

Cluster Server 7.4 Agent for Hitachi TrueCopy/HP-XP Continuous Access Configuration Guide - Windows

Windows

Last updated: 2019-02-01

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas 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. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

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

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introducing the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	6
	About the agent for Hitachi TrueCopy/HP-XP Continuous Access	6
	Supported software	7
	Supported hardware	7
	Typical Hitachi TrueCopy/Hewlett-Packard XP Continuous Access setup in a VCS cluster	8
	Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent functions	9
	About the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent's online function	12
Chapter 2	Configuring the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	14
	Configuration concepts for the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent	14
	Resource type definition for the Hitachi TrueCopy agent	14
	Attribute definitions for the TrueCopy/HP-XP-CA agent	15
	Sample configuration for the TrueCopy/HP-XP-CA agent	25
	Before you configure the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	28
	About cluster heartbeats	29
	About configuring system zones in replicated data clusters	29
	About preventing split-brain	30
	Configuring the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	31
	Configuring the agent manually in a global cluster	31
	Configuring the agent manually in a replicated data cluster	32

Chapter 3	Testing VCS disaster recovery support with Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	33
	How VCS recovers from various disasters in an HA/DR setup with Hitachi TrueCopy/Hewlett-Packard XP Continuous Access	34
	Failure scenarios in global clusters	34
	Failure scenarios in replicated data clusters	39
	Replication link / Application failure scenarios	43
	Testing the global service group migration	44
	Testing disaster recovery after host failure	45
	Testing disaster recovery after site failure	46
	Performing failback after a node failure or an application failure	47
	Performing failback after a site failure	48
Chapter 4	Setting up fire drill	50
	About fire drills	50
	About the HTCSnap agent	51
	HTCSnap agent functions	51
	Resource type definition for the HTCSnap agent	52
	Attribute definitions for the HTCSnap agent	52
	About the Snapshot attributes	54
	Sample configuration for a fire drill service group	54
	Additional considerations for running a fire drill	55
	Before you configure the fire drill service group	55
	Configuring the fire drill service group	56
	About the Fire Drill wizard	56
	Verifying a successful fire drill	56

Introducing the 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 software](#)
- [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 Cluster Server 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.

Supported software

For information on the software versions that the agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access supports, see the Veritas Services and Operations Readiness Tools (SORT) site:

<https://sort.veritas.com/agents>.

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

The following figure 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:

Table 1-2 Agent Functions

Function	Description
online	<p>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.</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>For S-VOL devices in any state other than SSWS/SSUS/SMPL, the agent runs the <code>horctakeover</code> command and makes the devices writable. The time required for failover depends on the following conditions:</p> <ul style="list-style-type: none"> ■ The health of the original primary. ■ The RAID Manager timeouts as defined in the <code>horcm</code> configuration file for the device group. <p>The agent considers P-VOL devices writable and takes no action other than going online, regardless of their status.</p> <p>If the S-VOL devices are in the COPY state, the agent runs the <code>horctakeover</code> command after one of the following:</p> <ul style="list-style-type: none"> ■ The synchronization from the primary completes. ■ The <code>OnlineTimeout</code> period of the entry point expires, in which case the resource faults. <p>If the S-VOL devices are in the PAIR state, the agent runs the <code>pairdisplay</code> command without the <code>-l</code> option to retrieve the state of the remote site devices. If it finds that the P-VOL devices are in PAIR state, the agent proceeds with the failover. But if the remote RAID manager is down, then the agent honors the <code>SplitTakeover</code> attribute configuration before performing failover.</p> <p>See “About the Hitachi TrueCopy/Hewlett-Packard XP Continuous Access agent’s online function” on page 12.</p> <p>If the RAID manager is not up for an instance, the agent runs the <code>horcmstart</code> command to bring the RAID manager online.</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>

Table 1-2 Agent Functions (*continued*)

Function	Description
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>Based on other attribute values, the monitor entry point examines 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 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 it 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>pairedisplay</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.

Table 1-2 Agent Functions (*continued*)

Function	Description
action\PreSwitch	<p>Ensures that the remote site cluster can come online during a planned failover within a GCO configuration without data loss. The VCS engine on the remote cluster invokes the PreSwitch action on all the resources of the remote Service Group during a planned failover using the <code>hagrps -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>Cluster Server Administrator's Guide</i>.</p>
action\vxdiske	<p>Reports the mapping between the physical disk name and the volume manager disk name for all connected disks.</p>

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

If the state of all local devices is SMPL (Simplex), then the AllowOnlineOnSimplex attribute is honored to allow/disallow resource to come online.

