start only from the primary cluster. So, if you activate a protection group from the secondary cluster, data replication does not start.

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

### 2 Activate the protection group.

This command activates the protection group on the local cluster.

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

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

`-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 and the command performs the following operations on each device group in the protection group:
- Verifies that the role configured for the replication resource is the same as the role of the protection group on the local cluster.
  - Verifies that the role of the volume sets associated with the device group is the same as the role of the protection group on the local cluster.
  - If the role of the protection group on the local cluster is secondary, unmounts the local volumes defined in all volume sets associated with the device group.
  - If the role of the protection group on the local cluster is primary, enables the autosynchronization feature of the Sun StorageTek Availability Suite remote mirror software. Also, resynchronizes the volume sets associated with the device group.

*AVSprotectiongroup* Specifies the name of the protection group.

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

The `geopg start` command performs the following actions if the role of the protection group is primary on the local cluster:

- The command runs a script that is defined in the `RoleChange_ActionCmd`.
- The command brings the application resource groups in the protection group online on the local cluster.
- If the application resource group is a failover type resource group that shares affinities with a device group in the same protection group, the command adds strong, positive affinities and failover delegation between the application resource group and the lightweight resource group.

The application resource group must not have strong, positive affinities with failover delegation. Otherwise, the attempt to add strong, positive affinities with failover delegation with the lightweight resource group will fail.

- The command creates strong dependencies between the `HASStoragePlus` resource in the application resource group and the `HASStoragePlus` resource in the lightweight resource group for this device group.

If the command fails, the `Configuration` status might be set to `Error`, depending on the cause of the failure. The protection group remains deactivated, but data replication might be started and some resource groups might be brought online. Run the `geoadm status` command to obtain the status of your system.

If the `Configuration` status is set to `Error`, revalidate the protection group by using the procedures that are described in [“How to Validate a Sun StorageTek Availability Suite Protection Group” on page 37](#).

### Example 2–12 Activating a Sun StorageTek Availability Suite Protection Group Globally

This example activates a protection group globally.

```
# geopg start -e global avspg
```

### Example 2–13 Activating a Sun StorageTek Availability Suite 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 avspg
```

## ▼ How to Deactivate a Sun StorageTek Availability Suite Protection Group

**Before You Begin** You can deactivate a protection group in the following ways:

- Globally, meaning you deactivate a protection group on both the primary and the secondary cluster where the protection group is configured
- On the primary cluster only
- On the secondary cluster only

The result of deactivating a protection group on primary or secondary cluster depends on the type of data replication you are using. If you are using Sun StorageTek Availability Suite software, data replication can be stopped only from the primary cluster. So, when you deactivate a protection group on the secondary cluster, this deactivate command does not stop data replication.

## 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 “Sun Cluster Geographic Edition Software and RBAC” in *Sun Cluster Geographic Edition System Administration Guide*.

## 2 Deactivate the protection group.

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

When you deactivate a protection group, its application resource groups are also unmanaged.

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

-e *scope*

Specifies the scope of the command.

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

---

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

---

-D

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

If you omit this option, the data replication subsystem and the protection group are both stopped. If the role of the protection group on the local cluster is `primary`, omitting the `-d` option also results in the following actions:

- Removal of resource group affinities and resource dependencies between the application resource groups in the protection group and the internal resource group
- Taking the application resource groups offline and putting them in an Unmanaged state

*protectiongroupname*

Specifies the name of the protection group.

If the role of the protection group is `primary` on the local cluster, the `geopg stop` command disables the autosynchronization of each device group and places the volume sets into logging mode.

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

If the Configuration status is set to Error, revalidate the protection group by using the procedures described in [“How to Validate a Sun StorageTek Availability Suite Protection Group”](#) on page 37.

**Example 2–14** Deactivating a Sun StorageTek Availability Suite Protection Group on All Clusters

This example deactivates a protection group on all clusters.

```
# geopg stop -e global avspg
```

**Example 2–15** Deactivating a Sun StorageTek Availability Suite Protection Group on a Local Cluster

This example deactivates a protection group on the local cluster.

