VERITAS™

VERITAS Cluster Server 4.1 Enterprise Agent for Hitachi TrueCopy

Installation and Configuration Guide

HP-UX

N16820G

June 2005

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Legal Notice

Copyright © 2005 VERITAS Software Corporation. All rights reserved. VERITAS, VERITAS Software, VERITAS Cluster Server, the VERITAS Logo, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS, the VERITAS logo, and Cluster Server Reg. U.S. Pat. & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650–527–8000 Fax 650–527–2908
www.veritas.com

Contents

Preface

This document describes how to install and configure the VERITAS Cluster Server (VCS) enterprise agent for Hitachi TrueCopy.

If this document is dated more than six months prior to the date you are installing your enterprise agent, contact VERITAS Technical Support to confirm you have the latest supported versions of the application and operating system.

How This Guide is Organized

- Chapter 1. "Introduction" on page 1 introduces the VCS enterprise agent for Hitachi TrueCopy and describes its operations.
- Chapter 2. "Installing the Hitachi TrueCopy Agent" on page 5 describes provides instructions on installing the Hitachi TrueCopy agent.
- Chapter 3. "Configuring the Hitachi TrueCopy Agent" on page 7 describes key configuration concepts and provides instructions on configuring the agent.
- Chapter 4. "Managing and Testing Clustering Support for Hitachi TrueCopy" on page 21 provides test scenarios and expected outcomes. It also describes how to uninstall the agent.
- Chapter 5. "Setting Up a Fire Drill" on page 25 describes how you can test the fault-readiness of the disaster recovery environment by running a fire drill.

VERITAS Cluster Server Documentation

The following documents, along with the online help and the Release Notes, comprise the VCS documentation for this release:

Title	File Name
VERITAS Cluster Server Installation Guide	`vcs_install.pdf`
VERITAS Cluster Server User's Guide	`vcs_users.pdf`
VERITAS Cluster Server Bundled Agents Reference Guide	`vcs_bundled_agents.pdf`
VERITAS Cluster Server Agent Developer's Guide	`vcs_agent_dev.pdf`

See the Release Notes for a complete list of documents, including VCS enterprise agent guides.

Conventions

The following conventions apply throughout the documentation set.

Typeface/Font	Usage
bold	names of screens, windows, tabs, dialog boxes, options, buttons
italic	new terms, book titles, emphasis, variables in tables or body text
`Courier`	computer output, command references within text
`Courier` (bold)	command-line user input, keywords in grammar syntax
`Courier` (bold, italic)	variables in a command
#	UNIX superuser prompt (all shells)

Getting Help

For technical assistance, visit http://support.veritas.com and select phone or email support. This site also provides access to resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service. Use the Knowledge Base Search feature to access additional product information, including current and past releases of VERITAS documentation.

For license information, software updates and sales contacts, visit https://my.veritas.com/productcenter/ContactVeritas.jsp. For information on purchasing product documentation, visit http://webstore.veritas.com.

Documentation Feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to clusteringdocs@veritas.com. Include the title and part number of the document (located in the lower left corner of the title page), and chapter and section titles of the text on which you are reporting. Our goal is to ensure customer satisfaction by providing effective, quality documentation. For assistance with topics other than documentation, visit http://support.veritas.com.

Introduction 1

The VCS enterprise agent for Hitachi TrueCopy provides failover support and recovery in environments employing TrueCopy to replicate data between Hitachi disk arrays.

About the Hitachi TrueCopy Agent

The VCS enterprise agent for Hitachi TrueCopy monitors and manages the state of replicated devices attached to local hosts. The agent ensures that the system on which the TrueCopy resource is online has safe and exclusive access to the configured devices.

The agent can be used in single VCS replicated data clusters and multi-cluster environments set up using the VCS Global Cluster Option.

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

When replicating between Lightning arrays, the agent supports the following fence levels: *data*, *never*, and *async*.

When replicating between Thunder arrays, the agent supports the following fence levels: *data* and *never*.

Supported Software and Hardware

The agent supports all versions of the Hitachi RAID Manager. It supports TrueCopy on all microcode levels on all Lightning arrays, provided the host/HBA/array combination is in Hitachi's hardware compatibility list. The agent supports Sun StorEdge 9900 and Hewlett-Packard XP arrays with TrueCopy rebranded as Continuous Access. The agent supports all fence levels on 9900 arrays and supports synchronous replication on the 9500 series.

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

The agent supports VERITAS Cluster Server 4.1 and 3.5. The agent also supports VERITAS Global Cluster Manager 3.5.1. The agent does not support ShadowImage fire drills with VCS 3.5 or GCM 3.5.1.

