

Veritas™ Cluster Server Agent for Hitachi TrueCopy Configuration Guide

ESX

5.1 MP2

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

Agent version: 5.1 MP2

Document Version: 5.1.MP2.0

Legal Notice

Copyright © 2008 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo 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. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

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

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

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

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

<http://www.symantec.com>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Technical Support

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

Symantec's maintenance offerings include the following:

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

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

www.symantec.com/techsupp/

Contacting Technical Support

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

www.symantec.com/techsupp/

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

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

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

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

Licensing and registration

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

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

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

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

Maintenance agreement resources

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

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

Additional enterprise services

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

Enterprise services that are available include the following:

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

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

www.symantec.com

Select your country or language from the site index.

Technical Support	4
Chapter 1	Introducing the VCS agent for Hitachi TrueCopy
	TrueCopy
	About the agent for Hitachi TrueCopy
	What's new in this release
	Supported software and hardware
	Typical Hitachi TrueCopy setup in a VCS cluster
	Hitachi TrueCopy agent functions
	About the Hitachi TrueCopy agent's online function
	Additional considerations in an ESX environment
Chapter 2	Configuring the agent for Hitachi TrueCopy
	Configuration concepts for the Hitachi TrueCopy agent
	Resource type definition for the Hitachi TrueCopy agent
	Attribute definitions for the Hitachi TrueCopy agent
	Sample configuration for the Hitachi TrueCopy agent
	Before you configure the agent for TrueCopy
	About cluster heartbeats
	About configuring system zones in replicated data clusters
	About preventing split-brain
	Configuring the agent for TrueCopy
	Configuring the agent manually in a global cluster
	Configuring the agent manually in a replicated data cluster
Chapter 3	Managing and testing clustering support for Hitachi TrueCopy
	Typical test setup for the Hitachi TrueCopy agent
	Testing service group migration
	Testing host failure
	Performing a disaster test
	Performing the failback test
	Failure scenarios for Hitachi TrueCopy
	Site disaster

	All host or all application failure	27
	Replication link failure	27
	Split-brain in a TrueCopy environment	28
	Rescanning Host Bus Adapters (HBAs) on VCS nodes	29
Chapter 4	Setting up a fire drill	31
	About fire drills	31
	About the HTCSnap agent	32
	HTCSnap agent functions	32
	About the agent's online function	33
	Processing the snapshot	33
	Resource type definition for the HTCSnap agent	34
	HTCSnap agent attributes	34
	Sample configuration for a fire drill service group	36
	Before you configure the fire drill service group	36
	Additional requirements for running a fire drill in an ESX environment	37
	Configuring the fire drill service group	38
	Configuring the new virtual machine for the fire drill	38
	Running the fire drill	39
	Verifying a successful fire drill	39
	Index	41

Introducing the VCS agent for Hitachi TrueCopy

This chapter includes the following topics:

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

About the agent for Hitachi TrueCopy

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

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

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

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

The agent supports different fence levels for different arrays:

Table 1-1 Supported fence levels

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

What's new in this release

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

- The Veritas agent for Hitachi Truecopy supports replicated data clusters in this release.

Supported software and hardware

The Hitachi TrueCopy agent supports Veritas Cluster Server 5.1 MP2 for ESX.

The agent for Hitachi TrueCopy provides support for the following:

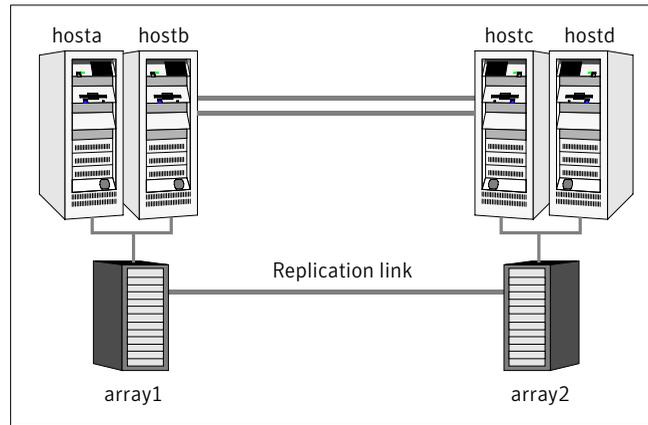
- All versions of the Hitachi RAID Manager
The agent supports TrueCopy on all microcode levels on all arrays, provided the host, HBA, array combination is in Hitachi's hardware compatibility list.
- Hewlett-Packard XP arrays with TrueCopy rebranded as Continuous Access
The agent for Hitachi TrueCopy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

Note: VMotion (i.e. the live migration of the virtual machine from the source ESX host to the target) is not supported across replicated system zones.