```
# geopg stop -e local avspg
```

**Example 2–16** Stopping Sun StorageTek Availability Suite Data Replication While Leaving the Protection Group Online

This example stops only data replication on a local cluster.

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

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 avspg
```

**Example 2–17** Deactivating a Sun StorageTek Availability Suite Protection Group While Keeping Application Resource Groups Online

This example keeps online two application resource groups, apprg1 and apprg2, while deactivating their protection group, avspg.

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

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

2. Deactivate the protection group.

```
# geopg stop -e global avspg
```



# Resynchronizing a Sun StorageTek Availability Suite Protection Group

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

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

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

## ▼ How to Resynchronize a Sun StorageTek Availability Suite Protection Group

**Before You Begin** The protection group must be deactivated on the cluster where you run the `geopg update` command.

### 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 “Sun Cluster Geographic Edition Software and RBAC” in *Sun Cluster Geographic Edition System Administration Guide*.

### 2 Resynchronize the protection group.

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

#### Example 2–18 Resynchronizing a Sun StorageTek Availability Suite Protection Group

This example resynchronizes a protection group.

```
# geopg update avspg
```

# Checking the Runtime Status of Sun StorageTek Availability Suite Data Replication

You can obtain an overall view of the status of replication, as well as a more detailed runtime status of the Sun StorageTek Availability Suite software from the status of the replication resource groups. The following sections describe the procedures for checking each status.

## Displaying a Sun StorageTek Availability Suite Runtime Status Overview

The status of each Sun StorageTek Availability Suite data replication resource indicates the status of replication on a particular device group. The status of all the resources under a protection group are aggregated in the replication status.

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 a cluster where the protection group is defined.

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

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

```
# geoadm status
```

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

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

#### 3 Check the runtime status of data replication for each Sun StorageTek Availability Suite protection group.

```
# clresource status -v avsdg-rep-rs
```

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

## Displaying a Detailed Sun StorageTek Availability Suite Runtime Status

One replication resource group exists for each protection group. The name of the replication resource group has the following format:

*avsprotectiongroupname*-rep-rg

If you add a Sun StorageTek Availability Suite device group to a protection group, the Sun Cluster Geographic Edition software creates a resource for each device group. This resource monitors the status of replication for its device group. The name of each resource has the following format:

*avsdevicegroupname*-rep-rs

You can monitor the state of the replication resource to give you the overall status of replication. Use the `clresource status` command as follows to obtain the State and Status Message values for the replication status of the device group:

```
# clresource status -v avsdg-rep-rs
```

The State is `Online` while the resource is online.

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

TABLE 2-1 Status and Status Messages of an Online Sun StorageTek Availability Suite Replication Resource Group

Status	Status Message
Faulted	Replication service disabled
Faulted	Incorrect role
Faulted	Volume failed
Faulted	Bitmap failed
Faulted	Queue failed
Faulted	Need sync
Faulted	Need reverse sync
Faulted	Reverse synching
Degraded	Synching

**TABLE 2-1** Status and Status Messages of an Online Sun StorageTek Availability Suite Replication Resource Group *(Continued)*

Status	Status Message
Degraded	Queuing
Degraded	Logging
Online	Replicating

For more details about these values, refer to the *Sun StorageTek Availability Suite Remote Mirror Software Administration and Operations Guide*.

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

# Migrating Services That Use Sun StorageTek Availability Suite Data Replication

---

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

- “Detecting Cluster Failure on a System That Uses Sun StorageTek Availability Suite Data Replication” on page 61
- “Migrating Services That Use Sun StorageTek Availability Suite With a Switchover” on page 62
- “Forcing a Takeover on Systems That Use Sun StorageTek Availability Suite” on page 65
- “Recovering Sun StorageTek Availability Suite Data After a Takeover” on page 68
- “Recovering From a Sun StorageTek Availability Suite Data Replication Error” on page 77

## Detecting Cluster Failure on a System That Uses Sun StorageTek Availability Suite Data Replication

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

### Detecting Primary Cluster Failure

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

The following actions occur when the overall state of a protection group changes to the Unknown state:

- Heartbeat failure is detected by a partner cluster.

