

Veritas™ Cluster Server Agent for Hitachi TrueCopy Installation and Configuration Guide

AIX, HP-UX, Linux, Solaris

5.0

Veritas Cluster Server Agent for Hitachi TrueCopy Installation and Configuration Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Agent version: 5.0.03.0

Document version: 5.0.03.0.1

Legal Notice

Copyright © 2008 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE 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. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

www.symantec.com/techsupp/

Contacting Technical Support

Customers with a current maintenance agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/assistance_care.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and maintenance contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to clustering_docs@symantec.com. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting.

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance agreement administration team for your region as follows:

Asia-Pacific and Japan	contractsadmin@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4	
Chapter 1	Introducing the Veritas agent for Hitachi TrueCopy	9
	About the agent for Hitachi TrueCopy	9
	What's new in this release	10
	Supported software and hardware	10
	Typical Hitachi TrueCopy setup in a VCS cluster	12
	Hitachi TrueCopy agent functions	13
	About the Hitachi TrueCopy agent's online function	14
Chapter 2	Installing and removing the agent for Hitachi TrueCopy	17
	Before you install the agent for TrueCopy	17
	Installing the agent for TrueCopy	17
	Configuring LVM on AIX	20
	Configuring LVM on HP-UX	20
	Upgrading the agent for TrueCopy	21
	Removing the agent for TrueCopy	22
Chapter 3	Configuring the agent for Hitachi TrueCopy	23
	Configuration concepts for the Hitachi TrueCopy agent	23
	Resource type definition for the Hitachi TrueCopy agent	23
	Attribute definitions for the Hitachi TrueCopy agent	24
	Sample configuration for the Hitachi TrueCopy agent	25
	Before you configure the agent for TrueCopy	26
	About cluster heartbeats	27
	About configuring a global cluster in an SF Oracle RAC environment	27
	About configuring system zones in replicated data clusters	28
	About preventing split-brain	28
	Configuring the agent for TrueCopy	29
	Configuring the agent manually in a global cluster	29
	Configuring the agent in an SF for Oracle RAC environment	30

	Configuring the agent manually in a replicated data cluster	30
Chapter 4	Testing VCS disaster recovery support with Hitachi TrueCopy	33
	How VCS recovers from various disasters in an HA/DR setup with Hitachi TrueCopy	33
	Failure scenarios in global clusters	34
	Failure scenarios in replicated data clusters	38
	Testing the service group migration	42
	Testing disaster recovery after host failure	43
	Testing disaster recovery after site failure	44
	Performing failback after a host failure or an application failure	45
	Performing failback after a site failure	46
Chapter 5	Setting up a fire drill	49
	About fire drills	49
	Fire drill configurations	50
	Note on the Gold configuration	51
	About the HTCSnap agent	51
	HTCSnap agent functions	51
	Resource type definition for the HTCSnap agent	53
	Attribute definitions for the HTCSnap agent	53
	About the Snapshot attributes	54
	Before you configure the fire drill service group	55
	Configuring the fire drill service group	56
	Creating the fire drill service group using Cluster Manager (Java Console)	56
	Creating the fire drill service group using the Fire Drill SetUp Wizard	58
	Verifying a successful fire drill	59
	Sample configuration for a fire drill service group	60
Index		61

Introducing the Veritas agent for Hitachi TrueCopy

This chapter includes the following topics:

- [About the agent for Hitachi TrueCopy](#)
- [What's new in this release](#)
- [Supported software and hardware](#)
- [Typical Hitachi TrueCopy setup in a VCS cluster](#)
- [Hitachi TrueCopy agent functions](#)

About the agent for Hitachi TrueCopy

The Veritas agent for Hitachi TrueCopy provides support for application failover and recovery. The agent provides this support in environments that use TrueCopy to replicate data between Hitachi arrays.

The agent monitors and manages the state of replicated Hitachi TrueCopy devices that are attached to VCS nodes. The agent ensures that the system that has the TrueCopy resource online also has safe and exclusive access to the configured devices.

You can use the agent in replicated data clusters and in global clusters that run VCS.

The agent supports TrueCopy in all fence levels that are supported on a particular array.

The agent supports different fence levels for different arrays:

Table 1-1 Supported fence levels

Arrays	Supported fence levels
Lightning	data, never, and async
Thunder	data and never

The agent also supports parallel applications, such as Veritas Storage Foundation for Oracle RAC.

The Hitachi TrueCopy agent also supports Hitachi Universal Replicator for asynchronous replication on two sites.

See the following Technical Support TechNote for the latest updates or software issues for this agent:

<http://seer.entsupport.symantec.com/docs/282004.htm>

What's new in this release

The Veritas agent for Hitachi TrueCopy includes the following new or enhanced features:

- Supports firedrill for Veritas Storage Foundation for Oracle RAC in this release.
- Supports the PreSwitch action to ensure that the remote site cluster can come online during a planned failover within a global cluster configuration.

Note: The agent features such as firedrill for SF Oracle RAC and the Preswitch action are supported only on VCS and SF Oracle RAC 5.0 MP3 versions.

Supported software and hardware

The agent for Hitachi TrueCopy supports the following software versions:

Veritas Cluster Server

- VCS 5.0 MP3 on AIX
- VCS 5.0 and 5.0 MP1 on HP-UX 11i v2
- VCS 5.0 MP3 on Red Hat Enterprise Linux
- VCS 5.0 MP3 on SUSE Linux Enterprise Server
- VCS 5.0 MP3 on Solaris SPARC
- VCS 5.0 MP3 on Solaris x64

On Linux, you must also install the 5.0MP3HF1 hotfix besides VCS 5.0 MP3.

See *Veritas Cluster Server Release Notes* for more details on the supported architectures and the operating system versions.

Veritas SF for Oracle RAC

- SF Oracle RAC 5.0 MP3 on AIX
- SF Oracle RAC 5.0 and 5.0 MP1 on HP-UX 11i v2
- SF Oracle RAC 5.0 and 5.0 MP1 on Red Hat Enterprise Linux
- SF Oracle RAC 5.0 and 5.0 MP1 on SUSE Linux Enterprise Server
- SF Oracle RAC 5.0 MP3 on Solaris SPARC
- SF Oracle RAC 5.0 MP3 on Solaris x64

See *Veritas Storage Foundation for Oracle RAC Release Notes* for more details on the supported architectures and the operating system versions.

Veritas Volume Manager

- VxVM 5.0 MP3 on AIX, Linux, and Solaris
- VxVM 5.0 and 5.0 MP1 on HP-UX 11i v2

On HP-UX, Symantec recommends using VxVM 5.0 MP1.

The agent for Hitachi TrueCopy provides support for the following:


- All versions of CCI.
 The agent supports TrueCopy on all microcode levels on all arrays, provided the host, HBA, array combination is in Hitachi's hardware compatibility list.
- Hewlett-Packard XP arrays with TrueCopy rebranded as Continuous Access
 The agent for Hitachi TrueCopy does not support other Hewlett-Packard replication solutions under the Continuous Access umbrella such as Continuous Access Storage Appliance (CASA). The agent only supports Continuous Access XP.

In environments using Veritas Storage Foundation for Oracle RAC, the arrays must support SCSI-3 persistent reservations.

Typical Hitachi TrueCopy setup in a VCS cluster

Figure 1-1 displays a typical cluster setup in a TrueCopy environment.

Figure 1-1 Typical clustering setup for the agent



Clustering in a TrueCopy environment typically consists of the following hardware infrastructure:

- The primary array (array1) has one or more P-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to the Hitachi TrueCopy array that contains the TrueCopy P-VOL devices.
- The secondary array (array2) has one or more S-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to a Hitachi TrueCopy array that contains the TrueCopy S-VOL devices. The S-VOL devices are paired with the P-VOL devices in the P-VOL array. The S-VOL hosts and arrays must be at a significant distance to survive a disaster that may occur at the P-VOL side.
- Network heartbeating between the two data centers to determine their health; this network heartbeating could be LLT or TCP/IP.
- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT.
In a global cluster environment, you must attach all hosts in a cluster to the same Hitachi TrueCopy array.
- In parallel applications like Veritas Storage Foundation for Oracle RAC, all hosts that are attached to the same array must be part of the same GAB membership. Veritas Storage Foundation for Oracle RAC is supported with

TrueCopy only in a global cluster environment and not in a replicated data cluster environment.