Typical Hitachi TrueCopy setup in a VCS cluster

[Figure 1-1](#) displays a typical cluster setup in a TrueCopy environment.

Figure 1-1 Typical clustering setup for the agent



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

- The primary array (array1) has one or more P-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to the Hitachi TrueCopy array that contains the TrueCopy P-VOL devices.
- The secondary array (array2) has one or more S-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to a Hitachi TrueCopy array that contains the TrueCopy S-VOL devices. The S-VOL devices are paired with the P-VOL devices in the P-VOL array. The S-VOL hosts and arrays must be at a significant distance to survive a disaster that may occur at the P-VOL side.
- Network heartbeating between the two data centers to determine their health; this network heartbeating could be LLT or TCP/IP.
 See “[About cluster heartbeats](#)” on page 19.
- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT.
 In a global cluster environment, you must attach all hosts in a cluster to the same Hitachi TrueCopy array.

Hitachi TrueCopy agent functions

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

The agent performs the following functions:

online	<p>If the state of all local devices is read-write enabled, the agent creates a lock file on the local host to indicate that the resource is online. This action makes the devices writable for the application.</p> <p>If one or more devices are not in a writable state, the agent runs the <code>horctakeover</code> command to enable read-write access to the devices.</p> <p>See “About the Hitachi TrueCopy agent's online function” on page 13.</p>
offline	<p>The agent removes the lock file that was created for the resource by the online entry point. The agent does not run any TrueCopy commands because taking the resource offline is not indicative of an intention to give up the devices.</p>
monitor	<p>Verifies the existence of the lock file to determine the resource status. If the lock file exists, the agent reports the status of the resource as online. If the lock file does not exist, the agent reports the status of the resource as offline.</p> <p>The monitor entry point does not examine the state of the devices or the state of the replication link between the arrays.</p>
open	<p>Removes the lock file from the host on which this entry point is called. This functionality prevents potential concurrency violation if the group fails over to another node.</p> <p>Note that the agent does not remove the lock file if the agent starts after the following command:</p> <pre>hastop -force</pre>
clean	<p>Determines whether 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>Resynchronizes the devices from the VCS command line after connectivity failures are detected and corrected.</p> <p>The agent supports the following actions:</p> <ul style="list-style-type: none">■ <code>pairdisplay</code>—Displays information about all devices.■ <code>pairresync</code>—Resynchronizes the S-VOLs.■ <code>pairresync-swaps</code>—Promotes the S-VOLs to P-VOLs and resynchronizes the original P-VOLs.■ <code>localtakeover</code>—Makes the local devices write-enabled.

About the Hitachi TrueCopy agent's online function

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

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

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

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

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

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

- The synchronization from the primary completes.
- The `OnlineTimeout` period of the entry point expires, in which case the resource faults.

Additional considerations in an ESX environment

The agent performs the following actions before coming online:

- Changes the following ESX server settings:
 - LVM.DisallowSnapshotLun=0.
 - LVM.EnableResignature=0.Refer to the VMware documentation for more information on these settings.
- Rescans all Host Bus Adapters (HBA).

Configuring the agent for Hitachi TrueCopy

This chapter includes the following topics:

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

Configuration concepts for the Hitachi TrueCopy agent

Review the configuration concepts and failure scenarios for the agent.

Resource type definition for the Hitachi TrueCopy agent

The resource type definition defines the agent in VCS.