- The heartbeat is activated in emergency mode to verify that the heartbeat loss is not transient and that the primary cluster has failed. The heartbeat remains in the OK state during this default timeout interval, while the heartbeat mechanism continues to retry the primary cluster. Only the heartbeat plug-ins appear in the Error state.

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

For more information about logging, see “Viewing the Sun Cluster Geographic Edition Log Messages” in *Sun Cluster Geographic Edition System Administration Guide*.

## Detecting Secondary Cluster Failure

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

During failure detection, the following actions occur:

- Heartbeat failure is detected by a partner cluster.
- The heartbeat is activated in emergency mode to verify that the secondary cluster is dead.
- The cluster notifies the administrator. The system detects all protection groups for which the cluster that failed was acting as secondary. The state of these protection groups becomes Unknown.

# Migrating Services That Use Sun StorageTek Availability Suite With a Switchover

You perform a switchover of a Sun StorageTek Availability Suite protection group when you want to migrate services to the partner cluster in an orderly fashion. A switchover consists of the following:

- Application services are unmanaged on the former primary cluster, `cluster-paris`.  
For a reminder of which cluster is `cluster-paris`, see “Example Sun Cluster Geographic Edition Cluster Configuration” in *Sun Cluster Geographic Edition System Administration Guide*.
- The data replication role is reversed and now continues to run from the new primary, `cluster-newyork`, to the former primary, `cluster-paris`.
- Application services are brought online on the new primary cluster, `cluster-newyork`.

## ▼ How to Switch Over a Sun StorageTek Availability Suite Protection Group From Primary to Secondary

**Before You Begin** For a switchover to occur, data replication must be active between the primary cluster and the secondary cluster. Additionally, the data volumes on the two clusters must be in a synchronized state.

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

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

### 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 “Sun Cluster Geographic Edition Software and RBAC” in *Sun Cluster Geographic Edition System Administration Guide*.

### 2 Initiate the switchover.

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

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

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

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

*protectiongroupname* Specifies the name of the protection group

### Example 3-1 Forcing a Switchover From Primary to Secondary

This example performs a switchover to the secondary cluster.

```
# geopg switchover -f -m cluster-newyork avspg
```

## Actions Performed by the Sun Cluster Geographic Edition Software During a Switchover

When you run the `geopg switchover` command, the software confirms that the volume sets that are associated with the device groups are in the `replicating` state. Then, the software performs the following actions on the original primary cluster:

- Removes affinities and resource dependencies between all the application resource groups in the protection group and the internal resource group, such as the lightweight resource groups
- Takes the application resource groups offline and places them in the Unmanaged state
- Waits for writes to complete
- Unmounts the primary volumes that correspond to the device groups in the protection group
- Stops data replication by placing all volume sets in logging mode
- Reverses the role of all volume sets

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

- Places all volume sets in logging mode
- Reverses the role of all volume sets
- Starts data replication by issuing update synchronization with the autosynchronization feature active
- Runs the script that is defined in the `RoleChange_ActionCmd` property
- Brings all application resource groups online and adds the affinities between the application resource groups and the internal resource groups, such as the lightweight resource group

If the command completes successfully, the secondary cluster, `cluster-newyork`, becomes the new primary cluster for the protection group. The original primary cluster, `cluster-paris`, becomes the new secondary cluster. Volume sets associated with a device group of the protection group have their role reversed according to the role of the protection group on the local cluster. The application resource group is online on the new primary cluster. Data replication from the new primary cluster to the new secondary cluster begins.

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

If the `Configuration` status of the protection group is set to `Error`, revalidate the protection group by using the procedures described in [“How to Validate a Sun StorageTek Availability Suite Protection Group” on page 37](#).



If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures described in [“How to Resynchronize a Sun StorageTek Availability Suite Protection Group” on page 57](#).