Hitachi TrueCopy agent functions

The VCS enterprise agent for Hitachi TrueCopy monitors and manages the state of replicated devices that are attached to VCS nodes.

The agent performs the following functions:

online	<p>If the state of all local devices is read-write enabled, the agent creates a lock file on the local host to indicate that the resource is online. This action makes the devices writable for the application.</p> <p>If one or more devices are not in a writable state, the agent runs the <code>horctakeover</code> command to enable read-write access to the devices.</p> <p>See “About the Hitachi TrueCopy agent's online function” on page 14.</p>
offline	<p>The agent removes the lock file that was created for the resource by the online entry point. The agent does not run any TrueCopy commands because taking the resource offline is not indicative of an intention to give up the devices.</p>
monitor	<p>Verifies the existence of the lock file to determine the resource status. If the lock file exists, the agent reports the status of the resource as online. If the lock file does not exist, the agent reports the status of the resource as offline.</p> <p>The monitor entry point does not examine the state of the devices or the state of the replication link between the arrays.</p>
open	<p>Removes the lock file from the host on which this entry point is called. This functionality prevents potential concurrency violation if the group fails over to another node.</p> <p>Note that the agent does not remove the lock file if the agent starts after the following command:</p> <pre>hastop -force</pre>
clean	<p>Determines whether it is safe to fault the resource if the online entry point fails or times out. The main consideration is whether a management operation was in progress when the online thread timed out and was killed. If a management operation was in progress, it could potentially leave the devices in an unusable state.</p>
info	<p>Reports the current role and status of the devices in the device group. This entry point can be used to verify the device state and to monitor dirty track trends.</p>

action	<p>Resynchronizes the devices from the VCS command line after connectivity failures are detected and corrected.</p> <p>The agent supports the following actions:</p> <ul style="list-style-type: none">■ <code>pairdisplay</code>—Displays information about all devices.■ <code>pairresync</code>—Resynchronizes the S-VOLs.■ <code>pairresync-swaps</code>—Promotes the S-VOLs to P-VOLs and resynchronizes the original P-VOLs.■ <code>localtakeover</code>—Makes the local devices write-enabled.
action\PreSwitch	<p>Ensures that the remote site cluster can come online during a planned failover within a GCO configuration. The VCS engine on the remote cluster invokes the PreSwitch action on all the resources of the remote site during a planned failover using the <code>hagr -switch</code> command. For this, the PreSwitch attribute must be set to 1. The option <code>-nopre</code> indicates that the VCS engine must switch the servicegroup regardless of the value of the PreSwitch service group attribute.</p> <p>If running the PreSwitch action fails, the failover should not occur. This minimizes the application downtime and data loss .</p> <p>For more information on the PreSwitch action and the PreSwitch feature in the VCS engine, refer to the <i>Veritas Cluster Server User's Guide</i>.</p>

About the Hitachi TrueCopy agent's online function

If the state of all local devices is read-write enabled, the agent makes the devices writable by creating a lock file on the local host.

If one or more devices are not in a writable state, the agent runs the `horctakeover` command to enable read-write access to the devices.

For S-VOL devices in any state other than SSWS or SSUS, the agent runs the `horctakeover` command and makes the devices writable. The time required for failover depends on the following conditions:

- The health of the original primary.
- The RAID Manager timeouts as defined in the `horcm` configuration file for the device group.

The agent considers P-VOL devices writable and takes no action other than going online, regardless of their status.

If the S-VOL devices are in the COPY state, the agent runs the `horctakeover` command after one of the following:

- The synchronization from the primary completes.

- The OnlineTimeout period of the entry point expires, in which case the resource faults.

Installing and removing the agent for Hitachi TrueCopy

This chapter includes the following topics:

- [Before you install the agent for TrueCopy](#)
- [Installing the agent for TrueCopy](#)
- [Upgrading the agent for TrueCopy](#)
- [Removing the agent for TrueCopy](#)

Before you install the agent for TrueCopy

Set up your cluster. For information about installing and configuring VCS, see the *Veritas Cluster Server Installation Guide*.

Set up replication and the required hardware infrastructure.

See “[Typical Hitachi TrueCopy setup in a VCS cluster](#)” on page 12.

Installing the agent for TrueCopy

You must install the Hitachi TrueCopy agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster.

These instructions assume that you have already installed VCS or SF for Oracle RAC.

To install the agent on AIX systems

- 1 Determine the device access name of the disc drive.

```
# cd /dev
# lsdev -C -c cdrom
```

The output resembles:

```
cd0 Available 10-60-00-4,0 16 Bit SCSI Multimedia CD-ROM Drive
```

In this example, the CD device access name is cd0.

- 2 Insert the disc into the system's drive.
- 3 Mount the disc.

```
# mkdir -p /cdrom
# mount -V cdrfs -o ro /dev/cd0 /cdrom
```

- 4 Navigate to the location of the agent packages:

```
# cd /cdrom/aix/replication/htc_agent/version/pkgs
```

The variable *version* represents the version of the agent.

- 5 Add the filesets for the software.

```
# installp -ac -d VRTSvcstc.rte.bff VRTSvcstc
```

To install the agent on HP-UX systems

- 1 Insert the disc into the system's drive.
- 2 Create a mount point directory. For example, /cdrom. The directory must have read-write permissions.
- 3 Determine the block device file for the disc drive.

```
# ioscan -fnC disk
```

For example, the listing may indicate the block device is /dev/dsk/c1t2d0.

- 4 Start the Portable File System (PFS).

```
# nohup pfs_mountd &
# nohup pfsd &
```

5 Mount the disc.

```
# /usr/sbin/pfs_mount -t rrip /dev/dsk/c#t#d# /cdrom
```

The variable */c#t#d#* represents the location of the drive.

6 Install the agent software. Type one of the following commands depending on the operating system on the node.

```
HP-UX (PA)      # swinstall -s /cdrom/hpux/replication\  
                /htc_agent/version/PA/depot VRTSvcstc
```

```
HP-UX (IA)      # swinstall -s /cdrom/hpux/replication\  
                /htc_agent/version/IA/depot VRTSvcstc
```

The variable *version* represents the version of the agent.

To install the agent on Linux systems

- 1 Log in as superuser.
- 2 Insert the disc into the system's drive.
- 3 Mount the disc, if the disc does not automatically mount.

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

4 Navigate to the */mnt* directory.

```
# cd /mnt/cdrom
```

5 Navigate to the location of the agent package.

```
# cd linux/generic/replication
```

6 Install the agent software:

```
# rpm -ivh agentrpm
```

The variable *agentrpm* represents the agent package in the *rpms* directory.

To install the agent on Solaris systems**1** Insert the disc into the system's drive.

```
# cd /cdrom/cdrom0
```

2 Navigate to the location of the agent package.

```
# cd solaris/platform/replication/htc_agent  
/version/pkgs/
```

The following are the *platform* values:

- x64
- sparc

The variable *version* represents the version of the agent.

- 3 Install the agent binaries.

```
# pkgadd -d . VRTSvcstc
```

Configuring LVM on AIX

To support failover of the LVM volume groups to the secondary site during a disaster or normal switch, you must have the AIX ODM repository at the secondary populated with the LVM volume group entries. This must be done as part of an initial setup process before VCS starts controlling the replication.

To configure LVM on AIX

- 1 Start the replication. Wait until it is in the synchronized state. Once it is synchronized, split the replication link.
- 2 At the secondary site, run the `chdev -l <diskname> -a pv=yes` command for each disk inside the replicated device group `lvmdg`. This gets the physical volume identity (PVID) from within the disk and updates the ODM with this value. Now, these disks have the same PVIDs as their counterparts at the primary site.
- 3 Run the `importvg -y <vgname> -n <diskname>` command for each volume group.
- 4 Resync the replication and start VCS.