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 the `horctakeover` command 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 and SSUS, the agent honors the SplitTakeover attribute and runs the `horctakeover` command to make 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.

If the S-VOL devices are in the SSUS state and if the RoleMonitor attribute is set to 1, the agent runs the `pairdisplay` command without the `-l` option, to determine if the S-VOL is in a writable state. The agent behavior when devices are in S-VOL SSUS state is as follows:

- If S-VOL devices are in SSUS writable state, the agent proceeds with online without failover.
- If S-VOL devices are in SSUS read only state, the agent honors the SplitTakeover attribute and accordingly proceeds with failover to make the devices writable.
- In case agent could not connect to remote RAID manager, the agent faults the resource.

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 `horctakeover` command will not be executed, in which case the resource faults.

If S-VOL devices are in the PAIR state, the agent issues the `pairdisplay` command without the `-l` option to get the replication link state. If it finds that the P-VOL devices are in the PAIR state, the agent proceeds with failover. If remote horcm is down, the SplitTakeover attribute is honored before issuing the `horctakeover` command. The agent validates the value of OnlineTimeout for the HTC type is sufficient to run the `horctakeover` command. If the agent finds this value of OnlineTimeout is insufficient, 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 keylist RegList = { SplitTakeover, LinkMonitor, RoleMonitor,
    FreezeSecondaryOnSplit, AllowOnlineOnSimplex }
    static i18nstr ArgList[] = { BaseDir, GroupName, Instance,
    SplitTakeover, LinkMonitor, RoleMonitor, FreezeSecondaryOnSplit,
    AllowOnlineOnSimplex, User, Domain, Password }
    static keylist SupportedActions = { localtakeover, pairresync,
    pairresync-swaps, pairedisplay, PreSwitch, vxdiske }
    str BaseDir = "C:\\\\HORCM\\etc"
    str GroupName
    int Instance
    int SplitTakeover
    int LinkMonitor
    int RoleMonitor
    int FreezeSecondaryOnSplit
    boolean AllowOnlineOnSimplex = 0
    str User
    str Domain
    str Password
    temp str TargetFrozen
    temp str OldState = ""
    str AdvancedOpts{} = { AllowAutoFailoverInterval="undefined" }
    temp str PVOLStateTime
)

```

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

[Table 2-1](#) lists the attributes associated with the agent:

Table 2-1 Attributes for the Hitachi TrueCopy agent

Attribute	Description
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

Table 2-1 Attributes for the Hitachi TrueCopy agent (*continued*)

Attribute	Description
Instance	<p>The Instance number of the device that the agent manages. Multiple device groups may have the same instance number.</p> <p>Do not define the attribute if the instance number is zero.</p> <p>Type-Dimension: integer</p>
SplitTakeover	<p>A flag that determines the following:</p> <ul style="list-style-type: none"> ■ Whether the agent permits a failover to S-VOL devices if the replication link is disconnected (that is, when P-VOL devices are in the PSUE state) ■ Whether the agent cannot connect to the remote site RAID manager ■ Whether the replication link is manually suspended (that is when P-VOL devices are in the PSUS state) <p>See “About the SplitTakeover attribute for the Hitachi TrueCopy agent” on page 19.</p> <p>Type-Dimension: integer-scalar</p> <p>Default: 0</p>
User	<p>The domain user account under which HORCM Manager is started.</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. You must encrypt this password by using the <code>vcscrypt</code> utility.</p> <p>Type-Dimension: string-scalar</p>

Table 2-1 Attributes for the Hitachi TrueCopy agent (*continued*)

Attribute	Description
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- "The state of <i>P-VOL/S-VOL</i> devices in device group <i>device group name</i> has changed from <i>previous state</i> to <i>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 Default: 0</p>
