

Veritas™ Cluster Server Agent for Hitachi TrueCopy/HP-XP Continuous Access Configuration Guide

Windows Server 2008 (x64), Windows
Server 2008 R2 (x64)

6.0.1

Veritas Cluster Server Agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access 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: 6.0.1

Document version: 6.0.1 Rev 0

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate 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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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
350 Ellis Street
Mountain View, CA 94043
<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 support 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/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

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

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

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

www.symantec.com/business/support/contact_techsupp_static.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/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as 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 support 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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec Web site.

<https://sort.symantec.com/documents>

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

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Technical Support	4	
Chapter 1	Introducing the Veritas agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	9
	About the agent for Hitachi TrueCopy/HP-XP Continuous Access	9
	Supported hardware	10
	Typical Hitachi TrueCopy/Hewlett-Packard XP Continuous Access setup in a VCS cluster	10
	Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent functions	12
	About the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent's online function	14
Chapter 2	Configuring the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	15
	Configuration concepts for the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent	15
	Resource type definition for the Hitachi TrueCopy agent	15
	Attribute definitions for the TrueCopy/HP-XP-CA agent	16
	Sample configuration for the TrueCopy/HP-XP-CA agent	19
	Before you configure the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	20
	About cluster heartbeats	21
	About configuring system zones in replicated data clusters	22
	About preventing split-brain	23
	Configuring the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	23
	Performing a manual Volume Manager rescan	23
	Configuring the agent manually in a global cluster	24
	Configuring the agent manually in a replicated data cluster	25

Chapter 3	Testing VCS disaster recovery support with Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	27
	How VCS recovers from various disasters in an HA/DR setup with Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	28
	Failure scenarios in global clusters	28
	Failure scenarios in replicated data clusters	33
	Replication link / Application failure scenarios	37
	Testing the global service group migration	38
	Testing disaster recovery after host failure	39
	Testing disaster recovery after site failure	40
	Performing failback after a node failure or an application failure	42
	Performing failback after a site failure	42
Chapter 4	Setting up fire drill	45
	About fire drills	45
	About the HTCSnap agent	46
	HTCSnap agent functions	46
	Resource type definition for the HTCSnap agent	47
	Attribute definitions for the HTCSnap agent	48
	About the Snapshot attributes	50
	Sample configuration for a fire drill service group	50
	Additional considerations for running a fire drill	50
	Before you configure the fire drill service group	51
	Configuring the fire drill service group	52
	About the Fire Drill wizard	52
	Verifying a successful fire drill	52
	Index	53

Introducing the Veritas agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access

This chapter includes the following topics:

- [About the agent for Hitachi TrueCopy/HP-XP Continuous Access](#)
- [Supported hardware](#)
- [Typical Hitachi TrueCopy/Hewlett-Packard XP Continuous Access setup in a VCS cluster](#)
- [Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent functions](#)

About the agent for Hitachi TrueCopy/HP-XP Continuous Access

The Veritas agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access provides support for application failover and recovery. The agent provides this support in environments that use TrueCopy/HP-XP-CA to replicate data between Hitachi TrueCopy/HP-XP arrays.

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

The VCS agent for TrueCopy/HP-XP-CA is enhanced to include additional monitoring and reporting capabilities. You can configure the agent to generate notifications based on the state of the configured HTC devices.

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

The agent supports TrueCopy/HP-XP-CA 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
Hitachi Lightning	data, never, and async
Hitachi Thunder	data and never

The Hitachi TrueCopy/HP-XP Continuous Access agent also supports Hitachi Universal Replicator for asynchronous replication on two sites.

Note: The terms Hitachi TrueCopy, TrueCopy/HP-XP-CA, and Hitachi TrueCopy/HP XP Continuous Access are all used interchangeably.

Supported hardware

The agent for Hitachi TrueCopy/HP-XP Continuous Access supports all versions of Command Control Interface (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.


The agent for Hitachi TrueCopy supports HP XP arrays with TrueCopy rebranded as Continuous Access.