Configuring LVM on HP-UX

To support failover of the LVM volume groups to the secondary site during a disaster or normal switch, create the LVM volume group on the primary site and export the volume group using the following command:

```
vgexport [[-p] [-v] [-s] [-m]/vg04map.map vg04].
```

Copy the map file to the secondary site and then import the volume group on the secondary using the map file. Use the following command:

```
vgimport [[-s] [-v] [-m]] /vg04map.map vg04
```

This must be done as part of an initial setup process before VCS starts controlling the replication.

To configure LVM on HP-UX

- 1 Configure the volume groups on a replicated primary lun.
- 2 Create the resources HTC, LVMGroup, LVMVolume and mount and bring them online on the primary site.
- 3 Then, bring the resources offline on the primary site and online on the secondary. The resources must be successfully brought online on the secondary site.

Upgrading the agent for TrueCopy

You must upgrade the agent on each node in the cluster.

To upgrade the agent software

- 1 Save the VCS configuration and stop the VCS engine.

```
# haconf -dump -makero  
# hastop -all -force
```

- 2 Remove the agent from the node.
See [“Removing the agent for TrueCopy”](#) on page 22.
- 3 Delete the file `/etc/VRTSvcs/conf/config/HTCTypes.cf`.
- 4 Install the current version of the agent.
See [“Installing the agent for TrueCopy”](#) on page 17.
- 5 Copy the file `HTCTypes.cf` from the directory `/etc/VRTSvcs/conf/` to the `/etc/VRTSvcs/conf/config` directory.
- 6 Repeat step 2 through step 5 on each node.
- 7 From a node in the cluster, edit your configuration file `/etc/VRTSvcs/conf/config/main.cf`.
Configure the new attributes, if applicable.
- 8 Verify the configuration

```
# hacf -verify config
```
- 9 Start VCS on local node first.
- 10 Start VCS on other nodes.

Removing the agent for TrueCopy

Before you attempt to remove the agent, make sure the application service group is not online. You must remove the agent from each node in the cluster.

To remove the agent from an AIX cluster

- ◆ Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# installp -u VRTSvcstc.rte
```

To remove the agent from an HP-UX cluster

- ◆ Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# swremove VRTSvcstc
```

To remove the agent from a Linux cluster

- ◆ Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# rpm -e VRTSvcstc
```

To remove the agent from a Solaris cluster

- ◆ Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# pkgrm VRTSvcstc
```

Configuring the agent for Hitachi TrueCopy

This chapter includes the following topics:

- [Configuration concepts for the Hitachi TrueCopy agent](#)
- [Before you configure the agent for TrueCopy](#)
- [Configuring the agent for TrueCopy](#)

Configuration concepts for the Hitachi TrueCopy agent

Review the configuration concepts and failure scenarios for the agent.

Resource type definition for the Hitachi TrueCopy agent

The resource type definition defines the agent in VCS.

```
type HTC (  
    static str ArgList[] = { BaseDir, GroupName, Instance,  
        SplitTakeover, LinkMonitor }  
    static keylist SupportedActions = { pairdisplay, pairresync,  
        pairresync-swaps, localtakeover, vxdiske, vxdctlenable }  
    str BaseDir = "/HORCM/usr/bin"  
    str GroupName  
    int Instance  
    int SplitTakeover  
    int LinkMonitor  
    temp str VCSResLock
```

```
temp str TargetFrozen  
)
```

Attribute definitions for the Hitachi TrueCopy agent

The descriptions of the agent attributes are as follows:

BaseDir	Path to the RAID Manager Command Line interface. Type-dimension: string-scalar Default: /HORCM/usr/bin.
GroupName	Name of the device group that the agent manages. Type-dimension: string-scalar
Instance	The Instance number of the device that the agent manages. Multiple device groups may have the same instance number. Do not define the attribute if the instance number is zero. Type-dimension: string-scalar
SplitTakeover	A flag that determines whether the agent permits a failover to S-VOL devices if the replication link is disconnected, that is, if P-VOL devices are in the PSUE state. See “About the SplitTakeover attribute for the Hitachi TrueCopy agent” on page 25. Type-dimension: integer-scalar Default: 0
LinkMonitor	A flag that defines whether the agent periodically attempts to resynchronize the S-VOL side if the replication link is disconnected. The agent uses the <code>pairresync</code> command to resynchronize arrays. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the <code>pairresync</code> command. Type-dimension: integer-scalar Default: 0
TargetFrozen	For internal use. Do not modify.
VCSResLock	The agent uses the VCSResLock attribute to guarantee serialized management in case of a parallel application. Type-dimension: temporary string

About the SplitTakeover attribute for the Hitachi TrueCopy agent

The SplitTakeover attribute determines whether the agent permits a failover to S-VOL devices if the replication link is disconnected, that is, if P-VOL devices are in the PSUE state.

The default value for this attribute is 0.

The default value indicates that the agent does not permit a failover to S-VOL devices if the P-VOL devices are in the PSUE state. If a failover occurs when the replication link is disconnected, data loss may occur because the S-VOL devices may not be in sync.

In this scenario, the agent attempts to contact the RAID manager at the P-VOL side to determine the status of the arrays. If the P-VOL side is down, the agent attempts to go online.

In a global cluster environment, if the agent at the P-VOL side detects the PSUE state locally, it freezes the service group at the S-VOL side to prevent a failover. The agent unfreezes the service group after the link is restored and the devices are resynchronized.


If a device group is made up of multiple devices, then, in case of a link failure, the state of each device changes on an individual basis. This change is not reflected on the device group level. Only those devices to which an application made a write after a link failure change their state to PSUE. Other devices in the same device group retain their state to PAIR.

Note: Setting LinkMonitor does not affect the SplitTakeover behavior. However you can minimize the time during which the P-VOL is in the PSUE by setting the LinkMonitor attribute.

Sample configuration for the Hitachi TrueCopy agent

[Figure 3-1](#) shows a dependency graph of a VCS service group that has a resource of type HTC.

Figure 3-1 VCS service group with resource type HTC



You can configure a resource of type HTC in the main.cf file as:

```
HTC DG (  
    GroupName = DG  
    Instance = 1  
)
```

Before you configure the agent for TrueCopy

Before you configure the agent, review the following information:

- Review the configuration concepts, which describe the agent's type definition and attributes.
See [“Configuration concepts for the Hitachi TrueCopy agent”](#) on page 23.
- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
See [“Typical Hitachi TrueCopy setup in a VCS cluster”](#) on page 12.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 27.
See [“About preventing split-brain”](#) on page 28.
- Set up system zones in replicated data clusters.
See [“About configuring system zones in replicated data clusters”](#) on page 28.
- Verify that the clustering infrastructure is in place.
 - If you plan to configure the agent in a global cluster, make sure the global service group for the application is configured.

For more information, see the *Veritas Cluster Server User's Guide*.

- If you want to configure the agent in an SF Oracle RAC environment, verify that the SF Oracle RAC global cluster infrastructure is in place. See [“About configuring a global cluster in an SF Oracle RAC environment”](#) on page 27.
- If you plan to configure the agent in a replicated data cluster, make sure the required replication infrastructure is in place and that the application is configured.

For more information, see the *Veritas Cluster Server User's Guide*.

About cluster heartbeats

In a replicated data cluster, ensure robust heartbeating by using dual, dedicated networks over which the Low Latency Transport (LLT) runs. Additionally, you can configure a low-priority heartbeat across public networks.

In a global cluster, VCS sends ICMP pings over the public network between the two sites for network heartbeating. To minimize the risk of split-brain, VCS sends ICMP pings to highly available IP addresses. VCS global clusters also notify the administrators when the sites cannot communicate.

Hitachi arrays do not support a native heartbeating mechanism between the arrays. The arrays send a support message on detecting replication link failure. You can take appropriate action to recover from the failure and to keep the devices in a synchronized state. The TrueCopy agent supports those actions that can automate the resynchronization of devices after a replication link outage is corrected.

About configuring a global cluster in an SF Oracle RAC environment

Global cluster for Oracle RAC requires setting up different components. The tasks required to set up a global cluster are as follows:

- Configure an SF Oracle RAC cluster at the primary site
- Configure an SF Oracle RAC cluster at the secondary site
- Configure replication on clusters at both sites
- Configure VCS service groups for replication
- Test the HA/DR configuration
- Upon successful testing, bring the environment into production

Some SF Oracle RAC HA/DR configuration tasks may require adjustments depending upon your particular starting point, environment, and configuration.