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

```
temp str TargetFrozen  
)
```

Attribute definitions for the Hitachi TrueCopy agent

The descriptions of the agent attributes are as follows:

BaseDir	Path to the RAID Manager Command Line interface. Type-dimension: string-scalar Default: /HORCM/usr/bin.
GroupName	Name of the device group that the agent manages. Type-dimension: string-scalar
Instance	The Instance number of the device that the agent manages. Multiple device groups may have the same instance number. Do not define the attribute if the instance number is zero. Type-dimension: string-scalar
SplitTakeover	A flag that determines whether the agent permits a failover to S-VOL devices if the replication link is disconnected, that is, if P-VOL devices are in the PSUE state. See “About the SplitTakeover attribute for the Hitachi TrueCopy agent” on page 17. Type-dimension: integer-scalar Default: 0
LinkMonitor	A flag that defines whether the agent periodically attempts to resynchronize the S-VOL side if the replication link is disconnected. The agent uses the <code>pairresync</code> command to resynchronize arrays. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the <code>pairresync</code> command. Type-dimension: integer-scalar Default: 0
TargetFrozen	For internal use. Do not modify.

About the SplitTakeover attribute for the Hitachi TrueCopy agent

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

The default value for this attribute is 0.

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

In this scenario, with the SplitTakeover attribute set to 0, the agent attempts to contact the RAID manager at the P-VOL side to determine the status of the arrays. If the P-VOL side is not PSUE or not reachable, the agent proceeds with failover.

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

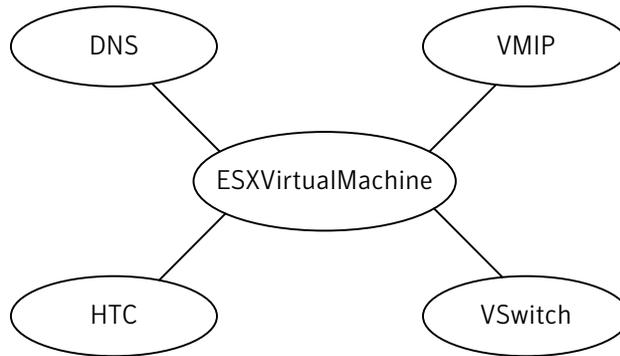
If a device group is made up of multiple devices, then, in case of a link failure, the state of each device changes on an individual basis. This change is not reflected on the device group level. Only those devices to which an application made a write after a link failure change their state to PSUE.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

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

Sample configuration for the Hitachi TrueCopy agent

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

Figure 2-1 VCS service group configuration for VMware ESX environment



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

```
HTC sgl_htc_res (  
    GroupName = VG01  
    Instance = 1  
)
```

Before you configure the agent for TrueCopy

Before you configure the agent, review the following information:

- Verify that the clustering infrastructure is in place.
If you plan to configure the agent in a global cluster, make sure the global service group for the application is configured.
For more information, see the *Veritas Cluster Server User's Guide*.
- Review the configuration concepts, which describe the agent's type definition and attributes.
See [“Configuration concepts for the Hitachi TrueCopy agent”](#) on page 15.
- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
See [“Typical Hitachi TrueCopy setup in a VCS cluster”](#) on page 10.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 19.
- Set up system zones in replicated data clusters.
See [“About configuring system zones in replicated data clusters”](#) on page 19.

About cluster heartbeats

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

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

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

About configuring system zones in replicated data clusters

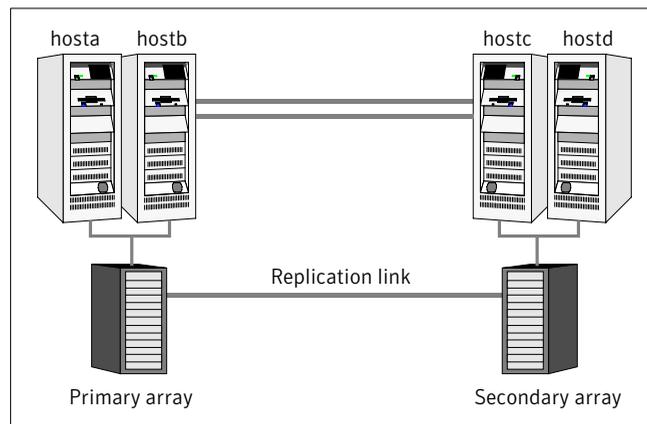
In a replicated data cluster, you can prevent unnecessary TrueCopy failover or failback by creating system zones. VCS attempts to fail over applications within the same system zone before failing them over across system zones.

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

Figure 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



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

About preventing split-brain

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

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

Configuring the agent for TrueCopy

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

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

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

See [“Sample configuration for the Hitachi TrueCopy agent”](#) on page 17.