Typical Setup

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

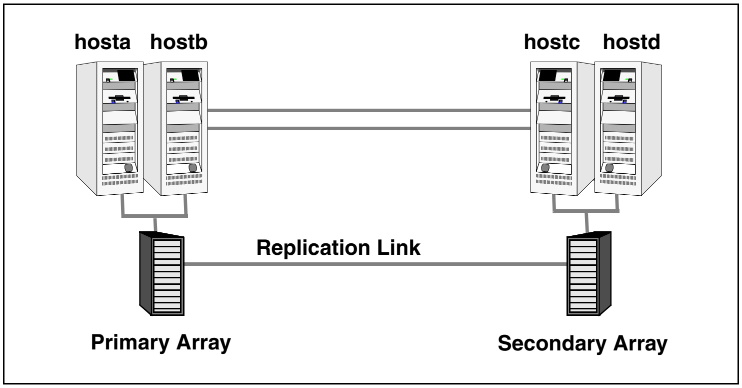

- ✔ The *primary array* comprising one or more *P-VOL hosts* directly attached via SCSI or Fibre Channel to a Hitachi array containing TrueCopy P-VOL volumes.
- ✔ The *secondary array* comprising one or more *S-VOL hosts* directly attached via SCSI or Fibre Channel to a second Lightning array containing TrueCopy S-VOL devices. These devices are paired with the P-VOL devices in the primary array.

  These hosts and the array must be at a significant distance apart from the primary side to survive a disaster that may occur there.
- ✔ Network heartbeats, using LLT or TCP/IP, between the two data centers to determine their health.
- ✔ In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them by dual, dedicated networks that support LLT.
- ✔ In a global cluster environment, you must attach all hosts in a cluster to the same array.

Agent Operations

The Hitachi TrueCopy agent performs the following operations:

Operation (Entry Point)	Description
online	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 makes the devices writable for the application. 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 health of the original primary and the RAID Manager timeouts 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 S-VOL devices are in the COPY state, the agent waits until the synchronization from the primary has completed before running the `horctakeover` command or until the OnlineTimeout period of the entry point has expired, in which case the resource faults.
offline	The agent removes the lock file on the device. The agent does not run any TrueCopy commands because taking the resource offline is not indicative of an intention to give up the devices.
monitor	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. The monitor entry point does not examine the state of the devices or the state of the replication link between the arrays.
open	Removes the lock file on the system on which this entry point is called. This prevents potential concurrency violation if the group fails over to another node. **Note** The agent does not remove the lock file if the agent was started after an `hastop -force` command.
clean	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, potentially leaving the devices in an unusable state.

Operation (Entry Point)	Description
info	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.
action	Resynchronizes devices from the VCS command line after various connectivity failures are detected and corrected. The agent supports the following actions: ◆ pairdisplay—Displays information about all devices. ◆ pairresync—Resynchronizes the S-VOLs. ◆ pairresync-swaps—Promotes the S-VOLs to P-VOLs and resynchronizes the original P-VOLs. ◆ localtakeover—Makes the local devices write-enabled.

Installing the Hitachi TrueCopy Agent 2

You must install the Hitachi TrueCopy enterprise agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster. These instructions assume that you have already installed VCS.

▼ To install the agent

1. Insert the disc into a drive connected to the host.

2. Create a mount point directory, `/cdrom`, if it does not exist. The directory must have read-write permissions.

3. Determine the block device file for the disc drive:

   ```
   # ioscan -fnC disk
   ```

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

4. Start the Portable File System (PFS).

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

5. Mount the disc:

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

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

6. Install the agent software:

   ```
   # swinstall -s /cdrom/depot VRTSvcstc
   # swinstall -s /cdrom/depot VRTScstcw
   # swinstall -s /cdrom/depot VRTScsfdw
   ```

Configuring the Hitachi TrueCopy Agent 3

You can adapt most applications configured in VCS to a disaster recovery environment by:

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

Configure the volumes of a TrueCopy device group as resources of type HTC.

Before Configuring the TrueCopy Agent

- ✔ Verify the agent is installed on all nodes in the cluster.
- ✔ Verify the hardware infrastructure required for the agent is in place. See "Typical Setup" on page 2 for more information.
- ✔ Make sure the cluster has an effective heartbeat mechanism in place. See "Cluster Heartbeats" on page 9 for more information.
- ✔ Review the agent's resource type definition and its attribute definitions.
- ✔ Review the section "Configuration Concepts" on page 10, which presents information about how VCS behaves during failover and how you can set attributes to customize VCS behavior.

Resource Type Definition

```
type HTC (
  static str ArgList[] = { BaseDir, GroupName, Instance }
    static int NumThreads = 1
    static keylist SupportedActions = { localtakeover, pairresync,
      pairresync-swaps, pairdisplay }
    NameRule = resource.GroupName
    str BaseDir = "/HORCM/usr/bin"
    str GroupName
    int Instance
    )
```

Attribute Definitions

Attribute	Type-Dimension	Description
BaseDir	string-scalar	Path to the RAID Manager Command Line interface. Default is `/HORCM/usr/bin`.
GroupName	string-scalar	Name of the device group managed by the agent.
Instance	integer-scalar	The Instance number of the device group managed by the agent. Multiple device groups may have the same instance number. Since the default value of any integer attribute in VCS is 0 (zero), do not define the attribute if the instance number is zero.

Sample Configuration

The following dependency graph shows a VCS service group that has a resource of type TrueCopy. The DiskGroup resource depends on the TrueCopy resource.