Review the installation requirements and sample cluster configuration files for primary and secondary clusters.

See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for information on requirements, installation instructions, and sample configuration files.

About configuring system zones in replicated data clusters


In a replicated data cluster, you can prevent unnecessary TrueCopy failover or failback by creating system zones. VCS attempts to fail over applications within the same system zone before failing them over across system zones.

Configure the hosts that are attached to an array as part of the same system zone to avoid unnecessary failover.

Figure 3-2 depicts a sample configuration where `hosta` and `hostb` are in one system zone, and `hostc` and `hostd` are in another system zone.

Use the `SystemZones` attribute to create these zones.

Figure 3-2 Example system zone configuration



Global clusters do not require system zones because failover occurs on a remote cluster if all local targets have been exhausted.

About preventing split-brain

Split-brain occurs when all heartbeat links between the primary and secondary hosts are cut. In this situation, each side mistakenly assumes that the other side is down. You can minimize the effects of split-brain by ensuring that the cluster heartbeat links pass through a similar physical infrastructure as the replication

links. When you ensure that both pass through the same infrastructure, if one breaks, so does the other.

Sometimes you cannot place the heartbeats alongside the replication links. In this situation, a possibility exists that the cluster heartbeats are disabled, but the replication link is not. A failover transitions the original P-VOL to S-VOL and vice-versa. In this case, the application faults because its underlying volumes become write-disabled, causing the service group to fault. VCS tries to fail it over to another host, causing the same consequence in the reverse direction. This phenomenon continues until the group comes online on the final node. You can avoid this situation by setting up your infrastructure such that loss of heartbeat links also mean the loss of replication links.

Configuring the agent for TrueCopy

If you plan to configure the agent in a replicated data cluster:

- Make sure the required replication infrastructure is in place. Also make sure that you have configured the application.

You can adapt most clustered applications to a disaster recovery environment by:

- Converting their devices to TrueCopy devices
- Synchronizing the devices
- Adding the Hitachi TrueCopy agent to the service group

See *Veritas Cluster Server User's Guide* for more information.

After configuration, the application service group must follow the dependency diagram.

See [“Sample configuration for the Hitachi TrueCopy agent”](#) on page 25.

Configuring the agent manually in a global cluster

Configuring the agent manually in a global cluster involves the following tasks:

To configure the agent in a global cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 If the agent resource type (HTC) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types** and select:
`/etc/VRTSvcs/conf/HTCtypes.cf.`
- 3 Click **Import**.
- 4 Save the configuration.

- 5 Add a resource of type HTC at the bottom of the service group.
Link the VMDg and HTC resources so that the VMDg resources depend on HTC.
- 6 Configure the attributes of the HTC resource.
- 7 If the service group is not configured as a global group, configure the service group using the Global Group Configuration Wizard.
See the *Veritas Cluster Server User's Guide* for more information.
- 8 Change the ClusterFailOverPolicy from the default, if necessary. Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.
- 9 Repeat step 5 through step 8 for each service group in each cluster that uses replicated data.
- 10 The configuration must be identical on all cluster nodes, both primary and disaster recovery.

Configuring the agent in an SF for Oracle RAC environment

To configure the agent to manage the volumes that Veritas Storage Foundation for Oracle RAC uses, do the following:

To configure the agent in a Storage Foundation for Oracle RAC environment:

- 1 Configure the SupportedActions attribute for the CVMVolDg resource.
- 2 Add the following keys to the list: import, deport, vxdctlenable.

Note that SupportedActions is a resource type attribute and defines a list of action tokens for the resource.

Configuring the agent manually in a replicated data cluster

Configuring the agent manually in a replicated data cluster involves the following tasks:

To configure the agent in a replicated data cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 If the agent resource type (HTC) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types** and select:
`/etc/VRTSvcs/conf/HTCTypes.cf.`
- 3 Click **Import**.
- 4 Save the configuration.

- 5** In each service group that uses replicated data, add a resource of type HTC at the top of the service group.

Link the VMDg and HTC resources so that VMDg resources depend on Hitachi Truecopy.
- 6** Configure the attributes of the HTC resource. Note that some attributes must be localized to reflect values for the hosts that are attached to different arrays.
- 7** Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array.

Testing VCS disaster recovery support with Hitachi TrueCopy

This chapter includes the following topics:

- [How VCS recovers from various disasters in an HA/DR setup with Hitachi TrueCopy](#)
- [Testing the service group migration](#)
- [Testing disaster recovery after host failure](#)
- [Testing disaster recovery after site failure](#)
- [Performing failback after a host failure or an application failure](#)
- [Performing failback after a site failure](#)

How VCS recovers from various disasters in an HA/DR setup with Hitachi TrueCopy

This section covers the failure scenarios and how VCS responds to the failures for the following DR cluster configurations:

Global clusters When a site-wide global service group or system faults occur, VCS failover behavior depends on the value of the ClusterFailOverPolicy attribute for the faulted global service group. The VCS agent for Hitachi TrueCopy ensures safe and exclusive access to the configured Hitachi TrueCopy devices.

See [“Failure scenarios in global clusters”](#) on page 34.

Replicated data clusters When service group or system faults occur, VCS failover behavior depends on the value of the AutoFailOver attribute for the faulted service group. The VCS agent for Hitachi TrueCopy ensures safe and exclusive access to the configured Hitachi TrueCopy devices.

See [“Failure scenarios in replicated data clusters”](#) on page 38.

See the *Veritas Cluster Server User's Guide* for more information on the DR configurations and the service group attributes.

Failure scenarios in global clusters

[Table 4-1](#) lists the failure scenarios in a global cluster configuration, and describes the behavior of VCS and the agent in response to the failure.

Table 4-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy

Failure	Description and VCS response
Application failure	<p>Application cannot start successfully on any hosts at the primary site.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the global group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the ClusterFailOverPolicy service group attribute: <ul style="list-style-type: none"> ■ Auto or Connected—VCS automatically brings the faulted global group online at the secondary site. ■ Manual—No action. You must bring the global group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ The agent does the following: <ul style="list-style-type: none"> ■ Swaps the P-VOL/S-VOL role of each device in the device group. ■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site. <p>See “Performing failback after a host failure or an application failure” on page 45.</p>

Table 4-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy (*continued*)

Failure	Description and VCS response
Host failure	<p>All hosts at the primary site fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the cluster fault. ■ Does the following based on the ClusterFailOverPolicy service group attribute: <ul style="list-style-type: none"> ■ Auto—VCS automatically brings the faulted global group online at the secondary site. ■ Manual or Connected—No action. You must bring the global group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ The agent does the following: <ul style="list-style-type: none"> ■ Swaps the P-VOL/S-VOL role of each device in the device group. ■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site. <p>See “Performing failback after a host failure or an application failure” on page 45.</p>
Site failure	<p>All hosts and the storage at the primary site fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the cluster fault. ■ Does the following based on the ClusterFailOverPolicy service group attribute: <ul style="list-style-type: none"> ■ Auto—VCS automatically brings the faulted global group online at the secondary site. ■ Manual or Connected—No action. You must bring the global group online at the secondary site. <p>Agent response: The agent does the following based on the SplitTakeover attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>horctakeover</code> command to make the HTC devices write-enabled. The HTC devices go into the SSWS (Suspend for Swapping with S-VOL side only) state. If the original primary site is restored, you must execute the <code>pairresync-swaps</code> action on the secondary site to establish reverse replication. ■ 0—The agent faults the HTC resource. <p>See “Performing failback after a site failure” on page 46.</p>

Table 4-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy (*continued*)