Configuring the agent manually in a global cluster

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

To configure the agent in a global cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 If the agent resource type (HTC) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types** and select:
`/etc/VRTSvcs/conf/HTCTypes.cf.`
- 3 Click **Import**.
- 4 Save the configuration.
- 5 Add a resource of type HTC at the bottom of the service group.
Link the VMDg and HTC resources so that the VMDg resources depend on HTC.
- 6 Configure the attributes of the HTC resource.
- 7 If the service group is not configured as a global group, configure the service group using the Global Group Configuration Wizard.
See the *Veritas Cluster Server User's Guide* for more information.
- 8 Change the ClusterFailOverPolicy from the default, if necessary. Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.
- 9 Repeat step 5 through step 8 for each service group in each cluster that uses replicated data.
- 10 The configuration must be identical on all cluster nodes, both primary and disaster recovery.

Configuring the agent manually in a replicated data cluster

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

To configure the agent in a replicated data cluster

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

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

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

Managing and testing clustering support for Hitachi TrueCopy

This chapter includes the following topics:

- [Typical test setup for the Hitachi TrueCopy agent](#)
- [Testing service group migration](#)
- [Testing host failure](#)
- [Performing a disaster test](#)
- [Performing the failback test](#)
- [Failure scenarios for Hitachi TrueCopy](#)
- [Rescanning Host Bus Adapters \(HBAs\) on VCS nodes](#)

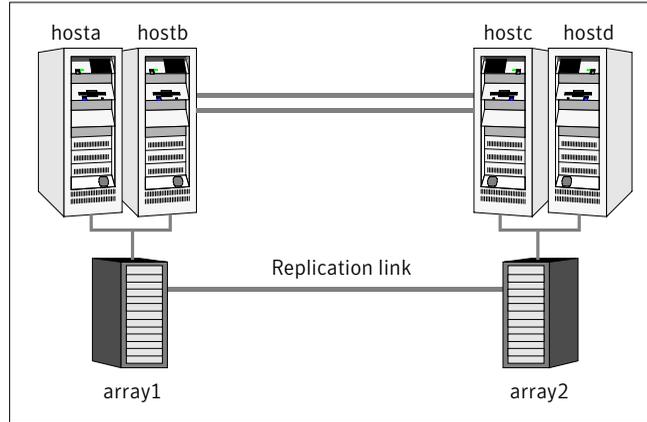
Typical test setup for the Hitachi TrueCopy agent

A typical test environment includes the following characteristics:

- Two hosts (hosta and hostb) are attached to the P-VOL Hitachi TrueCopy array.
- Two hosts (hostc and hostd) are attached to the S-VOL Hitachi TrueCopy array.
- The application runs on hosta and devices in the local array are P-VOLs in the PAIR state.
- A replicated data cluster has two dedicated heartbeat links.
A global cluster has one network heartbeat.

Figure 3-1 depicts a typical test environment.

Figure 3-1 Typical test setup



Testing service group migration

Verify the service group can migrate to different hosts in the cluster and across clusters.

To perform the service group migration test

- 1 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
Migrate the service group to a host that is attached to the same array.
- 2 Click **Switch To**, and click the system that is attached to the same array (hostb) from the menu.
The service group comes online on hostb and local volumes remain in the P-VOL/PAIR state.
- 3 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
Migrate the service group to a host that is attached to a different array.
- 4 Click **Switch To**, and click the system that is attached to the another array (hostc) from the menu.
The service group comes online on hostc and volumes there transition to the P-VOL/PAIR state, changing the original P-VOLs to S-VOLs.

- 5 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
 Migrate the service group back to its original host.
- 6 Click **Switch To**, and click the system on which the group was initially online (hosta).
 The group comes online on hosta. The devices return to the original state in step 1.

Testing host failure

In this scenario, the host where the application runs is lost. Eventually all the hosts in the system zone or cluster are lost.

To perform the host failure test

- 1 Halt or shut down the host where the application runs (hosta).
 The service group fails over to hostb and devices are in the P-VOL/PAIR state.
- 2 Halt or shut down hostb.
 In a replicated data cluster, the group fails over to hostc or hostd depending on the FailOverPolicy in the cluster.
 In a global cluster, a cluster down alert appears and gives you the opportunity to fail over the service group manually.
 The devices on the target array remain S-VOLs. They remain S-VOLs because they cannot communicate with the original primary's RAID manager, but they transition to the writable SSWS status. The failover can take some time as the RAID manager connection times out.
- 3 Reboot the two hosts that were shut down. A swap resynchronization is required to demote the original P-VOLs:

```
hares -actionHTCRes pairresync-swaps -sys system
```
- 4 Switch the service group to its original host when VCS starts.
 Do the following:
 - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
 - Click **Switch To**, and click the system on which the service group was initially online (hosta).
 The service group comes online on hosta and devices swap roles again.