RoleMonitor	<p>Determines if the agent must perform detailed monitoring of HTC volumes.</p> <p>If this attribute is set to 0, the agent does not perform detailed monitoring. This attribute is disabled by default.</p> <p>If this attribute is set to 1, the agent monitors the status of the HTC volumes everytime a monitor cycle runs. In addition, the HTC resource comes online only when any of the following conditions are met:</p> <ul style="list-style-type: none"> ■ When the volume is P-VOL ■ When the volume is S-VOL and the status is SSWS ■ When the volume is S-VOL, the status is SSUS, and the M flag of the corresponding P-VOL is set to W. <p>Type-Dimension: integer-scalar Default: 0</p>

Table 2-1 Attributes for the Hitachi TrueCopy agent (*continued*)

Attribute	Description
FreezeSecondaryOnSplit	<p>A flag that determines if the agent must freeze the service group in the remote cluster when the TrueCopy replication link is either split or suspended.</p> <p>The value 1 indicates that the agent must freeze the service group in the remote cluster when the replication link is split or suspended.</p> <p>Type-Dimension: integer-scalar</p> <p>Default: 0</p>
AllowOnlineOnSimplex	<p>A flag that determines if the agent must allow a resource to come online when the TrueCopy devices are in SMPL (Simplex) state. This attribute is honored only when the agent attempts to bring a resource online.</p> <p>The value false indicates that the agent must not allow a resource to come online when TrueCopy devices are in SMPL (Simplex) state.</p> <p>Type-Dimension: boolean-scalar</p> <p>Default: false</p>
TargetFrozen	<p>For internal use. Do not modify.</p>
AdvancedOpts	<p>Used at the time of monitoring. This attribute enables the agent to execute a custom script during the monitor cycle of the resource.</p> <p>Use the AllowAutoFailoverInterval attribute with this attribute. The agent automatically fails over if certain conditions are met and AllowAutoFailoverInterval is set to 0 (zero) or a positive integer.</p> <p>To disable the execution of the custom script, set AllowAutoFailoverInterval to <code>undefined</code> or remove it from the AdvancedOpts attribute. For example:</p> <pre>AdvancedOpts{ } = { AllowAutoFailoverInterval="undefined" }</pre> <p>Type-Dimension: string-association</p>

Table 2-1 Attributes for the Hitachi TrueCopy agent (*continued*)

Attribute	Description
AllowAutoFailoverInterval	<p>The agent uses this attribute to fail over automatically only if all the following conditions are met:</p> <ul style="list-style-type: none"> ■ This attribute is set to 0 (zero) or a positive integer. ■ The fence level is NEVER. ■ The remote RAID manager is not reachable. <p>The failover takes place only if the value of this attribute is greater than the last registered PAIR state time difference.</p> <p>Note: This attribute is applicable only when the fence level is NEVER.</p> <p>See “Special consideration for fence level NEVER” on page 21.</p> <p>Default: <code>undefined</code></p> <p>Type-Dimension: string-association</p>
PVOLStateTime	<p>This is an internal attribute that is used to maintain the P-VOL state and the timestamp when the instance was last registered as PAIR.</p> <p>Note: Do not modify this attribute.</p>

About the SplitTakeover attribute for the Hitachi TrueCopy agent

The SplitTakeover attribute determines the following:

- Whether the agent permits a failover to S-VOL devices if the replication link is disconnected (that is, when P-VOL devices are in the PSUE state).
- Whether the agent cannot connect to the remote site RAID manager.
- Whether the replication link is manually suspended (that is when P-VOL devices are in the PSUS 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 the agent cannot connect to the remote site RAID manager, or if the S-VOL devices are in the 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.

If the S-VOL devices are in the PAIR state, the agent attempts to contact the RAID manager at the P-VOL side to determine the status of the arrays.

If the P-VOL devices are in the PAIR state, the agent proceeds with failover. But if the P-VOL side is down, the agent attempts to honor the SplitTakeover attribute configuration before proceeding with failover.