## Forcing a Takeover on Systems That Use Sun StorageTek Availability Suite

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

The following steps occur after a takeover is initiated:

- If the former primary cluster, `cluster-paris`, can be reached and the protection group is not locked for notification handling or some other reason, the protection group is deactivated.

For a reminder of which cluster is `cluster-paris`, see “Example Sun Cluster Geographic Edition Cluster Configuration” in *Sun Cluster Geographic Edition System Administration Guide*.

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

---

**Note** – This data might not be consistent with the original primary volumes. Data replication from the new primary cluster, `cluster-newyork`, to the former primary cluster, `cluster-paris`, is stopped.

---

- The protection group is activated without data replication.

For details about the possible conditions of the primary and secondary cluster before and after takeover, see Appendix C, “Takeover Postconditions,” in *Sun Cluster Geographic Edition System Administration Guide*.

The following procedures describe the steps you must perform to force a takeover by a secondary cluster, and how to recover data afterward.

## ▼ How to Force Immediate Takeover of Sun StorageTek Availability Suite Services by a Secondary Cluster

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

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

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

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

### 2 Initiate the takeover.

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

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

*protectiongroupname* Specifies the name of the protection group

### Example 3–2 Forcing a Takeover by a Secondary Cluster

This example forces the takeover of avspg by the secondary cluster, cluster-newyork.

phys-newyork-1 is the first node of the secondary cluster. For a reminder of which node is phys-newyork-1, see “Example Sun Cluster Geographic Edition Cluster Configuration” in *Sun Cluster Geographic Edition System Administration Guide*.

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

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

## Actions Performed by the Sun Cluster Geographic Edition Software During a Takeover

When you run the `geopg takeover` command, the software confirms that the volume sets are in a `Replicating` or `Logging` state on the secondary cluster.

If the original primary cluster, `cluster-paris`, can be reached, the software performs the following actions:

- Removes affinities and resource dependencies between all the application resource groups in the protection group and the internal resource group if the protection group was active
- Takes the application resource groups offline and places them in an `Unmanaged` state
- Unmounts the primary volumes that correspond to the device groups in the protection group
- Stops data replication by placing all volume sets in logging mode
- Reverses the role of all volume sets

On the original secondary cluster, `cluster-newyork`, the software performs the following actions:

- Places all volume sets into logging mode
- Reverses the role of all volume sets
- Runs the script that is specified in the `RoleChange_ActionCmd` property
- If the protection group was active on the original secondary cluster before the takeover, brings all application resource groups online and adds affinities and resource dependencies between the application resource group and the internal resource group

If the command completes successfully, the secondary cluster, `cluster-newyork`, becomes the new primary cluster for the protection group. Volume sets associated with a device group in the protection group have their role reversed according to the role of the protection group on the local cluster. If the protection group was active on the original secondary cluster before the takeover, the application resource groups are brought online on the new primary cluster. If the original primary cluster can be reached, it becomes the new secondary cluster of the protection group. Replication of all volume sets that are associated with the device groups of the protection group is stopped.



---

**Caution** – After a successful takeover, data replication is stopped. If you want to continue to suspend replication, specify the `-n` option when you use the `geopg start` command. This option prevents the start of data replication from the new primary cluster to the new secondary cluster.

---

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

If the `Configuration` status of the protection group is set to `Error`, revalidate the protection group by using the procedures described in [“How to Validate a Sun StorageTek Availability Suite Protection Group” on page 37](#).

If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures described in [“How to Resynchronize a Sun StorageTek Availability Suite Protection Group” on page 57](#).

## Recovering Sun StorageTek Availability Suite Data 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, the services can be brought online again on the original primary by using a process called *failback*.

Sun Cluster Geographic Edition software supports the following two kinds of failback:

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

For a reminder of which clusters are `cluster-paris` and `cluster-newyork`, see “Example Sun Cluster Geographic Edition Cluster Configuration” in *Sun Cluster Geographic Edition System Administration Guide*.

- **Failback-takeover.** During a failback-takeover, applications are brought online again on the original primary cluster and use the current data on the primary cluster. Any updates that occurred on the secondary cluster are discarded.