Performing a disaster test

Test how robust your cluster is in case of a disaster.

To perform a disaster test

- 1 Shut down all hosts on the source side and shut down the source array.
If you can not shut down the primary array, disconnect the replication link between the two arrays and simultaneously shut down the hosts. This action mimics a disaster scenario to the secondary side.
- 2 In a replicated data cluster, the service group fails over to `hostc` or `hostd` in the following conditions:
 - All devices were originally in the PAIR state.
 - No synchronization was in progress at the time of disaster.
- 3 In a global cluster, the administrator is notified of the failure. The administrator can then initiate the failover.

Performing the failback test

You can set up your cluster for a failback test.

The failback test verifies the application can fail back to its original host after a failover to a remote site.

To perform a failback test

- 1 Reconnect the replication link and reboot the original P-VOL hosts.
- 2 Take the service group offline.
- 3 Write-disable both sides.
- 4 Manually resynchronize the device.
- 5 After the resynchronization is complete, migrate the application back to the original primary side.

Failure scenarios for Hitachi TrueCopy

Review the failure scenarios and agent behavior in response to failure.

Site disaster

In a total site failure, all hosts and the array are completely disabled, either temporarily or permanently.

In a replicated data cluster, site failure is detected the same way as a total host failure, that is, the loss of all LLT heartbeats.

In a global cluster environment, VCS detects the failure by the loss of the ICMP heartbeat between the clusters.

If a failover occurs, the online entry point of the TrueCopy agent runs the `horctakeover` command. The RAID manager waits for the timeout in trying to contact its peer RAID manager daemon before taking over the disks. This wait can cause delay in the failover. This timeout is defined in the device group's instance's configuration file. Make sure the value of the `OnlineTimeout` entry point of the HTC type is greater than the RAID manager timeout.

The online entry point detects whether any synchronization was in progress when the source array was lost. Since the target devices are inconsistent until the synchronization completes, the agent does not write-enable the devices, but it times out and faults. You must restore consistent data from a snapshot or tape backup.

All host or all application failure

Even if both arrays are operational, the service group fails over in the following conditions:

- All hosts on the P-VOL side are disabled.
- The application cannot start successfully on any P-VOL host.

In replicated data cluster environments, the failover can be automatic, whereas in global cluster environments failover requires user confirmation by default.

In both replicated data cluster and global cluster environments, multiple service groups can fail over in parallel.

TrueCopy does not provide any serialization restrictions on simultaneous device group failover. However, the `horctakeover` command makes an attempt to contact the RAID manager on the original P-VOL when performing a failover. In such a case, if the RAID manager is inaccessible, failover is delayed until the surviving RAID manager's connect timeout expires. This timeout is defined in the configuration file for the particular instance.

Replication link failure

Hitachi arrays send an alert in the following situations:

- When the array detects a replication link failure
- When any P-VOLs, on which data has been written, transition from the PAIR state to the PSUE state

In fence levels never and async, a replication link failure does not compromise the application’s ability to write to its local devices. The arrays start tracking changed regions on disk in preparation for resynchronization when the link is restored.

The devices do not automatically resynchronize when the link is restored, nor do they change state when the restoration is detected. An administrator can resynchronize the devices, either from the command line or by running a configured action using the agent’s action entry point.

[Table 3-1](#) shows the situations that require administrative action after you repair a link failure.

These actions depend on the fence level and any events that occurred during the failure.

Table 3-1 Replication link 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.
Application 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.

Split-brain in a TrueCopy environment

When split-brain occurs in a replicated database cluster, VCS assumes a total disaster because the P-VOL side hosts and array are unreachable. VCS attempts to start the application on the secondary site. Once the heartbeats are restored, VCS stops the applications on one side and restarts the VCS engine (HAD). This action eliminates concurrency violation of the same group being online at two places simultaneously.