The agent does not support other Hewlett-Packard replication solutions under the Continuous Access umbrella such as Continuous Access Storage Appliance (CASA).

Typical Hitachi TrueCopy/Hewlett-Packard XP Continuous Access setup in a VCS cluster

Figure 1-1 displays a typical cluster setup in a TrueCopy/HP-XP-CA environment.

Figure 1-1 Typical clustering setup for the agent



Clustering in a TrueCopy/HP-XP-CA 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/HP-XP array that contains the TrueCopy/HP-XP-CA 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/HP-XP array that contains the TrueCopy/HP-XP-CA 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/HP-XP array.

Hitachi TrueCopy/Hewlett-Packard XP Continuous Access 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

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/SSUS/SMPL, 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.

See [“About the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent's online function”](#) on page 14.

offline

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.

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<-all -local> -force</pre>
clean	<p>Determines whether if 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>The agent supports the following actions using the <code>hares -action</code> command from the command line:</p> <ul style="list-style-type: none"> ■ <code>pairdisplay</code>—Displays information about all devices. ■ <code>pairresync</code>—Resynchronizes the S-VOL devices from the VCS command line after connectivity failures are detected and corrected. ■ <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.

About the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access 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. If `horctakeover` exits with an error (exit code > 5), for example due to a timeout, then the agent flushes and freezes the group to indicate that user-intervention is required to identify the cause of the error.

For S-VOL devices in any state other than SSWS, 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.
- When the `OnlineTimeout` period of the entry point expires, the command `horctakeover` will not be executed, in which case the resource faults. The agent validates that the value of `OnlineTimeout` for the HTC type is ≥ 10 seconds. If this value is set to less than 10 seconds, the agent logs an appropriate error message.

Configuring the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access

This chapter includes the following topics:

- [Configuration concepts for the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent](#)
- [Before you configure the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access](#)
- [Configuring the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access](#)

Configuration concepts for the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent

Review the resource type definition and attribute definitions 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}  
    static int Numthreads = 1  
    static keylist SupportedActions = {localtakeover, pairresync,  
    pairresync-swaps, pairedisplay}  
    NameRule = resource.Groupname  
    str BaseDir = "C:\\\\HORCM\\etc"  
    str GroupName  
    int Instance  
    int SplitTakeover = 1  
    int LinkMonitor = 0  
)
```

Attribute definitions for the TrueCopy/HP-XP-CA agent

Following are the attributes associated with the agent:

BaseDir	Path to the RAID Manager Command Line interface. Type-dimension: string-scalar Default: C:\\\\HORCM\\etc.
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 (i.e. when S-VOL devices are in the SSUS state) or the replication link is manually suspended (i.e. when P-VOL devices are in the PSUS state). See “About the SplitTakeover attribute for the Hitachi TrueCopy agent” on page 18. Type-dimension: integer-scalar Default: 0

User	<p>The domain user account under which HORCM Manager is started, if it is not running.</p> <p>Type-dimension: string-scalar</p>
Domain	<p>The domain for the account specified in the User field</p> <p>This user must have sufficient privileges to perform the HORCM commands.</p> <p>Type-dimension: string-scalar</p>
Password	<p>The password for the user account specified in the User field. This password must be encrypted using the encryption tool provided by VCS i.e. <code>vcscrypt -agent .</code></p> <p>Type-dimension: string-scalar</p>
LinkMonitor	<p>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.</p> <p>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.</p> <p>If the value of the LinkMonitor attribute is set to 2, the agent generates SNMP traps or alerts when the status of the attached P-VOL and S-VOL changes. If the status of the configured HTC device changes to PSUE, PSUS, SSUS, or SSWS, the agent generates an SNMP trap indicating that the resource health has gone down. For all other types of status changes of the configured HTC devices, the agent generates an SNMP trap indicating that the resource health has improved. An error or information type message is logged by the agent in the VCS High Availability engine log- "<i>The state of P-VOL/S-VOL devices in device group device group name has changed from previous state to current state.</i>"</p> <p>Setting LinkMonitor does not affect the SplitTakeover behavior. However, you can minimize the time during which the P-VOL is in the PSUE state by setting the LinkMonitor attribute.</p> <p>Type-dimension: integer-scalar</p> <p>Default: 0</p>

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) or if the replication is manually suspended, if SVOL devices are in SSUS state.

SplitTakeover attribute = 0

The default value of the SplitTakeover 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, or if if SVOL devices are in SSUS 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 a global cluster environment, if the agent at the P-VOL side detects the PSUE or PSUS 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.

SplitTakeover attribute = 1

If the value of SplitTakeover is 1, the agent tries to make the SVOL devices writable, irrespective of the state of PVOL devices. Hence, even if there is a replication link failure, or the primary array fails, the agent attempts to failover to the S-VOL devices.

About the HTC configuration parameters

The TrueCopy/HP-XP-CA agent uses RAID manager to interact with Hitachi devices. All information about the remote site is exchanged mainly over the network.

To obtain information on the remote cluster of the pair, mention the details of the remote site in the instance configuration file.

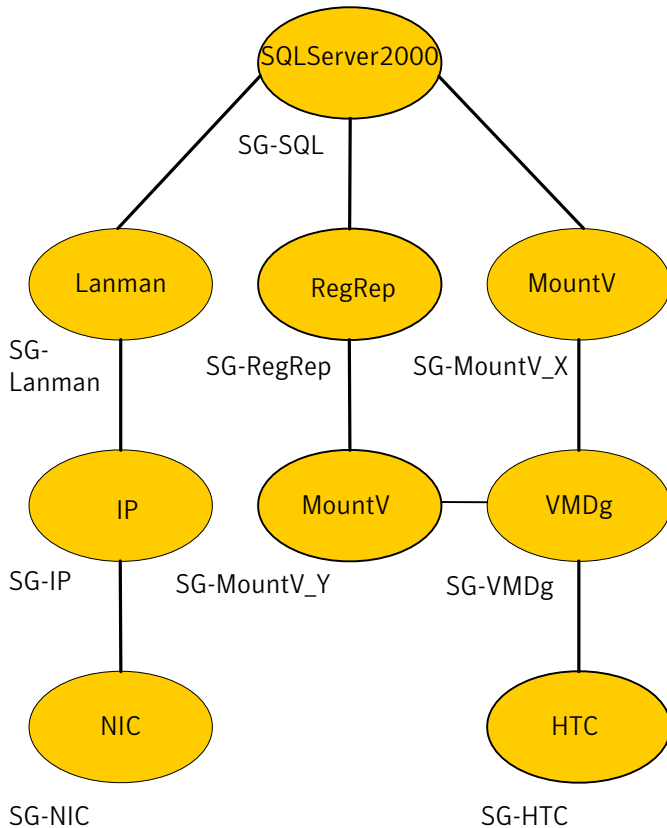
Update the HORCM_INST section of the configuration file.

Specify the value of the ClusterAddress attribute of the remote cluster in the lp_address field against the device group. Symantec recommends that you keep the ClusterService service group online on the same node, where the application service group is online.

Sample configuration for the TrueCopy/HP-XP-CA agent

[Figure 2-1](#) shows a dependency graph of a VCS service group that has a resource of type HTC.

Figure 2-1 VCS service group with resource type HTC



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

```
HTC SQLDG (  
  GroupName = SQLDG  
  Instance = 1  
  BaseDir = C:\\HORCM\\etc.  
)
```

Before you configure the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access

Before you configure the agent, review the following information:

- Verify that you have installed the agent on all systems in the cluster.

- Verify the hardware setup for the agent.
See “[Typical Hitachi TrueCopy/Hewlett-Packard XP Continuous Access setup in a VCS cluster](#)” on page 10.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See “[About cluster heartbeats](#)” on page 21.
See “[About preventing split-brain](#)” on page 23.
- Set up system zones in replicated data clusters.
See “[About configuring system zones in replicated data clusters](#)” on page 22.
- 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 Administrator's Guide*.
 - 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.
- Ensure that the HORC manager is configured to access the device groups.
 - Verify that the HTC instance is configured appropriately and is in a running state.
 - Verify that the HORC manager CLIs execute successfully. This is essential for the HTC and the HTCSnap agents to be able to fetch HTC-related data and to successfully perform failover, switchover and other operations.

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 TrueCopy/HP-XP 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/HP-XP-CA agent supports those actions that can automate the resynchronization of devices after a replication link outage is corrected.

About configuring system zones in replicated data clusters


In a replicated data cluster, you can prevent unnecessary TrueCopy/HP-XP-CA 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 2-2 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 2-2 Example system zone configuration



Modify the `SystemZones` attribute using the following command:

```
hagrpl -modify grpname SystemZones hosta 0 hostb 0 hostc 1 hostd 1
```

The variable `grpname` represents the service group in the cluster.

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 Hitachi TrueCopy/Hewlett-Packard XP Continuous Access

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

- Converting their devices to TrueCopy/HP-XP-CA devices
- Synchronizing the devices
- Adding the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent to the service group

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

See [“Sample configuration for the TrueCopy/HP-XP-CA agent”](#) on page 19.

Note: You must not change the replication state of devices primary to secondary and viceversa, outside of a VCS setup. The agent for Hitachi TrueCopy/HP-XP Continuous Access fails to detect a change in the replication state if the role reversal is done externally.

Performing a manual Volume Manager rescan

If you configure Volume Manager diskgroups on the disks that are replicated, the diskgroups do not come online the first time after failover on the secondary node.

You must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the diskgroups online. This rescans all Volume Manager objects and must be performed only once after which the failover works uninterrupted.

To perform a manual Volume Manager rescan

- 1 Bring all the resources in the service group offline on the primary node.
- 2 Bring the TrueCopy resource online on all the secondary nodes.
- 3 Run VM rescan on all the secondary nodes.
- 4 Bring all the resources (e.g. DiskGroup, Mount, and Application) online on the secondary nodes.
- 5 Fail over the service group to the primary node.

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:

```
systemdrive\Program Files\Veritas\cluster server\conf\  
Sample_HTC\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 service group, configure the service group using the Global Group Configuration Wizard.
See the *Veritas Cluster Server Administrator's Guide* for more information.
- 8 Change the ClusterFailOverPolicy attribute 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.

Note: The Disaster Recovery wizard configures the required settings for the HTC resource in the VCS application service group. Optional settings are left in the default state. The wizard creates a complete disaster recovery setup using the TrueCopy/HP-XP-CA replication and validates the replication setup. For information on using the Disaster Recovery wizard, see the Solutions guides chapters on disaster recovery.