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.

SplitTakeover attribute = 1

If there is a replication link failure, or if the primary array fails, or if a pair is suspended, the agent allows failover to the S-VOL devices.

About the FreezeSecondaryOnSplit attribute for the Hitachi TrueCopy agent

In a global cluster environment, if the agent at the P-VOL side detects the PSUE or PSUS state locally and FreezeSecondaryOnSplit is set to 1, then the agent 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.

FreezeSecondaryOnSplit attribute = 0

If the value of the FreezeSecondaryOnSplit attribute is 0, the agent unfreezes the remote site service group if it is already frozen. Hence, even if there is a replication link failure, or if the primary array fails, or if a pair is suspended, the agent allows 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.

In a multi-node configuration, horcm instances can be configured in the following manner:

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

The agent honors the default value of the remote RAID manager communication timeout (30sec) and poll (10sec) of the horcm configuration file. If the user modifies the remote RAID manager timeout value and the agent finds it insufficient for online operation, the agent logs an appropriate error message and faults the resource.

The recommended values of the agent attributes, if the value of remote RAID manager timeout is modified, are as follows:

- The OnlineTimeout value of HTC type should be four times more than the value of remote RAID manager timeout with some additional buffer time (~10sec).
- The MonitorTimeout value of HTC type should be more than twice the value of remote RAID manager timeout with some additional buffer time (~10sec).
- The ActionTimeout value of HTC type should be more than twice the value of remote RAID manager timeout.

Special consideration for fence level NEVER

During each monitor cycle, the VCS agent for HTC records the P-VOL status with the timestamp and propagates this information to the secondary site. The secondary site uses this information to keep track of the last known PAIR time of P-VOL.

Consider the following failure scenario:

- The primary site has failed.
- The status of P-VOL cannot be determined, because the RAID manager for that site is not reachable.
- The replication status of S-VOL is displayed as PAIR.

The agent provides the AllowAutoFailoverInternal attribute that lets you configure automatic failover in this scenario. The automatic failover allows for minimum downtime at the risk of data loss or corruption.

In this scenario, the agent allows a failover to happen only if

`AllowAutoFailoverInterval < (Event B - Event A)`, where:

- Event A is the last known PAIR status of P-VOL, which is a timestamp.
- Event B is the time at which the secondary site detects that the primary site has failed and the remote RAID manager is not reachable.

The AllowAutoFailoverInterval value is passed to the AdvancedOpts attribute.

Table 2-2 Failover scenarios for the various AllowAutoFailoverInterval values

Value	Conditions	Actions
0	<ul style="list-style-type: none"> ■ The fence level is NEVER. ■ The remote RAID manager is not reachable. 	The failover is triggered. The last known PAIR status of P-VOL is recorded but not used.
Greater than 0	<ul style="list-style-type: none"> ■ The fence level is NEVER. ■ The remote RAID manager is not reachable. ■ The last known remote state within τ seconds is PAIR, where τ is the AllowAutofailoverInterval value. 	The failover is triggered. The last known PAIR status of P-VOL is recorded and is used to determine the failover action.
Greater than 0	<ul style="list-style-type: none"> ■ The fence level is NEVER. ■ The remote RAID manager is not reachable. ■ The last known remote state is PAIR and is greater than τ seconds, where τ is the AllowAutofailoverInterval value. 	The failover is not triggered. The last known PAIR status of P-VOL is recorded and is used to determine the failover action.
Greater than 0	<ul style="list-style-type: none"> ■ The fence level is NEVER. ■ The remote RAID manager is not reachable. ■ The last known remote state is not PAIR. 	The failover is not triggered. The last known PAIR status of P-VOL is recorded and is used to determine the failover action.