Administrators must resynchronize the volumes manually using the `pairresync` commands.

In a global cluster, you can confirm the failure before failing over the service groups. You can check with the site administrator to identify the cause of the failure. If a fail over mistakenly occurs, the situation is similar to the replicated data cluster case. However, when the heartbeat is restored, VCS does not stop HAD at either site. VCS forces you to choose which group to take offline. You must resynchronize the data manually.

Rescanning Host Bus Adapters (HBAs) on VCS nodes

While executing rescan HBA operations on VCS nodes, use the following guidelines to make sure that the LVM settings conform to VCS requirements:

- Note the current values of the following LVM settings:
 - DisallowSnapshotLun
 - EnableResignature
- While rescanning snapshot LUNs for new datastores, set EnableResignature=1.
- If rescanning with DisallowSnapshotLun=0 and EnableResignature=0, make sure that snapshot LUNs that do not contain already resignatured datastores are not presented to the server.

Setting up a fire drill

This chapter includes the following topics:

- [About fire drills](#)
- [About the HTCSnap agent](#)
- [Before you configure the fire drill service group](#)
- [Additional requirements for running a fire drill in an ESX environment](#)
- [Configuring the fire drill service group](#)
- [Configuring the new virtual machine for the fire drill](#)
- [Running the fire drill](#)
- [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, 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 5.0 MP1 RP2, which is a component of Veritas Storage Foundation.

Note: The agent does not support fire drills for asynchronous configurations.

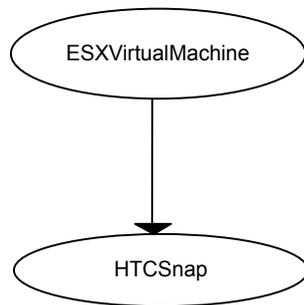
Note: In an ESX environment, you can also perform a fire drill locally to generate a last known good copy of your application data. If the replicated data produced by the fire drill tests successfully, it is the last known good copy.

About the HTCSnap agent

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

Refer to the [Figure 4-1](#) for the HTCSnap agent dependency tree.

Figure 4-1 HTCSnap agent dependency tree



HTCSnap agent functions

The HTCSnap agent performs the following functions:

- | | |
|---------|---|
| online | Destroys any existing snapshot and takes a new snapshot.
See “About the agent's online function” on page 33. |
| offline | Removes the lock file created by the online function. |
| 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.

About the agent's online function

The agent's online function performs the following actions:

- Resynchronizes snapshots LUNs with source LUNs.
- If TargetResName is of type HTC (remote fire drill), the agent suspends TrueCopy replication to get a consistent snapshot.
- Splits the ShadowImage replication into PSUS-SSUS state in order to obtain a new writeable snapshot.
- If TargetResName is of type HTC (remote fire drill), the agent resumes Truecopy replication between the arrays.
- Scans for new datastores with following VMware LVM settings.
EnableResignature = 1;
DisallowSnapshotLun = 0;
- Makes the snapshot of the virtual machine ready for the fire drill. Updates the dependent ESXVirtualMachine resource with the full path to the snapshot of the .vmx file.
See [“Processing the snapshot”](#) on page 33.
See [“Configuring the fire drill service group”](#) on page 38.
- Creates a lock file to indicate that the resource is online.

Processing the snapshot

Processing the snapshot requires the following actions;

- Temporarily sets the VMware LVM settings and scans for new datastores. Refer to the VMware documentation for more information on these settings.
- Registers the new virtual machine with the ESX host using the new snapshot datastore that was created.
- Populates the CfgFile attribute of the ESXVirtualMachine resource (parent resource) in the fire drill service group. The display name of the new virtual machine is the value of the vmname attribute, suffixed by _snapshot. The suffix clearly identifies the virtual machine as the snapshot virtual machine.

- Generates new MAC addresses for each vmnic and patches the configuration files of the new virtual machine with the new MAC addresses.
- Disconnects the new virtual machine from the public switches by setting the ethernet[n].networkName attribute in the vmx file to None. This action enables you to configure IP address and other parameters in the guest OS of the virtual machine after it boots up.