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

## ▼ How to Resynchronize and Revalidate the Protection Group Configuration

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

**Before You Begin** Before you resynchronize and revalidate the protection group configuration, a takeover has occurred on `cluster-newyork`. The clusters now have the following roles:

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

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

The cluster `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`, deactivate the protection group on the local cluster.

```
# geopg stop -e Local protectiongroupname
```

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

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

`protectiongroupname` Specifies the name of the protection group.

If the protection group is already deactivated, the state of the resource group in the protection group is probably `Error`. The state is `Error` because the application resource groups are managed and offline.

Deactivating the protection group results in the application resource groups no longer being managed, clearing the `Error` state.

#### b. On `cluster-paris`, resynchronize the partnership.

```
# geops update partnershipname
```

`partnershipname` Specifies the name of the partnership

---

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

---

For more information about synchronizing partnerships, see “Resynchronizing a Partnership” in *Sun Cluster Geographic Edition System Administration Guide*.

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

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

```
# geopg update protectiongroupname
```

*protectiongroupname* Specifies the name of the protection group

For more information about synchronizing protection groups, see “Resynchronizing a Sun StorageTek Availability Suite Protection Group” on page 57.

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

```
# geopg validate protectiongroupname
```

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

For more information, see “How to Validate a Sun StorageTek Availability Suite Protection Group” on page 37.

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

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

```
# geopg start -e Global protectiongroupname
```

-e Global Specifies the scope of the command.

By specifying a Global scope, the command operates on both clusters where the protection group is deployed.

*protectiongroupname* Specifies the name of the protection group.



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

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

For more information about the `geopg start` command, see “How to Activate a Sun StorageTek Availability Suite Protection Group” on page 52.

---

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

First, confirm that the state of the protection group on `cluster-newyork` is OK.

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

Refer to the Protection Group section of the output.

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

```
phys-newyork-1# clresource status -v AVSdevicegroupname-rep-rs
```

## ▼ How to Perform a Failback-Switchover on a System That Uses Sun StorageTek Availability Suite Replication

Use this procedure to restart an application on the original primary cluster, `cluster-paris`, after the data on the cluster has been resynchronized with the data on the current primary cluster, `cluster-newyork`.

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 now have the following roles:

- If the original primary cluster, `cluster-paris`, has failed, confirm that the cluster is booted and that the Sun Cluster Geographic Edition infrastructure is enabled on the cluster. For more information about booting a cluster, see “Booting a Cluster” in *Sun Cluster Geographic Edition System Administration Guide*.
- The protection group on `cluster-newyork` has the primary role.
- The protection group on `cluster-paris` has either the primary role or secondary role, depending on whether `cluster-paris` could be reached during the takeover from `cluster-newyork`.

**1 Resynchronize the original primary cluster, `cluster-paris`, with the current primary cluster, `cluster-newyork`.**

The cluster `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`, deactivate the protection group on the local cluster.**

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

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

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

*protectiongroupname* Specifies the name of the protection group.

If the protection group is already deactivated, the state of the resource group in the protection group is probably `Error`. The state is `Error` because the application resource groups are managed and offline.

Deactivating the protection group results in the application resource groups no longer being managed, clearing the `Error` state.

**b. 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 *Sun Cluster Geographic Edition System Administration Guide*.

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

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

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

*protectiongroupname* Specifies the name of the protection group

For more information about synchronizing protection groups, see “Resynchronizing a Sun StorageTek Availability Suite Protection Group” on page 57.

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

A protection group cannot be started when it is in an error state. Ensure that the protection group is not in an error state.

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

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

For more information, see “How to Validate a Sun StorageTek Availability Suite Protection Group” on page 37.



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

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

```
phys-paris-1# geogg start -e Global protectiongroupname
```

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

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

`protectiongroupname` Specifies the name of the protection group.



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

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

For more information about the `geogg start` command, see [“How to Activate a Sun StorageTek Availability Suite Protection Group” on page 52](#).

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

First, confirm that the state of the protection group on `cluster-newyork` is OK.