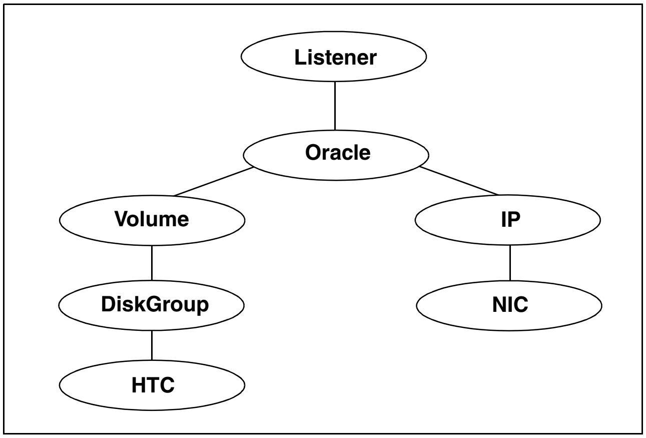

A resource of type TrueCopy may be configured as follows in `main.cf`:

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

Cluster Heartbeats

In a replicated data cluster, robust heartbeating is accomplished through dual, dedicated networks over which LLT runs. Additionally, a low-priority heartbeat may be configured across public networks.

In a global cluster, network heartbeating is accomplished by sending ICMP pings over the public network between the two sites. VCS global clusters minimize the risk of split-brain by sending ICMP pings to highly available IP addresses and by notifying administrators when the sites cannot communicate.

Hitachi arrays do not support a native heartbeating mechanism between the arrays. The default behavior of the arrays is to send a support message when a replication link failure is detected. Based on the type of failure and the state of the devices at the time the failure is corrected, you can take appropriate action to recover from the failure to keep the devices in a synchronized state. The TrueCopy agent supports various actions that can automate the resynchronization of devices after a replication link outage is corrected.

Configuration Concepts

This section describes some failure scenarios and provides guidelines on how to configure the agent in VCS.

Individual Component Failure

In a replicated data cluster, you can instruct VCS to give preference to hosts connected to the same array during failover. This avoids unnecessary TrueCopy failover or failback and is accomplished by creating system zones such that hosts attached to an array are part of the same system zone.

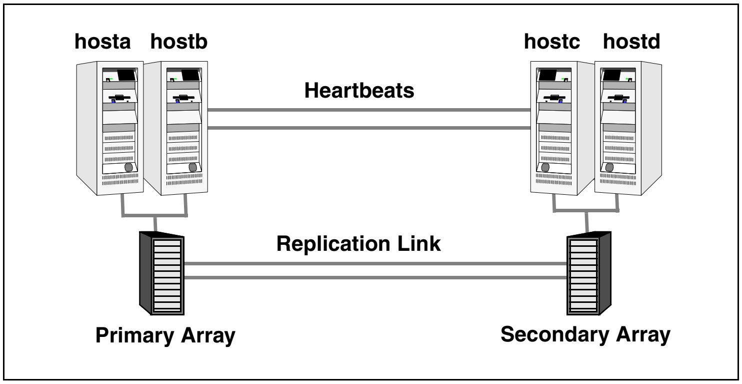

In this sample, hosta and hostb are in one system zone, and hostc and hostd are in another system zone. The SystemZones attribute enables you to create these zones.

You can modify the SystemZones attribute using the following command:

```
# hagrp -modify grpname SystemZones hosta 0 hostb 0 hostc 1 hostd 1
```

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

This command creates two system zones: zone 0 with hosta and host b, zone 1 with hostc and hostd.

System zones are not required in global clusters because failover will occur on a remote cluster if all local targets have been exhausted.

All Host or All Application Failure

If all hosts on the P-VOL side are disabled or if the application cannot start successfully on any P-VOL hosts, but both arrays are operational, the service group fails over.

In replicated data cluster environments, the failover can be automatic, whereas in global cluster environments, failover by default requires user confirmation. Multiple service groups can fail over in parallel; TrueCopy does not provide any serialization restrictions on simultaneous device group failover.

However, since the `horctakeover` command makes an attempt to contact the RAID manager on the original P-VOL when performing a failover, if the RAID manager is inaccessible, failover will be delayed until the surviving RAID manager's connect timeout expires. This timeout is defined in the configuration file for the particular instance.

Total Site Disaster

In a total site failure, all hosts and the Hitachi 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 Icmp heartbeat between the clusters.

If a failover occurs, the online entry point of the TrueCopy agent runs the horctakeover command; the failover may be delayed because the RAID manager waits for the timeout in trying to contact its peer RAID manager daemon before taking over the disks. 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 TrueCopy 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 ShadowImage or tape backup.

Replication Link Failure

Hitachi arrays send an alert in the following situations:

- When a replication link failure is detected
- 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 once 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. The following situations require administrative action after a link failure is repaired. These actions depend on the fence level and any events that may have occurred during the failure.