Failure	Description and VCS response
Replication link failure	<p>Replication link between the arrays at the two sites fails.</p> <p>The volume state on the primary site becomes PSUE.</p> <p>VCS response: No action.</p> <p>Agent response: The agent does the following based on the LinkMonitor attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 1—When the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the <code>pairresync</code> command. ■ 0—No action. <p>If the value of the LinkMonitor attribute is not set to 1, you must manually resynchronize the HTC devices after the link is restored.</p> <p>To manually resynchronize the HTC devices after the link is restored:</p> <ol style="list-style-type: none"> 1 Before you resync the S-VOL device, you must split off the Shadow Image device from the S-VOL device at the secondary site. 2 You must initiate resync of S-VOL device using the agent's update action. 3 After P-VOL and S-VOL devices are in sync, reestablish the mirror relationship between the Shadow Copy and the S-VOL devices. <p>If you initiate a failover to the secondary site when resync is in progress, the online function of the Hitachi TrueCopy agent waits for the resync to complete and then initiates a takeover of the S-VOL devices.</p> <p>Note: If you did not configure Shadow Copy devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Symantec recommends configuring Shadow Copy devices at both the sites.</p>

Table 4-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy (*continued*)

Failure	Description and VCS response
Network failure	<p>The network connectivity (wide-area connection (WAC) and the replication link) between the sites fails.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ VCS at each site concludes that the remote cluster has faulted. ■ Does the following based on the ClusterFailOverPolicy service group attribute: <ul style="list-style-type: none"> ■ Manual or Connected—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue. ■ Auto—VCS brings the global group online at the secondary site which may lead to a site-wide split brain. This causes data divergence between the devices on the primary and the secondary arrays. <p>When the network (wac and replication) connectivity restores, you must manually resync the data.</p> <p>Note: Symantec recommends that the value of the ClusterFailOverPolicy attribute is set to Manual for all global groups to prevent unintended failovers due to transient network failures.</p> <p>To resynchronize the data after the network link is restored:</p> <ol style="list-style-type: none"> 1 Take the service groups offline at both the sites. 2 Manually resync the data. <p>Depending on the site whose data you want to retain use the <code>pairresync</code> or the <code>pairresync swap</code> commands.</p> <ol style="list-style-type: none"> 3 Bring the service group online on one of the sites. <p>Agent response: Similar to the site failure</p>
Storage failure	<p>The array at the primary site fails.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the global group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the ClusterFailOverPolicy service group attribute: <ul style="list-style-type: none"> ■ Auto or Connected—VCS automatically brings the faulted global group online at the secondary site. ■ Manual—No action. You must bring the global group online at the secondary site. <p>Agent response: The agent does the following based on the SplitTakeover attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>horctakeover</code> command to make the HTC devices write-enabled. The S-VOL devices go into the SSWS state. ■ 0—The agent faults the HTC resource.

Failure scenarios in replicated data clusters

Table 4-2 lists the failure scenarios in a replicated data cluster configuration, and describes the behavior of VCS and the agent in response to the failure.

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy

Failure	Description and VCS response
Application failure	<p>Application cannot start successfully on any hosts at the primary site.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ The agent does the following: <ul style="list-style-type: none"> ■ Swaps the P-VOL/S-VOL role of each device in the device group. ■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site. <p>See “Performing failback after a host failure or an application failure” on page 45.</p>
Host failure	<p>All hosts at the primary site fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ The agent does the following: <ul style="list-style-type: none"> ■ Swaps the P-VOL/S-VOL role of each device in the device group. ■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site. <p>See “Performing failback after a host failure or an application failure” on page 45.</p>

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy (*continued*)

Failure	Description and VCS response
Site failure	<p>All hosts and the storage at the primary site fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site. <p>Agent response: The agent does the following based on the SplitTakeover attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>horctakeover</code> command to make the HTC devices write-enabled. The HTC devices go into the SSWS (Suspend for Swapping with S-VOL side only) state. If the original primary site is restored, you must execute the <code>pairresync-swaps</code> action on the secondary site to establish reverse replication. ■ 0—The agent faults the HTC resource. <p>See “Performing failback after a site failure” on page 46.</p>

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy (*continued*)

Failure	Description and VCS response
Replication link failure	<p>Replication link between the arrays at the two sites fails.</p> <p>VCS response: No action.</p> <p>Agent response: The agent does the following based on the LinkMonitor attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 1—When the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the <code>pairresync</code> command. ■ 0—No action. <p>If the value of the LinkMonitor attribute is not set to 1, you must manually resynchronize the HTC devices after the link is restored.</p> <p>To manually resynchronize the HTC devices after the link is restored:</p> <ol style="list-style-type: none"> 1 Before you resync the S-VOL device, you must split off the Shadow Image device from the S-VOL device at the secondary site. 2 You must initiate resync of S-VOL device using the agent's update action. 3 After P-VOL and S-VOL devices are in sync, reestablish the mirror relationship between the Shadow Copy and the S-VOL devices. <p>If you initiate a failover to the secondary site when resync is in progress, the online function of the Hitachi TrueCopy agent waits for the resync to complete and then initiates a takeover of the R2 devices.</p> <p>Note: If you did not configure Shadow Copy devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Symantec recommends configuring Shadow Copy devices at both the sites.</p> <p>If you did not configure BCV devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Symantec recommends configuring BCV devices at both the sites.</p>

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy (*continued*)

Failure	Description and VCS response
Network failure	<p>The LLT and the replication links between the sites fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ VCS at each site concludes that the nodes at the other site have faulted. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 2—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue. ■ 1—VCS brings the service group online at the secondary site which leads to a cluster-wide split brain. This causes data divergence between the devices on the arrays at the two sites. When the network (LLT and replication) connectivity is restored, VCS takes all the service groups offline on one of the sites and restarts itself. This action eliminates concurrency violation where in the same group is online at both the sites. After taking the service group offline, you must manually resync the data using the <code>symrdf establish</code> or the <code>symrdf restore</code> commands. <p>Note: Symantec recommends that the value of the AutoFailOver attribute is set to 2 for all service groups to prevent unintended failovers due to transient network failures.</p> <p>To resynchronize the data after the network link is restored:</p> <ol style="list-style-type: none"> 1 Take the service groups offline at both the sites. 2 Manually resync the data. <p>Depending on the site whose data you want to retain use the <code>pairresync</code> or the <code>pairresync swap</code> commands.</p> <ol style="list-style-type: none"> 3 Bring the service group online on one of the sites. <p>Agent response: Similar to the site failure</p>
Storage failure	<p>The array at the primary site fails.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site. <p>Agent response: The agent does the following based on the SplitTakeover attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>horctakeover</code> command to make the HTC devices write-enabled. The S-VOL devices go into the SSWS state. ■ 0—The agent faults the HTC resource.

Testing the service group migration

After you configure the VCS agent for Hitachi TrueCopy, verify that the service group can migrate to hosts across the sites. Depending on your DR configuration, perform one of the following procedures.

To test service group migration in global cluster setup

- 1 Fail over the global group from the primary site to the secondary site.

Perform the following steps:

- Switch the service group from the primary site to any node in the secondary site.

```
hagrp -switch global_group -any -clus cluster_name
```

VCS brings the global group online on a node at the secondary site.

- Verify that the HTC devices at the secondary site are write-enabled, and the device state is SSWS.

- 2 Fail back the global group from the secondary site to the primary site.

Perform the following steps:

- Switch the service group from the secondary site to any node in the primary site.

```
hagrp -switch global_group -any -clus cluster_name
```

VCS brings the service group online on a node at the primary site.

- Verify that the HTC devices at the primary site are write-enabled, and the device state is PAIR.

To test service group migration in replicated data cluster setup

- 1 Fail over the service group from the primary site to the secondary site.

Perform the following steps:

- Switch the service group from the primary site to any node in the secondary site.

```
hagrp -switch service_group -to sys_name
```

VCS brings the service group online on a node at the secondary site.

- Verify that the HTC devices at the secondary site are write-enabled, and the device state is SSWS.

- 2 Fail back the service group from the secondary site to the primary site.

Perform the following steps:

- Switch the service group from the secondary site to any node in the primary site.

```
hagrp -switch service_group -to sys_name
```

VCS brings the service group online on a node at the primary site.

- Verify that the HTC devices at the primary site are write-enabled, and the device state is PAIR.

Testing disaster recovery after host failure

Review the details on host failure and how VCS and the agent for Hitachi TrueCopy behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 34.

See [“Failure scenarios in replicated data clusters”](#) on page 38.