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

Refer to the Protection Group section of the output.

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

```
phys-newyork-1# clresource status -v AVSdevicegroupname-rep-rs
```

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

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

For more information, see [“How to Switch Over a Sun StorageTek Availability Suite Protection Group From Primary to Secondary” on page 63](#).

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

## 6 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 Sun StorageTek Availability Suite protection group.

```
# clresourcegroup status -v resourcegroupname
# clresource status -v AVSdevicegroupname-rep-rs
```

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

For more information about the runtime status of data replication, see “[Checking the Runtime Status of Sun StorageTek Availability Suite Data Replication](#)” on page 58.

## ▼ How to Perform a Failback-Takeover on a System That Uses Sun StorageTek Availability Suite 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`.

---

**Before You Begin** Before you begin the failback-takeover operation, the clusters have the following roles:

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

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

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

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

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

*partnershipname* Specifies the name of the partnership

---

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

---

For more information about synchronizing partnerships, see “Resynchronizing a Partnership” in *Sun Cluster Geographic Edition System Administration Guide*.

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

If the protection group has been activated, deactivate the protection group by using the `geopg stop` command. For more information about deactivating a protection group, see “How to Deactivate a Sun StorageTek Availability Suite Protection Group” on page 54.

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

*protectiongroupname* Specifies the name of the protection group

For more information about synchronizing protection groups, see “How to Resynchronize a Sun StorageTek Availability Suite Protection Group” on page 57.

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

Ensure that the protection group is not in an error state. A protection group cannot be started when it is in an error state.

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

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

For more information, see “How to Validate a Sun StorageTek Availability Suite Protection Group” on page 37.

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

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

```
phys-paris-1# geopg start -e local -n protectiongroupname
```

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

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

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

---

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

---

*protectiongroupname* Specifies the name of the protection group.

For more information, see [“How to Activate a Sun StorageTek Availability Suite Protection Group” on page 52.](#)

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

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

phys-paris-1# **geopg takeover** [-f] *protectiongroupname*

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

*protectiongroupname* Specifies the name of the protection group

For more information about the geopg takeover command, see [“How to Force Immediate Takeover of Sun StorageTek Availability Suite Services by a Secondary Cluster” on page 66.](#)

The protection group on cluster-paris now has the primary role, and the protection group on cluster-newyork has the secondary role.

#### 5 On cluster-newyork, activate each protection group.

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

phys-newyork-1# **geopg start -e local** [-n] *protectiongroupname*

-e local Specifies the scope of the command.

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

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

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

*protectiongroupname* Specifies the name of the protection group.

For more information about the geopg start command, see [“How to Activate a Sun StorageTek Availability Suite Protection Group” on page 52.](#)

**6 Start data replication.**

To start data replication, activate the protection group on the primary cluster, `cluster-paris`.  
`phys-paris-1# geopg start -e local protectiongroupname`

For more information about the `geopg start` command, see [“How to Activate a Sun StorageTek Availability Suite Protection Group”](#) on page 52.

**7 For each cluster, verify that the protection groups are set correctly and that the application resource group status and the data replication status are okay.**

**a. Verify that the protection group is now primary on `cluster-paris` and secondary on `cluster-newyork`. Run the following command from one node on each cluster:**

```
# geoadm status
```

**b. Check the runtime status of the application resource group and data replication for each Sun StorageTek Availability Suite protection group. Run the following commands from one node on each cluster:**

```
# clresourcegroup status -v resourcegroupname
# clresource status -v AVSdevicegroupname-rep-rs
```

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

For more information about the runtime status of data replication, see [“Checking the Runtime Status of Sun StorageTek Availability Suite Data Replication”](#) on page 58.

## Recovering From a Sun StorageTek Availability Suite Data Replication Error

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

For example, suppose a device group controlled by Sun StorageTek Availability Suite that is called `avsdg` changes to a Volume failed state, `VF`. This state is reflected in the following resource status:

```
Resource Status = "FAULTED"
Resource status message = "FAULTED : Volume failed"
```

---

**Note** – The Resource State remains Online because the probe is still running correctly.

---

Because the resource status has changed, the protection group status also changes. In this case, the local Data Replication state, the Protection Group state on the local cluster, and the overall Protection Group state become Error.