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:

```
systemdrive\Program Files\Veritas\Cluster Server\conf\
config\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 bottom 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/Hewlett-Packard XP Continuous Access

This chapter includes the following topics:

- [How VCS recovers from various disasters in an HA/DR setup with Hitachi TrueCopy/Hewlett-Packard XP Continuous Access](#)
- [Testing the global service group migration](#)
- [Testing disaster recovery after host failure](#)
- [Testing disaster recovery after site failure](#)
- [Performing failback after a node 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/Hewlett-Packard XP Continuous Access

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 fault occurs, VCS failover behavior depends on the value of the ClusterFailOverPolicy attribute for the faulted global service group. The VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access ensures safe and exclusive access to the configured Hitachi TrueCopy/Hewlett-Packard XP Continuous Access devices.

See [“Failure scenarios in global clusters”](#) on page 28.

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/Hewlett-Packard XP Continuous Access ensures safe and exclusive access to the configured Hitachi TrueCopy/Hewlett-Packard XP Continuous Access devices.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

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

Failure scenarios in global clusters

[Table 3-1](#) lists the failure scenarios in a global cluster configuration and describes the behavior of VCS and the agent in response to the failure.

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

Table 3-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access

Failure	Description and VCS response
Application failure	<p>Application cannot start successfully on any hosts at the primary site.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Causes global service group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the ClusterFailOverPolicy global 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 node failure or an application failure” on page 42.</p> <p>See “Replication link / Application failure scenarios” on page 37.</p>
Host failure	<p>All hosts at the primary site fail.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the primary cluster fault. ■ Does the following based on the ClusterFailOverPolicy global 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 node failure or an application failure” on page 42.</p>

Table 3-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access
(continued)

Failure	Description and VCS response
Site failure	<p>All hosts and the storage at the primary site fail.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the cluster fault. ■ Does the following based on the ClusterFailOverPolicy global 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 on the secondary site in case of a manual failover based on the value of the SplitTakeover attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 1—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—Since the agent cannot reach the Raid Manager on the PVOL side, the agent issues the <code>horctakeover</code> command to make the HTC devices write-enabled. The HTC devices go into the SSWS state. If the original primary is restored, you must execute the <code>pairresync-swaps</code> action on the secondary site to restore reverse replication. <p>See “Performing failback after a site failure” on page 42.</p>

Table 3-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access (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> <ul style="list-style-type: none"> ■ Before you resync the S-VOL device, you must split off the Shadow Image device from the S-VOL device at the secondary site. ■ You must initiate resync of S-VOL device using the agent's <code>pairresync</code> action. ■ 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/Hewlett-Packard XP Continuous Access 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> <p>See “Replication link / Application failure scenarios” on page 37.</p>

Table 3-1 Failure scenarios in a global cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access
(continued)

Failure	Description and VCS response
Network failure	<p>The network connectivity and the replication link between the sites fail.</p> <p>VCS response at the secondary site:</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 global 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> <ul style="list-style-type: none"> ■ Take the global service group offline at both the sites. ■ Manually resync the data. Use the <code>pairresync swap</code> command to resynchronize from the secondary. ■ Bring the global service group online on the secondary site. <p>Agent response: Similar to the site failure.</p>
Storage failure	<p>The array at the primary site fails.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Causes the global service group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ Auto or Connected—VCS automatically brings the faulted global service 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—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 3-2](#) 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 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access

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 global service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted global service group online at the secondary site. ■ 2—You must bring the global 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 node failure or an application failure” on page 42.</p> <p>See “Replication link / Application failure scenarios” on page 37.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access
(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"> ■ 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 node failure or an application failure” on page 42.</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"> ■ 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— 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—Since the agent cannot reach the Raid Manager on the PVOL side, the agent issues the <code>horctakeover</code> command to make the HTC devices write-enabled. The HTC devices go into SSWS state. If the original primary is restored, you must execute the <code>pairresync-swaps</code> action on the secondary site to restore reverse replication. <p>See “Performing failback after a site failure” on page 42.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access (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 <code>pairresync</code> 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/Hewlett-Packard XP Continuous Access 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> <p>See “Replication link / Application failure scenarios” on page 37.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access
(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. <p>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.</p> <p>After taking the service group offline, you must manually resync the data.</p> <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> 3 Bring the service group online on one of the sites. <p>Agent response: Similar to the site failure.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access (continued)

Failure	Description and VCS response
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—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.

Replication link / Application failure scenarios

Table 3-3 shows the link failure scenarios and recommended actions:

Table 3-3 Replication link / Application failure scenarios

Event	Fence level	Recommended action
Link fails and is restored, but application does not fail over.	never, async	Run the <code>pairresync</code> action to resynchronize the S-Vols.
Link fails and application fails to the S-VOL side.	never, async, or data	Run the <code>pairresync-swaps</code> action to promote the S-VOLs to P-VOLs, and resynchronize the original P-VOLs.
Action faults due to I/O errors.	data	Run the <code>localtakeover</code> action to write enable the local devices. Clear faults and restart service group.

Testing the global service group migration

After you configure the VCS agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access, verify that the global service group can migrate to hosts across the sites. Depending on your DR configuration, perform one of the following procedures.

To test the global service group migration in global cluster setup

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

Perform the following steps:

- Switch the global 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 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 PAIR.

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

Perform the following steps:

- Switch the global service group from the secondary site to the primary site.

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

VCS brings the global service group online at the primary site.

- Verify that the HTC devices at the secondary 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 PAIR.
- 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 secondary 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/Hewlett-Packard XP Continuous Access behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 28.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

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 service group online at the secondary site.
 - Manual or Connected–You must bring the global service group online at the secondary site.
On a node in the secondary site, run the following command:

```
hagrpsvc -online -force global_group -any
```

- 2 Verify that the global service group is online at the secondary site.