Depending on the DR configuration, perform one of the following procedures to test how VCS recovers after all hosts at the primary site fail.

To test disaster recovery for host failure in global cluster setup

- 1 Halt the hosts at the primary site.

The value of the ClusterFailOverPolicy attribute for the faulted global group determines the VCS failover behavior.

- Auto—VCS brings the faulted global group online at the secondary site.
- Manual or Connected—You must bring the global group online at the secondary site.

On a node in the secondary site, run the following command:

```
hagrp -online -force global_group -any
```

- 2 Verify that the HTC devices at the secondary site are write-enabled, and the device state is SSWS.
- 3 Verify that the global group is online at the secondary site.

```
hagrp -state global_group
```

To test disaster recovery for host failure in replicated data cluster setup

- 1 Halt the hosts at the primary site.

The value of the `AutoFailOver` attribute for the faulted service group determines the VCS failover behavior.

- 1—VCS brings the faulted service group online at the secondary site.
- 2—You must bring the service group online at the secondary site. On a node in the secondary site, run the following command:

```
hagrp -online service_group -to sys_name
```

- 2 Verify that the HTC devices at the secondary site are write-enabled, and the device state is SSWS.
- 3 Verify that the service group is online at the secondary site.

```
hagrp -state global_group
```

Testing disaster recovery after site failure

Review the details on site failure and how VCS and the agent for Hitachi TrueCopy behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 34.

See [“Failure scenarios in replicated data clusters”](#) on page 38.

Depending on the DR configuration, perform one of the following procedures to test the disaster recovery in the event of site failure.

To test disaster recovery for site failure in global cluster setup

- 1 Halt all hosts and the arrays at the primary site.

If you cannot halt the array at the primary site, then disable the replication link between the two arrays.

The value of the `ClusterFailOverPolicy` attribute for the faulted global group determines the failover behavior of VCS.

- Auto—VCS brings the faulted global group online at the secondary site.
- Manual or Connected—You must bring the global group online at the secondary site. On a node in the secondary site, run the following command:

```
hagrp -online -force global_group -any
```

- 2 Verify that the HTC devices at the secondary site are write-enabled, and the device state is SSUS.
- 3 Verify that the global group is online at the secondary site.

```
hagrp -state global_group
```

To test disaster recovery for site failure in replicated data cluster setup

- 1 Halt all hosts and the arrays at the primary site.

If you cannot halt the array at the primary site, then disable the replication link between the two arrays.

The value of the AutoFailOver attribute for the faulted service group determines the VCS failover behavior.

- 1—VCS brings the faulted service group online at the secondary site.
- 2—You must bring the service group online at the secondary site.
On a node in the secondary site, run the following command:

```
hagrp -online service_group -sys sys_name
```

- 2 Verify that the HTC devices at the secondary site are write-enabled, and the device state is SSWS.
- 3 Verify that the service group is online at the secondary site.

```
hagrp -state global_group
```

Performing failback after a host failure or an application failure

Review the details on host failure and application failure and how VCS and the agent for Hitachi TrueCopy behave in response to these failures.

See [“Failure scenarios in global clusters”](#) on page 34.

See [“Failure scenarios in replicated data clusters”](#) on page 38.

After the hosts at the primary site are restarted, you can perform a failback of the service group to the primary site. Depending on your DR configuration, perform one of the following procedures.

To perform failback after a host failure or an application failure in global cluster

- 1 Switch the global group from the secondary site to any node in the primary site.

```
hagr -switch global_group -any -clus cluster_name
```

VCS brings the global group online on a node at the primary site.

The VCS agent for Hitachi TrueCopy does the following based on

- 1—Write enables the devices at the primary site, swaps the P-VOL/S-VOL personality of each device in the device group or the consistency group, and restarts replication from P-VOL devices on the primary site to the S-VOL devices at the secondary site.
 - 0—Issues the `pairresync swaps` command to resync the P-VOL devices and to write enable the P-VOL devices at the primary site.
- 2 Verify that the HTC devices at the primary site are write-enabled and the device state is PAIR.

To perform failback after a host failure or an application failure in replicated data cluster

- 1 Switch the service group from the secondary site to any node in the primary site.

```
hagr -switch service_group -to sys_name
```

VCS brings the service group online on a node at the primary site.

The VCS agent for Hitachi TrueCopy does the following based on

- 1—Write enables the devices at the primary site, swaps the P-VOL/S-VOL personality of each device in the device group or the consistency group, and restarts replication from P-VOL devices on the primary site to the S-VOL devices at the secondary site.
 - 0—Issues the `pairresync swaps` command to resync the P-VOL devices and to write enable the P-VOL devices at the primary site.
- 2 Verify that the HTC devices at the primary site are write-enabled and the device state is PAIR.

Performing failback after a site failure

After a site failure at the primary site, the hosts and the storage at the primary site are down. VCS brings the service group online at the secondary site and the

Hitachi TrueCopy agent write enables the S-VOL devices. The device state is SSUS. Review the details on site failure and how VCS and the agent for Hitachi TrueCopy behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 34.

See [“Failure scenarios in replicated data clusters”](#) on page 38.

When the hosts and the storage at the primary site are restarted and the replication link is restored, the HTC devices attain PSUS/SSWS state at both the sites. The devices are write-enabled at both sites. You can now perform a failback of the service group to the primary site.

To perform failback after a site failure in global cluster

- 1 Take the global group offline at the secondary site. On a node at the secondary site, run the following command:

```
hagrp -offline global_group -any
```

- 2 Resync the devices using the `pairresync` command.

After the resync is complete, the device state is PAIR at both the sites. The devices are write-enabled at the primary site and write-disabled at the secondary site.

- 3 Bring the global group online at the primary site. On a node in the primary site, run the following command:

```
hagrp -online global_group -any
```

To perform failback after a site failure in replicated data cluster

- 1 Take the service group offline at the secondary site. On a node in the secondary site, run the following command:

```
hagrp -offline service_group -sys sys_name
```

- 2 Resync the devices using the `pairresync` command.

After the resync is complete, the device state is PAIR at both the sites. The devices are write-enabled at the primary site and write-disabled at the secondary site.

- 3 Bring the service group online at the primary site. On a node in the primary site, run the following command:

```
hagrp -online service_group -sys sys_name
```


Setting up a fire drill

This chapter includes the following topics:

- [About fire drills](#)
- [Fire drill configurations](#)
- [About the HTCSnap agent](#)
- [Before you configure the fire drill service group](#)
- [Configuring the fire drill service group](#)
- [Verifying a successful fire drill](#)
- [Sample configuration for a fire drill service group](#)

About fire drills

A fire drill procedure verifies the fault-readiness of a disaster recovery configuration. This procedure is done without stopping the application at the primary site and disrupting user access.

A fire drill is performed at the secondary site using a special service group for fire drills. The fire drill service group is identical to the application service group, but uses a fire drill resource in place of the replication agent resource. The fire drill service group uses a copy of the data that is used by the application service group.

In clusters employing Hitachi TrueCopy, the HTCSnap resource manages the replication relationship during a fire drill.

Bringing the fire drill service group online demonstrates the ability of the application service group to come online at the remote site when a failover occurs.

The HTCSnap agent supports fire drills for storage devices that are managed using Veritas Volume Manager, which is a component of Veritas Storage

Foundation. The agent supports fire drills in a Storage Foundation for Oracle RAC environment.