To recover from an error state, complete the relevant steps in the following procedure.

## ▼ How to Recover From a Data Replication Error

- 1 Use the procedures in the Sun StorageTek Availability Suite documentation to determine the causes of the FAULTED state. This state is indicated as VF.**

- 2 Recover from the faulted state by using the Sun StorageTek Availability Suite procedures.**

If the recovery procedures change the state of the device group, this state is automatically detected by the resource and is reported as a new protection group state.

- 3 Revalidate the protection group configuration.**

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

*protectiongroupname* Specifies the name of the Sun StorageTek Availability Suite protection group

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

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

*protectiongroupname* Specifies the name of the Sun StorageTek Availability Suite protection group

# Sun Cluster Geographic Edition Properties for Sun StorageTek Availability Suite

---

This appendix provides the properties of Sun Cluster Geographic Edition data replication device groups.

This appendix contains the following sections:

- “Sun StorageTek Availability Suite Properties” on page 79
- “Sun Cluster Geographic Edition Resource Properties for Sun StorageTek Availability Suite That Must Not Be Changed” on page 80

## Sun StorageTek Availability Suite Properties

The following table describes the Sun StorageTek Availability Suite properties that the Sun Cluster Geographic Edition software defines.

TABLE A-1 Sun StorageTek Availability Suite Properties

Property	Description
Data Replication Property: NodeList (string array)	Lists the host names of the machines that can be primary for the device group in the protection group. Device groups in the protection group must share the same ordered nodelist. This list is comma delimited.  Tuning recommendations: This property can be tuned only when the protection group is offline.  Category: Optional  Default: All nodes in the cluster

TABLE A-1 Sun StorageTek Availability Suite Properties (Continued)

Property	Description
Device Group Property: Enable_volume_set (Boolean)	<p>Defines whether the volume sets that are defined in the file (/var/cluster/geo/avs/&lt;AVSdevicegroupname&gt;-volset.ini) are enabled when the device group is added. Set to either true or false.</p> <p>Tuning recommendations: This property cannot be tuned after the device group is added to a protection group.</p> <p>Category: Optional</p> <p>Default: false</p>
Device Group Property: Local_logical_host (string)	<p>Defines the local logical hostname that is used for the replication of the device group. Do not use an underscore (_) character in the logical hostname.</p> <p>Tuning recommendations: This property cannot be tuned after the device group is added to a protection group.</p> <p>Category: Required</p> <p>Default: None</p>
Data Replication Property: Remote_logical_host (string)	<p>Defines the remote logical hostname that is used for the replication of the device group. Do not use an underscore (_) character in the logical hostname.</p> <p>Tuning recommendations: This property cannot be tuned after the device group is added to a protection group.</p> <p>Category: Required</p> <p>Default: None</p>

## Sun Cluster Geographic Edition Resource Properties for Sun StorageTek Availability Suite That Must Not Be Changed

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

For Sun StorageTek Availability Suite, do not edit the following properties:

- Remote\_logical\_host – Defines the remote logical hostname that is used for the replication of the device group.
- Device\_group – Specifies the Sun Cluster device group that contains the volumes that are being replicated.
- Role – local data replication role.



# Index

---

## A

- activating, protection group, 52-54
- administering
  - data replication with Sun StorageTek Availability Suite, 11-24, 25-60
  - device groups, 44-50
- application resource groups
  - administering, 40-44
  - creating, 41-43
  - removing, 43-44

## C

- configuration summary, 11-12
- configuring
  - device groups, 22
  - HAStoragePlus resource, 23-24
  - local file system, 23-24
  - protection groups, 34-36
  - Sun StorageTek Availability Suite software, 14-16
  - Sun StorageTek Availability Suite volume, 17-18
- creating
  - application resource group, 41-43
  - protection groups, 34-36
  - replication device group, 45-46