Event	**Fence Level**	**Recommended Action**
Link fails and is restored, but application does not fail over	never, async	Run the `pairresync` action to resynchronize the S-VOLs
Link fails and application fails to the S-VOL side	never, async, or data	Run the `pairresync-swaps` 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 `localtakeover` action to write-enable the local devices. Clear faults and restart service group.

Split-brain

Split-brain occurs when all heartbeat links between the source and target hosts are cut and each side mistakenly thinks the other side is down. To minimize the effects of split-brain, it is best if the cluster heartbeat links pass through similar physical infrastructure as the replication links so that if one breaks, so does the other.

In a replicated data cluster, VCS attempts to start the application assuming a total disaster because the P-VOL hosts and array are unreachable. Once the heartbeats are restored, VCS stops the applications on one side and restarts the VCS engine (HAD) there to eliminate concurrency violation of the same group being online at two places simultaneously. Administrators must resynchronize the volumes manually using the `pairresync` commands.

In global cluster environments, administrators 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 you do mistakenly fail over, 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. Again, resynchronization must be performed manually. If it is physically impossible to place the heartbeats alongside the replication links, there is a possibility that the cluster heartbeats are disabled, but the replication link is not. A failover transitions the original P-VOLs to S-VOLs and vice-versa. In this case, the original running application faults because its underlying volumes become write-disabled. This causes the service group to fault and VCS tries to fail it over to another host, causing the same consequence in the reverse direction. This phenomenon, sometimes called *ping-pong*, continues until the group comes online on the final node. This situation can be avoided by setting up your infrastructure such that loss of heartbeat links also mean the loss of replication links.

Configuring the Agent Using the Wizard

This section describes how to use the wizard to configure the Hitachi TrueCopy agent in an application service group.

1. Run the wizard on a system attached to the array. Verify Hitachi RAID Manager is installed on the system where you run the wizard.

2. Set the *DISPLAY* variable and start the HTC Configuration wizard as `root`.

   ```
   # hawizard htc
   ```

3. Read the information on the Welcome screen and click **Next**.

4. In the Wizard Options dialog box, select the application service group to which you want to add an HTC resource.

   Note The wizard displays service groups having disk group resources; it does not display service groups having HTC resources.

5. In the Select Device Group dialog box, specify the device group from the Hitachi array for which the HTC resource is to be added.

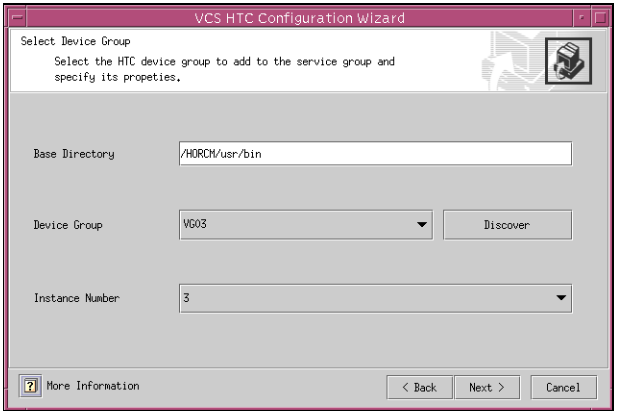

VCS HTC Configuration Wizard

Select Device Group

Select the HTC device group to add to the service group and specify its propeties.

Base Directory: /HORCM/usr/bin

Device Group: VG03

Instance Number: 3

   a. In the **Base Directory** field, specify the path where the CLI package for the Hitachi array is installed. The default location is `/HORCM/usr/bin`.

   b. From the **Device Group** list, select a device group.

      If the wizard does not display the required device groups, verify the HTC instance is running and click **Discover**.

   c. Select the instance number for the device group.

   d. Click **Next**.

6. Select the replicated diskgroups and click **Next**.

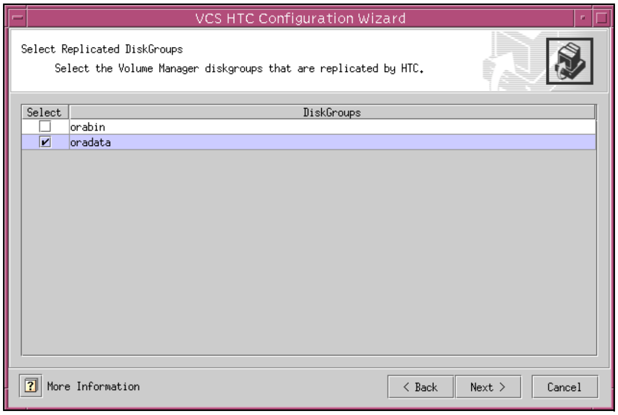

Select Replicated DiskGroups

Select the Volume Manager diskgroups that are replicated by HTC.

Select	DiskGroups
☐	orabin
☑	oradata

If you are adding an HTC resource in a service group configured in a replicated data cluster, proceed to the next step. Otherwise, proceed to step 8 on page 17.