undefined or The value is not passed to the AdvancedOpts attribute.	Any	The automatic failover is not enabled. The last known PAIR status of P-VOL is not recorded. Manual intervention is required to restore the operations in this scenario.

Consider the following before using the AllowAutoFailoverInterval attribute:

- This attribute can allow for an automatic failover only when all the following conditions are met:
 - The fence level is NEVER.

- The remote HORCM connection has failed.
- The SplitTakeover attribute is set to 0 (zero).
- The use of this attribute provides a tradeoff between minimum downtime and data consistency. You may achieve a smaller downtime at the cost of possible data loss or corruption. The tradeoff exists because, in the fence level NEVER, if the remote HORCM is down, there is no way to figure out whether the replication link is healthy and the latest data is available for failover.
- In this scenario if takeover has failed, the service group goes into Freeze state.
- If the SplitTakeover attribute is set to 1, the agent triggers a failover regardless of the AllowAutoFailoverInterval value.

See [“Considerations for calculating the AllowAutoFailoverInterval attribute value”](#) on page 23.

Considerations for calculating the AllowAutoFailoverInterval attribute value

You can configure the VCS agent for HTC to trigger an automatic failover when a primary site failure occurs in a Global Cluster Option (GCO) environment. Such a configuration comes into effect after a certain time has elapsed, which is defined by the AllowAutoFailoverInterval attribute.

The value of AllowAutoFailoverInterval is determined based on the following events:

- The time when the latest PAIR status of P-VOL is propagated to the secondary site
- The time when the secondary site detects the primary site failure
- The time when the remote RAID manager is longer reachable

Table 2-3 Variables used to calculate the value of AllowAutoFailoverInterval

Variable	Source	Default value	Usage
A = MonitorInterval attribute value	HTC agent	60 seconds	Specifies how frequently the agent polls and records the PAIR status for P-VOL.

Table 2-3 Variables used to calculate the value of AllowAutoFailoverInterval (continued)

Variable	Source	Default value	Usage
B = AYAIInterval attribute value	Heartbeat agent	60 seconds	The interval between two heartbeats in the global cluster. You can modify this value using the <code>hahb</code> command, for example: <pre>hahb -modify ICMP AYAIInterval 45</pre>
C = AYARetryLimit attribute value	Heartbeat agent	3 attempts	The maximum number of lost heartbeats before the agent reports that heartbeat to the cluster is down. You can modify this value using the <code>hahb</code> command, for example: <pre>hahb -modify ICMP AYARetryLimit 2</pre>
D = Timeout value specified for the cluster node at the Primary site	HORCM file at the Secondary site	120 seconds	The HTC agent at secondary site attempts to get the replication link state using the <code>pairdisplay</code> command. This operation times out after the specified interval.

Considering the default values, the time interval is calculated as follows:

$$\begin{aligned}
 &A + (B \times C) + D + BufferTime \\
 &= 60 + (60 \times 3) + 120 + 40 \\
 &= 400 \text{ seconds}
 \end{aligned}$$

You can modify these attribute values (A to D) to reduce the effective failover time. For example, the turnaround time can be reduced to 180 seconds by tweaking the attributes values as follows:

- A = 30 seconds
- B = 45 seconds
- C = 2 attempts
- D = 30 seconds

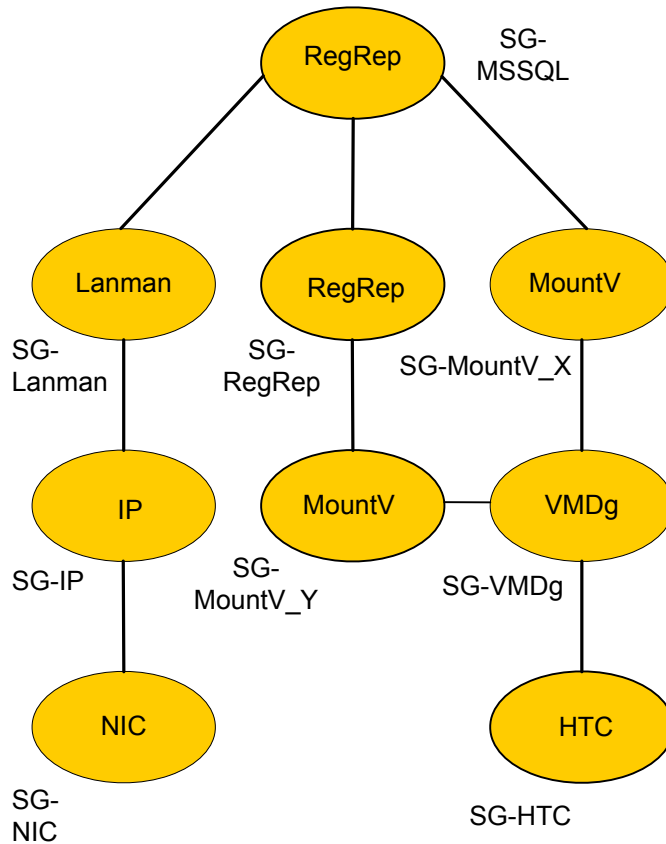
Caution: AYAIInterval and AYARetryLimit are responsible for GCO link monitoring and toleration of intermittent network failures. Significantly reducing this value may falsely flag intermittent network issues as network failures, which may trigger a failover.

Note: AYAIInterval and AYARetryLimit are not used in a replicated data cluster (RDC) environment, so the effective time for failover in that environment is greatly reduced.

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.
)
  