```
hagrpsvc -state global_group
```

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

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:

```
hagrpsvc -online service_group -to sys_name
```

- 2 Verify that the service group is online at the secondary site.

```
hagrpsvc -state global_group
```

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

Testing disaster recovery after site failure

Review the details on site failure and how VCS and the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 28.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

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 nodes 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 SSWS.
- 3 Verify that the global service 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 global service group determines the VCS failover behavior.

- 1—VCS brings the faulted global service group online at the secondary site.
- 2—You must bring the global 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 global service group is online at the secondary site.

```
hagrp -state global_group
```

Performing failback after a node failure or an application failure

Review the details on node failure and application failure and how VCS and the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access behave in response to these failures.

See [“Failure scenarios in global clusters”](#) on page 28.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

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

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

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

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

VCS brings the global service group online 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 global service group from the secondary site to any node in the primary site.

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

VCS brings the global service group online on a node 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 global service group online at the secondary site and the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent write enables the S-VOL devices.

The device state is SSWS.

Review the details on site failure and how VCS and the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 28.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

When the hosts and the storage at the primary site are restarted and the replication link is restored, you can perform a failback of the global service group to the primary site.

To perform failback after a site failure in global cluster

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

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

- 2** Since the application has made writes on the secondary due to a failover, resync the primary from the secondary site and reverse the PVOL/SVOL roles with pairresync-swaps action on the secondary site .

After the resync is complete, the devices in the secondary are PVOL and the devices in the primary are SVOL .The device state is PAIR at both the sites.

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

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

This again swaps the role of PVOL and SVOL.

To perform failback after a site failure in replicated data cluster

- 1 Take the global 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 Since the application has made writes on the secondary due to a failover, resync the primary from the secondary site and reverse the PVOL/SVOL roles with pairresync-swaps action on the secondary site .

After the resync is complete, the devices in the secondary are PVOL and the devices in the primary are SVOL .The device state is PAIR at both the sites.

- 3 Bring the global 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
```

This again swaps the roles of PVOL and SVOL.

Setting up fire drill

This chapter includes the following topics:

- [About fire drills](#)
- [About the HTCSnap agent](#)
- [Additional considerations for running a fire drill](#)
- [Before you configure the fire drill service group](#)
- [Configuring the fire drill service group](#)
- [Verifying a successful fire drill](#)

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/Hewlett-Packard XP Continuous Access, 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.

About the HTCSnap agent

The HTCSnap agent is the fire drill agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access. 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

Gold Configuration

- 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.

Silver Configuration

- 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.

Bronze Configuration

- 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  
    str DiskGroupSnapList  
    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.</p> <p>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>
MountSnapshot	<p>Specifies whether the resource uses the snapshot to bring the service group online. Set this attribute to 1.</p> <p>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if UseSnapshot is set to 1.</p>
UseSnapshot	<p>Specifies whether the HTCSnap resource takes a local snapshot of the target array. Set this attribute to 1.</p> <p>Type-Dimension: integer-scalar</p> <p>See “About the Snapshot attributes” on page 50.</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>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if UseSnapshot is set to 1.</p>
DiskGroupSnapList	<p>This is an optional attribute that lists the original disk group names and the fire drill disk group names.</p> <p>Type-dimension: string-scaler</p>
CopyMode	<p>Indicates the array snapshot technology to be used.</p> <p>0 indicates TimeFinder/Mirror (BCV), 1 indicates TimeFinder/Clone.</p> <p>Type-dimension: integer-scalar</p> <p>Default value: 0</p>
UseTgt	<p>Indicates whether the agent should use target devices or BCVs in the device group.</p> <p>0 indicates BCV devices, 1 indicates target devices.</p> <p>Type-dimension: integer-scalar</p> <p>Default value: 0</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 4-1 lists the snapshot attribute values for fire drill configurations:

Table 4-1 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.