Resource type definition for the HTCSnap agent

Following are the resource type definitions for the HTCSnap agent:

```
type HTCSnap (  
    static keylist RegList = { MountSnapshot, UseSnapshot }  
    static str ArgList[] = { TargetResName, MountSnapshot,  
        UseSnapshot, RequireSnapshot, ShadowInstance }  
    str TargetResName  
    int ShadowInstance  
    int MountSnapshot = 1  
    int UseSnapshot = 1  
    int RequireSnapshot = 1  
    temp str Responsibility  
    temp str FDFile  
)
```

HTCSnap agent attributes

Configure the following attributes to customize the behavior of the HTCSnap agent:

- | | |
|----------------|--|
| 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 (for remote fire drill) or the TrueCopy P-VOL group (for local fire drill). Type-dimension: integer-scalar |
|----------------|--|

TargetResName	<p>Name of the resource managing the LUNs to be snapshot.</p> <p>Set this attribute to the name of HTC resource for running the remote firedrill.</p> <ul style="list-style-type: none"> ■ Set this attribute to the name of the ESXVirtualMachine resource for running the local fire drill. The local fire drill runs on the primary site and the remote fire drill runs on the secondary site. <p>Type-dimension: string-scalar</p>
UseSnapshot	<p>Specifies whether the HTCSnap resource takes a local snapshot of the target array.</p> <p>Default Value : 1</p> <p>Set this attribute to 1. The current release of the agent supports only the default value.</p> <p>Type-Dimension: integer-scalar</p>
RequireSnapshot	<p>Specifies whether the HTCSnap resource must take a snapshot before coming online.</p> <p>Default Value : 1</p> <p>Set this attribute to 1. The current release of the agent supports only the default value.</p> <p>Type-Dimension: integer-scalar</p>
MountSnapshot	<p>Specifies whether the resource uses the snapshot to bring the service group online.</p> <p>Default Value : 1</p> <p>Set this attribute to 1. The current release of the agent supports only the default value.</p> <p>Type-Dimension: integer-scalar</p>

Internal attributes

Do not modify internal attributes:

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	Do not modify. For internal use only. Used by the agent to locate the latest fire drill report. Type-dimension: temporary string
--------	--

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 as shown in the main.cf file.

See [“Resource type definition for the HTCSnap agent”](#) on page 34..

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

```
HTCSnap sgl_htcsnap_res{
    TargetResName = sgl_htc
    ShadowInstance = 14
    UseSnapshot = 1
    RequireSnapshot = 1
    MountSnapshot = 1
}
```

Before you configure the fire drill service group

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

- Make sure the application service group is configured with a HTC resource.
- Make sure the infrastructure to take snapshots is properly configured. This process involves creating the Shadow Image pairs.
- Make sure the infrastructure to configure hardware snapshots is properly configured.
- If you plan to use Gold or Silver configuration, make sure ShadowImage for TrueCopy is installed and configured at the target array.
- For the Gold configuration, you must use Veritas Volume Manager to import and deport the storage.
- You can use the Silver configuration only with ShadowImage pairs that are created with the `-m noread` flag to the `paircreate` command. A fire drill uses the `-E` flag to split the pairs, which requires a 100% resynchronization. The Silver mode that preserves the snapshots as `noread` after a split.

- The name of the ShadowImage device group must be the same as the replicated device group for both replicated and non-replicated LUNs that are to be snapshot. The instance number may be different.
- Make sure the HORC instance managing the S-VOLs runs continuously; the agent does not start this instance.
- For non-replicated devices:
 - You must use Veritas Volume Manager.
On HP-UX, you must use Veritas Volume Manager 5.0 MP1.
 - For Gold configuration to run without the Bronze mode, set the RequireSnapshot attribute to 1.

Additional requirements for running a fire drill in an ESX environment

Follow these guidelines for fire drills in an ESX environment:

- Ensure that each service group contains only one virtual machine resource to be able to support fire drill.
- Symantec recommends dedicating one node in the cluster for fire drills.
- Configure fire drill service groups on this node. Do not configure replication resources on this node.
- Configure your array such that snapshot LUNs (i.e. target LUNs on which hardware snapshots are taken) are visible to this node only; the other nodes must not see snapshot LUNs. The restriction occurs because other nodes in the cluster may have LVM settings `LVM.EnableResignature = 0` and `LVM.DisAllowSnapshotLUN = 0` set by the replication agents, in which case VMWare recommends that no snapshot LUNs should be exposed to such ESX hosts.
- Install VMware tools to all virtual machines in the fire drill configuration. Otherwise, the missed virtual machine heartbeats cause the virtual machine resource to fault.
- Ensure that no virtual machines are configured with raw device mapping.
- Ensure that the resources configured within the fire drill service group are not marked as critical. Not marking the fire drill service group resources as critical prevents service group failover.