Fire drill configurations

VCS supports the following fire drill configurations for the agent:

- | | |
|--------|---|
| Gold | <p>Runs the fire drill on a snapshot of the target array. The replicated device keeps receiving writes from the primary.</p> <p>Symantec recommends this configuration because it does not affect production recovery.</p> <p>In the Gold configuration, VCS does the following:</p> <ul style="list-style-type: none">■ Suspends replication to get a consistent snapshot.■ Takes a snapshot of the target array on a ShadowImage device.■ Resumes replication.■ Modifies the disk group name in the snapshot.■ Brings the fire drill service group online using the snapshot data. <p>For Gold configurations, you must use Volume Manager to import and deport the storage.</p> <p>You can use the Gold configuration only with ShadowImage pairs created without the <code>-m noread</code> flag to the <code>paircreate</code> command.</p> |
| Silver | <p>VCS takes a snapshot, but does not run the fire drill on the snapshot data. VCS breaks replication and runs the fire drill on the replicated target device. If a disaster occurs while resynching data after running the fire drill, you must switch to the snapshot for recovery.</p> <p>In the Silver configuration, VCS does the following:</p> <ul style="list-style-type: none">■ Suspends replication to get a consistent snapshot.■ Takes a snapshot of the target array on a ShadowImage device.■ Resumes replication■ Modifies the disk name and the disk group name in the snapshot.■ Brings the fire drill service group online using the data on the target array; the agent does not use the snapshot data for the fire drill. <p>You can use the Silver configuration only with ShadowImage pairs created with the <code>-m noread</code> flag to the <code>paircreate</code> command.</p> |

Bronze

VCS breaks replication and runs the fire drill test on the replicated target. VCS does not take a snapshot in this configuration.

If a disaster occurs while resynching data after the test, it may result in inconsistent data as there is no snapshot data.

In the Bronze configuration, VCS does the following:

- Suspends replication.
- Brings the fire drill service group online using the data on the target array.

Note on the Gold configuration

You must perform several steps for a successful Gold configuration fire drill.

To create a Gold configuration fire drill

- 1 Bring the fire drill service group online in the DR cluster.
- 2 Take the fire drill service group offline in the DR cluster.
- 3 Bring the application group online in the DR cluster.
- 4 Migrate the application group (or failover/manually switch it) to the production cluster.
- 5 Bring the application group online on to the production cluster.

About the HTCSnap agent

The HTCSnap agent is the fire drill agent for Hitachi TrueCopy. The agent manages the replication relationship between the source and target arrays when running a fire drill. Configure the HTCSnap resource in the fire drill service group, in place of the HTC resource.

HTCSnap agent functions

The HTCSnap agent performs the following functions:

online	<p>Gold Configuration</p> <ul style="list-style-type: none">■ Suspends replication between the source and the target arrays.■ Takes a local snapshot of the target LUN.■ Resumes the replication between the arrays.■ Takes the fire drill service group online by mounting the replication target LUN.■ Creates a lock file to indicate that the resource is online. <p>Silver Configuration</p> <ul style="list-style-type: none">■ Suspends replication between the source and the target arrays.■ Takes a local snapshot of the target LUN.■ Takes the fire drill service group online by mounting the target LUN.■ Creates a lock file to indicate that the resource is online. <p>Bronze Configuration</p> <ul style="list-style-type: none">■ Suspends replication between the source and the target arrays.■ Takes the fire drill service group online using the target array.■ Creates a lock file to indicate that the resource is online.
offline	<p>Gold Configuration</p> <ul style="list-style-type: none">■ Destroys the snapshot by synchronizing data between the target array and the device on which snapshot was taken.■ Removes the lock file created by the online function. <p>Silver Configuration</p> <ul style="list-style-type: none">■ Resumes replication between the source and the target arrays.■ Synchronizes data between the target array and the device on which the snapshot was taken. Destroys the snapshot of the target array after the data is synchronized. <p>Bronze Configuration</p> <ul style="list-style-type: none">■ Resumes the replication between the source and the target arrays.■ Removes the lock file created by the Online operation.
monitor	Verifies the existence of the lock file to make sure the resource is online.
clean	Restores the state of the LUNs to their original state after a failed online function.
action	For internal use.

Resource type definition for the HTCSnap agent

Following is the resource type definition for the HTCSnap agent:

```
type HTCSnap (  
    static keylist RegList = { MountSnapshot, UseSnapshot }  
    static keylist SupportedActions = { clearvm }  
    static str ArgList[] = { TargetResName, MountSnapshot,  
        UseSnapshot, RequireSnapshot, ShadowInstance }  
    str TargetResName  
    int ShadowInstance  
    int MountSnapshot  
    int UseSnapshot  
    int RequireSnapshot  
    temp str Responsibility  
    temp str FDFFile  
)
```

Attribute definitions for the HTCSnap agent

To customize the behavior of the HTCSnap agent, configure the following attributes:

ShadowInstance	<p>The instance number of the ShadowInstance P-VOL group. The P-VOL group must include one of the following:</p> <ul style="list-style-type: none">■ The same LUNs as in the TrueCopy S-VOL group (if taking snapshots of replicated data)■ The same LUNs as in the VxVM disk group (if taking snapshots of non-replicated data). <p>Type-dimension: integer-scalar</p>
TargetResName	<p>Name of the resource managing the LUNs that you want to take snapshot of. Set this attribute to the name of the HTC resource if you want to take a snapshot of replicated data. Set this attribute to the name of the DiskGroup resource if the data is not replicated.</p> <p>For example, in a typical Oracle setup, you might replicate data files and redo logs, but you may choose to avoid replicating temporary tablespaces. The temporary tablespace must still exist at the DR site and may be part of its own disk group.</p> <p>Type-dimension: string-scalar</p>

UseSnapshot	<p>Specifies whether the HTCSnap resource takes a local snapshot of the target array. Set this attribute to 1 for Gold and Silver configurations. For Bronze, set this attribute to 0.</p> <p>Type-Dimension: integer-scalar</p> <p>See “About the Snapshot attributes” on page 54.</p>
RequireSnapshot	<p>Specifies whether the HTCSnap resource must take a snapshot before coming online.</p> <p>Set this attribute to 1 if you want the resource to come online only after it succeeds in taking a snapshot.</p> <p>Set this attribute to 0 if you do want the resource to come online even if it fails to take a snapshot. Setting this attribute to 0 creates the Bronze configuration.</p> <p>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if UseSnapshot is set to 1.</p>
MountSnapshot	<p>Specifies whether the resource uses the snapshot to bring the service group online. Set this attribute to 1 for Gold configuration. For Silver and Bronze configurations, set the attribute to 0.</p> <p>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if UseSnapshot is set to 1.</p>
Responsibility	<p>Do not modify. For internal use only.</p> <p>Used by the agent to keep track of resynchronizing snapshots.</p> <p>Type-Dimension: temporary string</p>
FDFile	<p>Do not modify. For internal use only.</p> <p>Used by the agent to store the absolute pathname to the file with the latest fire drill report on the local system.</p> <p>Type-Dimension: temporary string</p>

About the Snapshot attributes

The UseSnapshot, MountSnapshot, and RequireSnapshot attributes define the fire drill configuration.

[Table 5-1](#) lists the snapshot attribute values for fire drill configurations:

Table 5-1 Snapshot attribute values for fire drill configurations

Attribute	Gold	Silver	Bronze
MountSnapshot	1	0	0
UseSnapshot	1	1	0

Setting the RequireSnapshot attribute to 0 enables a Gold or Silver configuration to run in the Bronze mode if the snapshot operation fails.

Before you configure the fire drill service group

Before you configure the fire drill service group, follow the steps below:

- Make sure the application service group is configured with a HTC resource.
- Make sure the infrastructure to take snapshots is properly configured. This process involves creating the Shadow Image pairs.
- Make sure the infrastructure to configure hardware snapshots is properly configured.
- If you plan to use Gold or Silver configuration, make sure ShadowImage for TrueCopy is installed and configured at the target array.
- For the Gold configuration, you must use Veritas Volume Manager to import and deport the storage.
- You can use the Silver configuration only with ShadowImage pairs that are created with the `-m noread` flag to the `paircreate` command. A fire drill uses the `-E` flag to split the pairs, which requires a 100% resynchronization. The Silver mode that preserves the snapshots as `noread` after a split.
- The name of the ShadowImage device group must be the same as the replicated device group for both replicated and non-replicated LUNs that are to be snapshot. The instance number may be different.
- Make sure the HORC instance managing the S-VOLs runs continuously; the agent does not start this instance.
- For non-replicated devices:
 - You must use Veritas Volume Manager.
On HP-UX, you must use Veritas Volume Manager 5.0 MP1.
 - For Gold configuration to run without the Bronze mode, set the RequireSnapshot attribute to 1.

- Add vxdctlenable action in the list of SupportedActions for the CVMVolDg resource in an SF for Oracle RAC environment.