```

Sample main.cf configuration for CVM with the HTC resource:

```

group HTC (
  SystemList = { fred = 0, barney = 1 }
  Parallel = 2
  ClusterList = { clus1 = 0, clus2 = 1 }
  Authority = 1
)
  
```

```

AutoStartList = { fred, barney }
)

CFSMount htc_mnt (
    BlockDevice = "/dev/vx/dsk/TCdg/htcvol"
    MountPoint = "/htc"
)

CVMVolDg htc_dg (
    CVMVolume = { htcvol }
    CVMActivation = sw
    CVMDeportOnOffline = 1
    CVMDiskGroup = TCdg
    ClearClone = 1
)

HTC rep_htc (
    GroupName = vg1
    Instance = 1
)

requires group cvm online local firm
htc_dg requires rep_htc
htc_mnt requires htc_dg

group cvm (
    SystemList = { fred = 0, barney = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { fred, barney }
)

CFSfsckd vxfsckd (
)

CVMcluster cvm_clus (
    CVMTransport = gab
    CVMClustName = htc701
    CVMTimeout = 200
    CVMNodeId = { fred = 0, barney = 1 }
)

CVMVxconfigd cvm_vxconfigd (

```

```
Critical = 0  
CVMVxconfigdArgs = { syslog }  
)
```

```
cvm_clus requires cvm_vxconfigd  
vxfsckd requires cvm_clus
```

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 8.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 29.
See [“About preventing split-brain”](#) on page 30.
- Set up system zones in replicated data clusters.
See [“About configuring system zones in replicated data clusters”](#) on page 29.
- 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, refer to the *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:

```
hagrps -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 S-VOL to P-VOL. 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 means the loss of replication links.

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

You can configure clustered application in 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 25.

Note: You must not change the replication state of devices from primary to secondary and from secondary to primary, 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 and RoleMonitor is disabled.

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 (Java Console) 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.
Refer to the *Cluster Server Administrator's Guide* for more information.
- 8 Change the ClusterFailOverPolicy attribute from the default, if necessary. Veritas 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 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.
- 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 topic lists various failure scenarios and describes how VCS responds to the failures in 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 34.

Replicated data clusters

When service group faults or system faults occur, the VCS failover behavior depends on the value of the AutoFailOver attribute of 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 39.

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

Failure scenarios in global clusters

The following table lists the failure scenarios in a global cluster configuration and describes the behavior of VCS and the agent in response to the failure.

Table 3-1 Failure scenarios in a global cluster configuration with the 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, except when the link is manually suspended with the read-only option. ■ 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 47.</p> <p>See “Replication link / Application failure scenarios” on page 43.</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>The agent does the following:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site, except when the link is manually suspended with the read-only option. ■ 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 47.</p>

Table 3-1 Failure scenarios in a global cluster configuration with the 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—Agent does not perform failover to the secondary site. <p>See “Performing failback after a site failure” on page 48.</p>

Table 3-1 Failure scenarios in a global cluster configuration with the 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: When the replication link is disconnected, the agent does the following based on the value of LinkMonitor attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 0—No action. ■ 1—The agent periodically attempts to resynchronize the S-VOL side using the <code>pairresync</code> command. ■ 2—The agent periodically attempts to resynchronize the S-VOL side and also sends notifications about the disconnected link. Notifications are sent in the form of either SNMP traps or emails. For information about the VCS NotifierMngr agent, refer to the <i>Cluster Server Bundled Agents Reference Guide</i>. <p>If the value of the LinkMonitor attribute is not set to 1 or 2, 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 the S-VOL device using the agent's <code>pairresync</code> action. ■ After P-VOL and S-VOL devices are in sync, re-establish 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. Veritas recommends configuring Shadow Copy devices at both the sites.</p> <p>See "Replication link / Application failure scenarios" on page 43.</p>

Table 3-1 Failure scenarios in a global cluster configuration with the 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 is restored, you must manually resync the data.</p> <p>Note: Veritas 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 resynchronize 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

The following table 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 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>The agent does the following:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site, except when the link is manually suspended with the read-only option. ■ 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 47.</p> <p>See “Replication link / Application failure scenarios” on page 43.</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>The agent does the following:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site, except when the link is manually suspended with the read-only option. ■ 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 47.</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 — Agent does not perform failover to the secondary site. <p>See “Performing failback after a site failure” on page 48.</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: When the replication link is disconnected, the agent does the following based on the LinkMonitor attribute of the HTC resource:</p> <ul style="list-style-type: none"> ■ 0—No action. ■ 1—The agent periodically attempts to resynchronize the S-VOL side using the <code>pairresync</code> command. ■ 2—The agent periodically attempts to resynchronize the S-VOL side and also sends notifications about the disconnected link. Notifications are sent in the form of either SNMP traps or emails. For information about the VCS NotifierMngr agent, refer to the <i>Cluster Server Bundled Agents Reference Guide</i>. <p>If the value of the LinkMonitor attribute is not set to 1 or 2, 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. Veritas recommends configuring Shadow Copy devices at both the sites.</p> <p>See “Replication link / Application failure scenarios” on page 43.</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 resynchronize the data.</p> <p>Note: Veritas 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 resynchronize the data. <p>Depending on the site whose data you want to retain run the <code>pairresync</code> or the <code>pairresync-swap</code> command.</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>

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 <code>AutoFailOver</code> 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 <code>SplitTakeover</code> 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 does not perform failover to the secondary site.

Replication link / Application failure scenarios

The following table lists 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 Cluster Server 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.