7. In the System Zones Selection dialog box, specify the systems in each zone of a replicated data cluster.

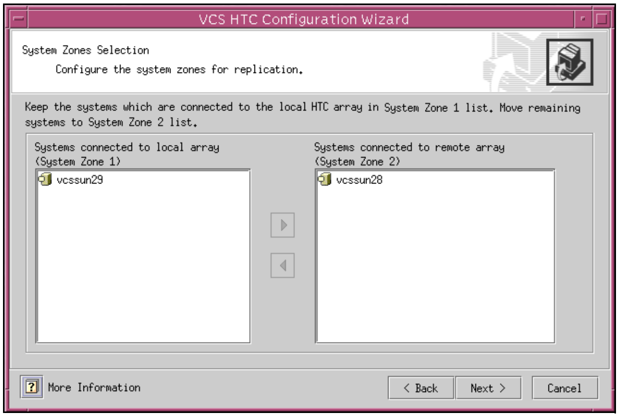

System Zones Selection

Configure the system zones for replication.

Keep the systems which are connected to the local HTC array in System Zone 1 list. Move remaining systems to System Zone 2 list.

Systems connected to local array (System Zone 1)	Systems connected to remote array (System Zone 2)
vcssun29	vcssun28

   a. If you had configured SystemZones in the application service group, verify the configuration. Use the arrows to move systems to their respective zones.

   b. Click **Next**.

8. In the Service Group Summary dialog box, review the service group configuration and change the name of the HTC resource, if desired.

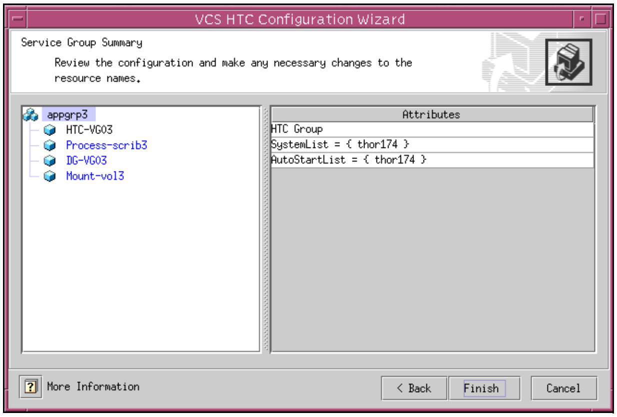

VCS HTC Configuration Wizard

Service Group Summary

Review the configuration and make any necessary changes to the resource names.

- appgrp3
  - HTC-VG03
  - Process-scrib3
  - DG-VG03
  - Mount-vol3

Attributes
HTC Group
SystemList = { thor174 }
AutoStartList = { thor174 }

   a. To change the name of the HTC resource, select the resource name and either click it or press the F2 key. Press Enter after editing the resource name. To cancel editing a resource name, press Esc.

   b. Click **Finish**.

   The wizard starts running commands to add the HTC resource to the service group. Various messages indicate the status of these commands.

9. In the Completing the HTC Configuration Wizard dialog box, select the check box to bring the service group online on the local system.

10. Click **Close**.

Configuring the Agent Manually

This section describes how to configure the agent using the Java Console in global and replicated data clusters.

Configuring the Agent in a Global Cluster

1. If the agent's resource type is not added to your configuration, add it.

    a. Start Cluster Manager and log on to the cluster.

    b. From the Cluster Explorer **File** menu, choose **Import Types** and select `/etc/VRTSvcs/conf/HTCTypes.cf.`

    c. Click **Import**.

    d. Save the configuration.

Note You can also add the resource type using the command `/etc/VRTSvcs/conf/sample_htc/addHTCType.sh`.

2. Perform the following tasks for each service group in each cluster that uses replicated data:

    a. Add a resource of Type HTC at the bottom of the service group. See "Sample Configuration" on page 9 for more information.

    b. Configure its attributes. See "Attribute Definitions" on page 8 for more information about these attributes.

    c. If the service group is not configured as a global group, configure it using the Global Group Configuration Wizard. See the *VERITAS Cluster Server User's Guide* for more information.

    d. Change the ClusterFailOverPolicy from the default. if necessary. VERITAS recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.

Configuring the Agent in a Replicated Data Cluster

1. If the agent's resource type is not added to your configuration, add it.

   a. Start Cluster Manager and log on to the cluster.

   b. From the Cluster Explorer **File** menu, choose **Import Types** and select `/etc/VRTSvcs/conf/HTCTypes.cf.`

   c. Click **Import**.

   d. Save the configuration.

Note You can also add the resource type using the command `/etc/VRTSvcs/conf/sample_htc/addHTCType.sh`.

2. Perform the following tasks for each service group that uses TrueCopy to replicate data:

   a. Add a resource of type HTC at the bottom of the service group. See "Sample Configuration" on page 9 for more information.

   b. Configure its attributes. See "Attribute Definitions" on page 8 for more information about these attributes. Note that some attributes may need to be localized to reflect values for hosts attached to different arrays.

   c. Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array. See "Individual Component Failure" on page 10 for more information.

Configuring the Agent in Clusters Running VCS 3.5

To configure a replicated data cluster with VCS 3.5, follow the instructions in“Configuring the Agent in a Replicated Data Cluster” on page 19.