## D

- data recovery, 68-77
  - failback-switchover, 71-74

## data recovery (*Continued*)

- failback-takeover, 74-77
- deactivating, protection groups, 54-56
- deleting
  - application resource group, 43-44
  - protection groups, 39-40
  - replication device group, 49-50
- detecting failure, 61-62
- device groups
  - adding to protection group, 45-46
  - administering, 44-50
  - configuring, 22
  - modifying, 48-49
  - removing, 49-50
- DID, with raw-disk device groups, 16-17

## E

- enabling
  - volume set, 18-22
    - manually, 21-22
    - raw device, 20
  - Solaris Volume Manager, 19
  - VERITAS Volume Manager, 19-20

## F

- failback-switchover, 71-74
- failback-takeover, 74-77
- failure
  - detecting, 61-62

failure (*Continued*)

- primary cluster, 61-62
- secondary cluster, 62

**H**

HASStoragePlus resource, configuring, 23-24

**L**

lightweight resource groups, 13  
local file system configuration, 23-24

**M**

migrating services, 61-78

- data recovery after, 68-77
- with a switchover, 62-65
- with a takeover, 65-68

modifying

- protection groups, 36-37
- replication device group, 48-49

**P**

primary cluster

- data recovery, 68-77
- failure detection, 61-62
- switchover, 62-65
- takeover, 65-68

properties, Sun StorageTek Availability Suite, 79-80  
protection groups

- activating, 52-54
- adding application resource group to, 41-43
- adding device group to, 45-46
- configuring, 34-36
- creating, 34-36
- creating while application offline, 26
- creating while application online, 26-33
  - example of, 27-33
- creation strategies, 25-33

protection groups (*Continued*)

- deactivating, 54-56
- deleting, 39-40
- modifying, 36-37
- modifying device group for, 48-49
- removing application resource group, 43-44
- removing device group from, 49-50
- replicating configuration of, 50-51
- resynchronizing, 57
- validating, 37-38

**R**

raw-disk device groups, 16-17  
recovery

- See data recovery
- from replication error, 77-78

replication

- adding device group, 45-46
- initial configuration of, 14-24
- migrating services, 61-78
- modifying device group, 48-49
- protection group configuration, 50-51
- recovering from errors, 77-78
- removing device group, 49-50
- resource groups, 13-14
- runtime status details, 59-60
- runtime status overview, 58
- Sun StorageTek Availability Suite, 11-24, 25-60

resource groups

- application, 40-44
- lightweight, 13
- replication, 13-14
- replication status, 59-60

resynchronizing, protection groups, 57  
runtime status

- replication, 58-60
- state and status messages, 59-60

**S**

secondary cluster

- failure detection, 62

- secondary cluster (*Continued*)
  - switchover, 62-65
  - takeover, 65-68
- Sun Cluster volume, configuring, 17-18
- Sun StorageTek Availability Suite
  - administering data replication with, 11-24, 25-60
  - configuring software, 14-16
  - detecting failure, 61-62
  - initial software configuration, 14-24
  - lightweight resource groups, 13
  - migrating services that use, 61-78
  - properties of, 79-80
  - replication resource groups, 13-14
  - runtime status, 58-60
    - overall, 58
  - supported versions, 7
- Sun StorageTek Availability Suite 4, 7
- Sun StorEdge Availability Suite 3.2.1, 7
- supported replication software, Sun StorageTek Availability Suite, 7
- switchover, 62-65
  - actions performed during, 64-65
  - primary to secondary, 63
- volume set, enabling (*Continued*)
  - VERITAS Volume Manager, 19-20

## T

- takeover, 65-68
  - actions performed during, 67-68
  - data recovery after, 68-77
  - failback-switchover, 71-74
  - failback-takeover, 74-77
  - how to force, 66

## V

- validating, protection groups, 37-38
- volset file, 14-16
- volume set
  - configuring, 17-18
  - enabling, 18-22
    - manually, 21-22
    - raw device, 20
    - Solaris Volume Manager, 19