Configuring the fire drill service group

This section describes how to use Cluster Manager (Java Console) to create the fire drill service group.

To create the fire drill service group

- 1 Open the Veritas Cluster Manager (Java Console).
- 2 Log on to the cluster and click **OK**.
- 3 Click the Service Group tab in the left pane and click the **Resources** tab in the right pane.
- 4 Right-click the cluster in the left pane and click **Add Service Group**.
- 5 In the Add Service Group dialog box, provide information about the new service group
- 6 In Service Group name, enter a name for the fire drill service group
- 7 From the Available Systems box, select the node that is dedicated for fire drills.
- 8 Click the fire drill service group in the left pane and click the Resources tab in the right pane.
- 9 Add a resource of type HTCSnap and configure its attributes.
- 10 Add a resource of type ESXVirtualMachine and configure its attributes.
See “[Sample configuration for a fire drill service group](#)” on page 36..
- 11 Link resources such that the ESXVirtualMachine resource depends on the HTCSnap resource.

Configuring the new virtual machine for the fire drill

To configure the new virtual machine for the fire drill

- 1 Bring the fire drill service group online on the node dedicated for fire drills.
Reboot the new virtual machine. This virtual machine now comes up without any connections to any public network. You may see some errors and the application may report a faulted status; ignore these errors.
See “[About the agent's online function](#)” on page 33.
- 2 Assign a unique IP address and host name to the new virtual machine.

- 3 Take the fire drill service group offline.
- 4 If you want to connect the new virtual machine to the internet, connect the virtual machine to at least one public switch while it is shut down.

Running the fire drill

Before running a fire drill on the secondary site, ensure that it does not have any global service groups that have failed or switched over to the secondary site in the online state.

Restart the two RAID managers configured for the ShadowImage pair so that the latest device bindings are obtained.

Additionally, ensure that you do not run a local fire drill on the secondary site, even if the secondary site has become the new primary site after a failover to the remote cluster. These restrictions occur because of ESX server behavior during datastore rescans with LVM settings: `EnableResignature = 1`; `DisAllowSnapshotLun = 0`; which may cause datastores to be resignatured. Refer to the *Release Notes* for more information.

You are now ready to run the fire drill.

To run the fire drill

- 1 Bring the fire drill service group online. The configured applications come up using snapshot data.
- 2 Ensure the validity of your application by using it to perform tasks that are appropriate for the application.

Being able to use the application indicates a successful fire drill and ensures that your replicated data is valid. The snapshot is stored in the current state until you run the next fire drill.

- 3 Take the fire drill service group offline.

Failing to take the fire drill service group 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.

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.

A

- application failure 27
- attribute definitions
 - Hitachi TrueCopy agent 16

C

- cluster
 - heartbeats 19

D

- disaster test 26

F

- failback test 26
- failure scenarios
 - all application failure 27
 - all host failure 27
 - replication link failure 27
 - total site disaster 26
- FDFile attribute 36
- fire drill
 - about 31
 - configuration wizard 36
 - running 39
 - service group for 36
 - SRDFSnap agent 32

H

- Hitachi TrueCopy agent
 - attribute definitions 16
 - type definition 15
- host failure 27

M

- migrating service group 24
- MountSnapshot attribute 35

R

- replication link failure 27

- RequireSnapshot attribute 35
- resource type definition
 - Hitachi TrueCopy agent 15
 - SRDFSnap agent 34
- Responsibility attribute 35

S

- sample configuration 17
- service group
 - migrating 24
- split-brain
 - handling in cluster 20
 - handling in clusters 28
- SRDFSnap agent
 - about 32
 - attribute definitions 34
 - operations 32
 - type definition 34
- SRDFSnap agent attributes
 - MountSnapshot 35
 - RequireSnapshot 35
 - UseSnapshot 35

T

- testing
 - disaster 26
 - failback 26
- total site disaster 26
- type definition
 - Hitachi TrueCopy agent 15
 - SRDFSnap agent 34

U

- UseSnapshot attribute 35