Configuring the fire drill service group

On the secondary site, the initial steps create a fire drill service group that closely follows the configuration of the original application service group. The fire drill service group uses a point-in-time copy of the production data. Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to fail over and come online at the secondary site, should the need arise.

See [“Sample configuration for a fire drill service group”](#) on page 60.

You can create the fire drill service group using one of the following methods:

- Cluster Manager (Java Console)
See [“Creating the fire drill service group using Cluster Manager \(Java Console\)”](#) on page 56.
- Fire Drill Setup wizard
This text-based wizard is available at `/opt/VRTSvcs/bin/fdsetup-htc`.
See [“Creating the fire drill service group using the Fire Drill Setup Wizard”](#) on page 58.

Note: If multiple disk groups are dependent on the HTC or the HTCSnap resources in the application service group, then you must use the text-based Fire Drill Setup wizard to create the fire drill service group.

Creating the fire drill service group using Cluster Manager (Java Console)

This section describes how to use Cluster Manager (Java Console) to create the fire drill service group. After creating the fire drill service group, you must set the failover attribute to false so that the fire drill service group does not fail over to another node during a test.

To create the fire drill service group

- 1 Open the Veritas Cluster Manager (Java Console).
- 2 Log on to the cluster and click **OK**.
- 3 Click the **Service Group** tab in the left pane and click the **Resources** tab in the right pane.

- 4 Right-click the cluster in the left pane and click **Add Service Group**.
- 5 In the Add Service Group dialog box, provide information about the new service group.
 - In Service Group name, enter a name for the fire drill service group
 - Select systems from the Available Systems box and click the arrows to add them to the Systems for Service Group box.
 - Click **OK**.

To disable the AutoFailOver attribute

- 1 Click the **Service Group** tab in the left pane and select the fire drill service group.
- 2 Click the **Properties** tab in the right pane.
- 3 Click the **Show all attributes** button.
- 4 Double-click the **AutoFailOver** attribute.
- 5 In the Edit Attribute dialog box, clear the **AutoFailOver** check box.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.

Adding resources to the fire drill service group

Add resources to the new fire drill service group to recreate key aspects of the application service group.

To add resources to the service group

- 1 In Cluster Explorer, click the **Service Group** tab in the left pane, click the application service group and click the **Resources** tab in the right pane.
- 2 Right-click the resource at the top of the tree, select **Copy > Self and Child Nodes**.
- 3 In the left pane, click the fire drill service group.
- 4 Right-click the right pane, and click **Paste**.
- 5 In the Name Clashes dialog box, specify a way for the resource names to be modified, for example, insert an FD_ prefix. Click **Apply**.
- 6 Click **OK**.

Configuring resources for fire drill service group

Edit the resources in the fire drill service group so they work properly with the duplicated data. The attributes must be modified to reflect the configuration at the remote site. Bringing the service group online without modifying resource attributes is likely to result in a cluster fault and interruption in service.

To configure the fire drill service group

- 1 In Cluster Explorer, click the **Service Group** tab in the left pane.
- 2 Click the fire drill service group in the left pane and click the **Resources** tab in the right pane.
- 3 Right-click the HTC resource and click **Delete**.
- 4 Add a resource of type HTCSnap and configure its attributes.
- 5 Right-click the resource to be edited and click **View > Properties View**. If a resource to be edited does not appear in the pane, click **Show All Attributes**.
- 6 Edit attributes to reflect the configuration at the remote site. For example, change the Mount resources so that they point to the volumes that are used in the fire drill service group.

Creating the fire drill service group using the Fire Drill SetUp Wizard

This section describes how to use the Fire Drill SetUp Wizard to create the fire drill service group.

See “[Fire drill configurations](#)” on page 50.

To create the fire drill service group

- 1 Start the Fire Drill SetUp Wizard.

```
/opt/VRTSvcs/bin/fdsetup-htc
```
- 2 Enter the name of the application service group for which you want to configure a fire drill service group.
- 3 Select the supported snapshot configurations:
Gold, Silver or Bronze
- 4 Choose whether to run a Bronze fire drill, if the snapshot fails with Gold or Silver configurations.

```
If snapshot fails, should bronze be used? [y,n,q](n)
```
- 5 Specify the ShadowImage instance.
- 6 Press **Return** to verify the snapshot infrastructure.

- 7 In the Snapshot Details, the wizard informs whether the device group on the target array has synchronized ShadowImage devices to take a snapshot. If the devices are synchronized, press **Return**.

If the devices are not synchronized, specify the correct ShadowImage instance.

If the ShadowImage instance is correct, make sure the data between the target array and the ShadowImage device is synchronized and rerun the wizard.

- 8 Enter **y** to create the fire drill service group.

The wizard runs various commands to create the fire drill service group.

- 9 In Linux clusters, verify that the StartVolumes attribute for each DiskGroup type resource in the fire drill group is set to 1. If not, modify the resource to set the value to 1.

- 10 Schedule fire drill for the service group by adding the following command to the crontab to be run at regular intervals.

```
/opt/VRTSvcs/bin/fdsched-htc
```

- 11 Make fire drill highly available by adding the following command to the crontab on every node in this cluster.

```
fdsched-htc
```

Verifying a successful fire drill

Run the fire drill routine periodically to verify the application service group can fail over to the remote node.

To verify a successful fire drill

- 1 Bring the fire drill service group online on a node that does not have the application running. Verify that the fire drill service group comes online.

This action validates your disaster recovery configuration. The production service group can fail over to the secondary site in the event of an actual failure (disaster) at the primary site.

- 2 If the fire drill service group does not come online, review the VCS engine log for more information.

You can also view the fire drill log, which is located at `/tmp/fd-servicegroup`.

- 3 Take the fire drill offline after its functioning has been validated.

Failing to take the fire drill offline could cause failures in your environment. For example, if the application service group fails over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

Sample configuration for a fire drill service group

The sample configuration of a fire drill service group is identical to an application service group with a hardware replication resource. However, in a fire drill service group, the HTCSnap resource replaces the HTC resource.

The following configuration creates a Gold fire drill configuration.

You can configure a resource of type HTCSnap in the main.cf file as follows.

```
HTCSnap oradg_fd {  
    TargetResName = "DG"  
    ShadowInstance = 5  
    UseSnapshot = 1  
    RequireSnapshot = 0  
    MountSnapshot = 1  
}
```

Index

A

- attribute definitions
 - Hitachi TrueCopy agent 24

C

- cluster
 - heartbeats 27
- Configuring global clusters
 - for SF Oracle RAC 27

D

- disaster recovery 33

F

- failure scenarios 33
 - global clusters 34
 - application failure 34
 - host failure 34
 - network failure 34
 - replication link failure 34
 - site failure 34
 - storage failure 34
 - replicated data clusters 38
 - application failure 38
 - host failure 38
 - network failure 38
 - replication link failure 38
 - site failure 38
 - storage failure 38

FDFile attribute 54

fire drill

- about 49
- configuration wizard 55
- HTCSnap agent 51
- running 59
- service group for 55
- supported configurations 50

G

- global clusters
 - failure scenarios 34
 - overview 27

H

- Hitachi TrueCopy agent
 - attribute definitions 24
 - type definition 23
- HTCSnap agent
 - about 51
 - attribute definitions 53
 - operations 51
 - type definition 53
- HTCSnap agent attributes
 - FDFile 54
 - MountSnapshot 54
 - RequireSnapshot 54
 - Responsibility 54
 - UseSnapshot 54

I

- installing the agent
 - AIX systems 17
 - HP-UX systems 17
 - Linux systems 17
 - Solaris systems 17

M

- MountSnapshot attribute 54

R

- replicated data clusters
 - failure scenarios 38
- RequireSnapshot attribute 54
- resource type definition
 - Hitachi TrueCopy agent 23
 - HTCSnap agent 53
- Responsibility attribute 54

S

- sample configuration 25
- split-brain
 - handling in cluster 28

T

- type definition
 - Hitachi TrueCopy agent 23
 - HTCSnap agent 53

U

- uninstalling the agent
 - AIX systems 22
 - HP-UX systems 22
 - Linux systems 22
 - Solaris systems 22
- UseSnapshot attribute 54