To configure two separate VCS 3.5 clusters and manage them with GCM 3.5.1, run the wizard at `/opt/VRTSgcm/bin/HTC_config`. The wizard guides you through the process of configuring a GCM Global Application and GCM Events to manage application failover across clusters. Enter the name of the service group in each site representing the same application when prompted. The wizard builds the Global Application and Events for the application. Run the wizard for each application pair that uses Hitachi Truecopy as for replication. You must configure a resource of type HTC in your service groups to fail over the replicated devices.

Managing and Testing Clustering Support for Hitachi TrueCopy 4

After configuring the TrueCopy agent in a VCS environment, you can perform some basic tests to verify the implementation. This chapter describes some test scenarios and expected behavior.

These tests assume the following environment:

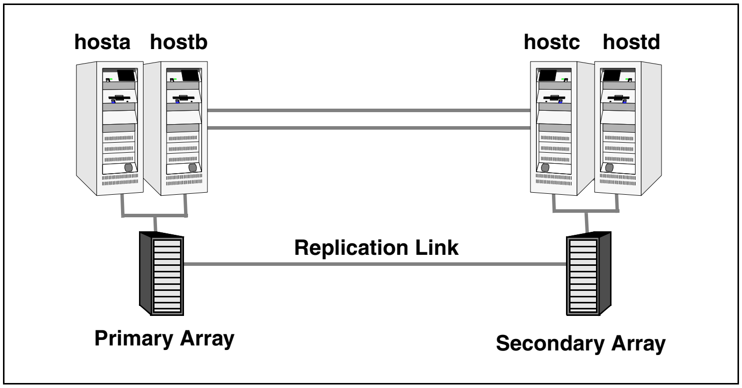

Two hosts (hosta and hostb) are attached to the primary array, and the other hosts are attached to the secondary array. The application is running on hosta and devices in the local array are P-VOLs in the PAIR state.

A replicated data cluster has two dedicated heartbeat links; a global cluster has one network heartbeat. The test scenario is similar for both environments.

Service Group Migration

Verify the service group can migrate to different hosts in the cluster by performing the following tests.

▼ To perform the service group migration test

1. Migrate the service group to a host attached to the same array:

   a. In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

   b. Click **Switch To** and click a system attached to the same array (hostb).

   The service group comes online on hostb and local volumes remain in the P-VOL/PAIR state.

2. Migrate the service group to a host attached to a different array:

   a. In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

   b. Click **Switch To**, and click **Remote switch**.

   c. Select a system attached to another array (hostc) and click **OK**.

   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.

3. Migrate the service group back to its original host:

   a. In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

   b. Click **Switch To**, and click **Remote switch**.

   c. Select the system on which the service group was originally online (hosta) from the menu.

   The group comes online on hosta. The devices return to the original state in step 1.

Host Failure

In this scenario, the host on which the application is running is lost and eventually all hosts in the system zone or cluster are lost.

▼ **To perform the host failure test**

1. Shut down the host on which the application is running:

   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 value of the FailOverPolicy attribute in the cluster.

   In a global cluster, a cluster down alert appears and gives you the opportunity to fail over the service group manually.

   In both environments, the devices on the target array remain S-VOLs because they cannot communicate with the original primary's RAID manager, but they transition to the writable SSWS status. Also, the failover may 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 -action HTCRes pairresync-swaps -sys system
   ```

4. Migrate back when the devices transition from COPY to PAIR:

   a. In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

   b. Click **Switch To**, and click the system where the service group was initially online (hosta).

   The service group comes online on hosta and devices swap roles again.

Disaster Test

Shut down all hosts on the source side and shut down the source array. If shutting down the primary array is not feasible, disconnect the replication link between the two arrays while simultaneously shutting down the hosts; this action mimics a disaster scenario to the secondary side.

In a replicated data cluster, the service group fails over to hostc or hostd if all devices were originally in the PAIR state, that is, no synchronization was in progress at the time of disaster.

In a global cluster, the administrator is notified of the failure. The administrator can then initiate the failover by declaring an outage.

Failback Test

Reconnect the replication link and reboot the original P-VOL hosts. You must manually resynchronize the devices using the same steps as the host failure test above.

Once the resynchronization is complete, migrate the application back to the original primary side:

```
# hagrp -online aqlagrp -sys hosta
```

The devices swap roles again and the environment state will be the same as when the test began.

Removing the Agent

Type the following command on each system to remove the agent. Answer prompts accordingly:

```
# swremove VRTSvcstc
# swremove VRTScstcw
# swremove VRTScsfdw
```

Setting Up a Fire Drill 5

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. Before setting up a fire drill for an application service group, make sure you have added the HTC resource to the service group.

A fire drill is performed at the secondary site. The initial steps involve configuring a fire drill service group, which is identical to the application service group, but uses HTCSnap, a fire drill resource, in place of the HTC resource. The fire drill service group uses a copy of the data used by the application service group.

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

VCS supports several fire drill configurations and provides the HTCSnap agent to manage the replication relationships during a fire drill. The Fire Drill Configuration wizard configures the fire drill service group.

Fire Drill Configurations

VCS supports the following fire drill configurations:

Fire Drill Configuration	Description
Gold Configuration	Runs the fire drill on a snapshot of the target array. Involves the following steps: ◆ Suspend replication to get a consistent snapshot. ◆ Take a snapshot of the target array on a ShadowImage device. ◆ Modify the disk name and the disk group name in the snapshot. ◆ Bring the fire drill service group online using the snapshotted data. **Note** For Gold configurations, you must use VERITAS Volume Manager to import and deport the storage.
Silver Configuration	Runs the fire drill on the target array after taking a snapshot. Involves the following steps: ◆ Suspend replication to get a consistent snapshot. ◆ Take a snapshot of the target array on a ShadowImage device. ◆ Bring the fire drill service group online using the data on the target array. The Silver configuration can only be used with ShadowImage pairs created with the `-m noread` flag to the `paircreate` command.
Bronze Configuration	Runs the fire drill on the target array. No snapshots are taken. Involves the following steps: ◆ Suspend replication to get a consistent snapshot. ◆ Bring the fire drill service group online using the data on the target array.

HTCSnap Agent

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

Agent Operations

The agent performs different functions depending on the fire drill configuration.

Operation (Entry Point)	Description
online	◆ **Gold** Suspends replication between the source and the target arrays, takes a local snapshot of the target LUN, resumes the replication between the arrays, and takes the fire drill service group online by mounting the snapshot. ◆ **Silver** Suspends replication between the source and the target arrays, takes a local snapshot of the target LUN, and takes the fire drill service group online by mounting the target LUN. ◆ **Bronze** Suspends replication between the source and the target arrays and takes the fire drill service group online using the target array. The operation also creates a lock file to indicate that the resource is online.
offline	◆ **Gold** Destroys the snapshot by synchronizing data between the target array and the device on which snapshot was taken. ◆ **Silver** Resumes replication between the source and the target arrays. Once the data is synchronized between the two arrays, the snapshot of the target array is destroyed by synchronizing data between the target array and the device where snapshot was taken. ◆ **Bronze** Resumes the replication between the source and the target arrays. The operation also removes the lock file created by the online operation.
monitor	Verifies the existence of the lock file to make sure the resource is online.

Operation (Entry Point)	Description
clean	Restores the state of the LUNs to their original state after a failed online operation.
action	For internal use.

Resource Type Definition

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

Attribute Definitions

Required Attributes	Type-Dimension	Description
TargetResName	string-scalar	For HTC - Name of the resource managing the LUNs to be snapshotted. The target resource is of type HTC if the data being snapshot is replicated; the resource is of type DiskGroup if the data is not replicated. For example, some applications like Oracle have data files and redo logs replicated, but temporary tablespace not replicated. The temporary tablespace must still exist at the DR site and may be part of its own disk group and is snapshotted independently.
ShadowInstance	integer-scalar	The instance number of the ShadowInstance P-VOL group. **Note** The P-VOL group must include the same LUNs as either the TrueCopy S-VOL group (if snapshotting replicated data) or the same LUNs as in the VxVM disk group (if snapshotting non-replicated data).
UseSnapshot	integer-scalar	Specifies whether the HTCSnap resource takes a local snapshot of the target array. Set this attribute to 1 for Gold and Silver configurations. For Bronze, set this attribute to 0. See “Configuring the Snapshot Attributes” on page 30.