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.

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

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

Additional considerations for running a fire drill

Follow these guidelines for fire drills in a Windows environment:

- The primary and secondary sites must be fully configured with TrueCopy/HP-XP-CA replication and the global cluster option. The configuration must follow the applicable instructions in the Veritas Storage Foundation HA for Windows documentation for configuring disaster recovery with TrueCopy/HP-XP-CA.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.

- Ensure that Hitachi RAID Manager / Command Control Interface (CCI) is installed.
- You must configure ShadowImage pairs (for Hitachi) before running the Fire Drill wizard.
- During a fire drill, Volume Manager commands like `vxassist rescan` can cause performance issues due to the number of Volume Manager objects like diskgroups, disks, and volumes. To avoid this, you must increase the timeout value of the agent for larger configurations.

For more information, refer to *Veritas Cluster Server Release Notes*.

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.
- 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.

Configuring the fire drill service group

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

About the Fire Drill wizard

Veritas Storage Foundation High Availability for Windows (SFW HA) provides a Fire Drill Wizard to help you set up and run a fire drill on a disaster recovery environment. You launch the Fire Drill Wizard from the Solutions Configuration Center.

The Fire Drill Wizard tests the fault readiness of a disaster recovery configuration by mimicking a failover from the primary site to the secondary site. The wizard does this without stopping the application at the primary site and disrupting user access.

The Fire Drill Wizard supports conducting a fire drill for a disaster recovery site that uses Veritas Volume Replicator (VVR) or that uses Hitachi TrueCopy/Hewlett-Packard XP Continuous Access hardware replication.

For more information on using the Fire Drill wizard, see the Solutions guides chapters on testing fault readiness and running a fire drill.

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.
- 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.

Index

A

- agent functions 12
 - action 12
 - clean 12
 - info 12
 - monitor 12
 - offline 12
 - online 12
 - open 12
- attribute definitions
 - Hitachi TrueCopy agent 16
- attributes
 - BaseDir 16
 - Domain 16
 - GroupName 16
 - Instance 16
 - LinkMonitor 16
 - Password 16
 - SplitTakeover 16
 - User 16

B

- BaseDir attribute 16

C

- cluster
 - heartbeats 21
- configuring
 - before 15
 - samples 19
- CopyMode attribute 49

D

- disaster recovery 28
- DiskGroupSnapList attribute 49
- Domain attribute 16

F

- failure scenarios 28
 - global clusters 28
 - application failure 28
 - host failure 28
 - network failure 28
 - replication link failure 28
 - site failure 28
 - storage failure 28
 - replicated data clusters 33
 - application failure 33
 - host failure 33
 - network failure 33
 - replication link failure 33
 - site failure 33
 - storage failure 33
- FDFile attribute 49
- fire drill
 - about 45
 - configuration wizard 51
 - HTCSnap agent 46
 - running 52
 - service group for 51

G

- global clusters
 - failure scenarios 28
- GroupName attribute 16

H

- Hitachi TrueCopy agent
 - attribute definitions 16
 - type definition 15
- HTCSnap agent
 - about 46
 - attribute definitions 48
 - operations 46
 - type definition 47
- HTCSnap agent attributes 48–49
 - FDFile 49

HTCSnap agent attributes *(continued)*

- MountSnapshot 48
- RequireSnapshot 49
- Responsibility 49
- UseSnapshot 48

I

Instance attribute 16

L

LinkMonitor attribute 16

M

MountSnapshot attribute 48

P

Password attribute 16

R

- replicated data clusters
 - failure scenarios 33
- RequireSnapshot attribute 49
- resource type definition
 - Hitachi TrueCopy agent 15
 - HTCSnap agent 47
- Responsibility attribute 49

S

- sample configuration 19
- ShadowInstance attribute 48
- split-brain
 - handling in cluster 23
- SplitTakeover attribute 16

T

- type definition
 - Hitachi TrueCopy agent 15
 - HTCSnap agent 47
- typical setup 10

U

- User attribute 16
- UseSnapshot attribute 48
- UseTGibute 49