```
hagrps -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 Cluster Server agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access behave in response to the failure.

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

See [“Failure scenarios in replicated data clusters”](#) on page 39.

Perform the procedure that is applicable to your DR configuration 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:

```
hagrps -online -force global_group -any
```

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

```
hagrps -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:

```
hagrps -online service_group -to sys_name
```

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

```
hagrps -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 Cluster Server agent for Hitachi TrueCopy/Hewlett-Packard XP Continuous Access behave in response to the failure.

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

See [“Failure scenarios in replicated data clusters”](#) on page 39.

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 34.

See [“Failure scenarios in replicated data clusters”](#) on page 39.

After the nodes at the primary site are restarted, you can perform a failback of the global service group to the primary site. Perform the procedure that applicable to your DR configuration.

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 34.

See [“Failure scenarios in replicated data clusters”](#) on page 39.

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, resynchronize the primary from the secondary site and reverse the P-VOL/S-VOL roles with the `pairresync-swaps` action on the secondary site.

After the resync is complete, the devices in the secondary are P-VOL and the devices in the primary are S-VOL. 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 P-VOL and S-VOL.

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 P-VOL/S-VOL roles with the `pairresync-swaps` action on the secondary site.

After the resync is complete, the devices in the secondary are P-VOL and the devices in the primary are S-VOL. 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 P-VOL and S-VOL.

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 drill for storage devices that are managed using Veritas Volume Manager.

The agent supports fire drill 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 maintains 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

- Splits 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

Gold Configuration

- 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.

Silver Configuration

- 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.

Bronze Configuration

- 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.

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,  
        DiskGroupSnapList, User, Domain, Password }  
    static int OpenTimeout = 180  
    str TargetResName  
    int ShadowInstance  
    int MountSnapshot  
    int UseSnapshot  
    int RequireSnapshot  
    str DiskGroupSnapList  
    str User  
    str Domain  
    str Password  
    temp str Responsibility  
    temp str FDFile  
)
```

Attribute definitions for the HTCSnap agent

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

ShadowInstance	The instance number of the ShadowInstance P-VOL group. The P-VOL group must include one of the following: <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). Type-dimension: integer-scalar
----------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

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 database 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 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>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-scalar</p>

Responsibility	Do not modify. For internal use only. Used by the agent to keep track of resynchronizing snapshots. Type-Dimension: temporary string
FDFile	Do not modify. For internal use only. Used by the agent to store the absolute pathname to the file with the latest fire drill report on the local system. Type-Dimension: temporary string

About the Snapshot attributes

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

The following table 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 sqldg_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 Storage Foundation and High Availability Solutions 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 the *Veritas InfoScale Release Notes*.

Before you configure the fire drill service group

Before you configure the fire drill service group, ensure that the following pre-requisites are met:

- 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 ShadowImage 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 HORCM 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

Storage Foundation and High Availability Solutions 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 Volume Replicator (VVR) or that uses Hitachi TrueCopy/Hewlett-Packard XP Continuous Access hardware replication.

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 at the secondary site that does not have the application running.

If the fire drill service group comes online, it 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.