RequireSnapshot	integer-scalar	Specifies whether the HTCSnap resource must take a snapshot before coming online. Set this attribute to 1 if you want the resource to come online only after it succeeds in taking a snapshot. Set this attribute to 0 if you do want the resource to come online even if it fails to take a snapshot. Setting this attribute to 0 creates the Bronze configuration. **Note** Set this attribute to 1 only if UseSnapshot is set to 1.

Required Attributes	Type-Dimension	Description
MountSnapshot	integer-scalar	Specifies whether the resource uses the snapshot to bring the service group online. Set this attribute to 1 for Gold configuration. For Silver and Bronze configurations, set the attribute to 0. **Note** Set this attribute to 1 only if UseSnapshot is set to 1.

Internal Attributes	Type-Dimension	Description
Responsibility	temporary string	For internal use only. Used by the agent to keep track of resynchonizing snapshots.
FDFile	temporary string	For internal use only. Used by the agent to locate the latest fire drill report.

Configuring the Snapshot Attributes

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

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

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 HTC resource is replaced by the fire drill resource HTCSnap.

The following configuration creates a Gold fire drill configuration, but allows VCS to run a Bronze fire drill if the snapshot does not complete successfully.

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

Configuring the Fire Drill Service Group

This section describes how to configure a fire drill service group using the Fire Drill Configuration wizard. Note that you can also use the text-based wizard, available at `/opt/VRTSvcs/bin/fdsetup-htc.`

Prerequisites

- ✔ Make sure the application service group is configured with an HTC resource.
- ✔ Make sure that the infrastructure to take snapshots is properly configured between the source and target arrays. This involves creating the Shadow Image pairs.
- ✔ When using Gold or Silver configuration, make sure you have ShadowImage for HTC installed and configured at the target array.
- ✔ For the Gold configuration, you must use VERITAS Volume Manager to import and deport the storage.
- ✔ The Silver configuration can only be used with ShadowImage pairs 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, since this is the only 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
  - ◆ You must use the Gold configuration without the option to run in the Bronze mode. This means the RequireSnapshot attribute must be set to 1.

Configuration Instructions

1. Set the *DISPLAY* variable and start the Fire Drill Configuration wizard as `root`.

   ```
   # hawizard firedrill
   ```

2. Read the information on the Welcome screen and click **Next**.

3. In the Wizard Options dialog box, select the application service group for which a fire drill service group is being configured.

   Note The wizard does not display service groups that do not have HTC resources.

4. In the Device Group Details dialog box, the wizard discovers and presents the device group from the application service group for which the fire drill service group is being configured. Verify the information and click **Next**.

5. In the Snapshot Methods dialog box, choose the configuration option for the fire drill service group.

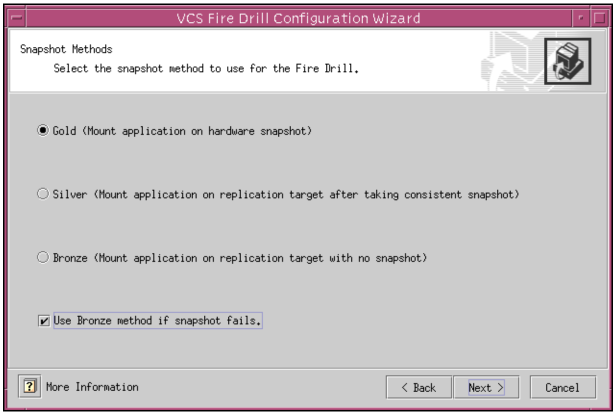

Snapshot Methods

Select the snapshot method to use for the Fire Drill.

● Gold (Mount application on hardware snapshot)

○ Silver (Mount application on replication target after taking consistent snapshot)

○ Bronze (Mount application on replication target with no snapshot)

☑ Use Bronze method if snapshot fails.

   a. Choose either a **Gold**, **Silver**, or **Bronze** configuration option. See "Fire Drill Configurations" on page 26 for more information.

   b. Select the **Use Bronze method if snapshot fails** check box if you want the fire drill service group to come online even if the resource fails to take a snapshot. This check box is enabled only if you choose the Gold or Silver configuration.

   c. Click **Next**.

6. Specify the ShadowImage instance.

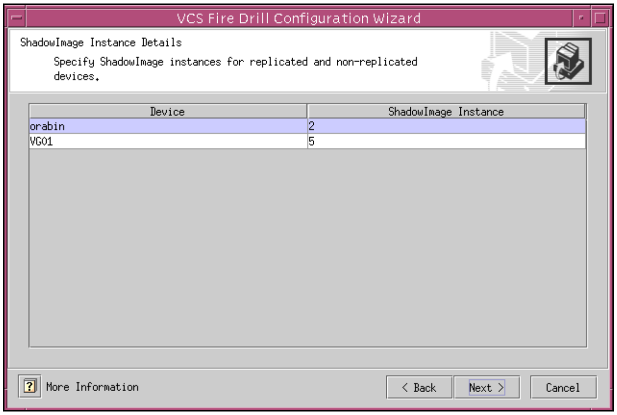

ShadowImage Instance Details

Specify ShadowImage instances for replicated and non-replicated devices.

Device	ShadowImage Instance
orabin	2
VG01	5

7. In the Snapshot Details dialog box, the wizard informs whether the device group on the target array has synchronized ShadowImage devices to take a snapshot. If the devices are synchronized, click **Next**.

   If the devices are not synchronized, click **Back** and verify whether you specified the correct ShadowImage instance.

   If the ShadowImage instance is correct, it is possible that data between the target array and the ShadowImage device, where the snapshot will be taken, is not synchronized. Quit the wizard, synchronize the data, and rerun the wizard.

8. In the Service Group Summary dialog box, review the service group configuration and change the resource names if desired.

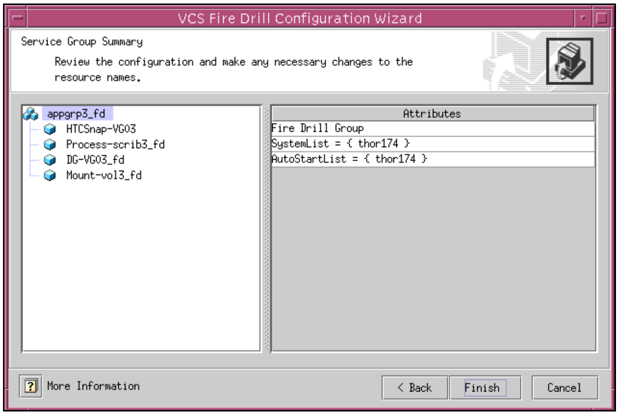

VCS Fire Drill Configuration Wizard

Service Group Summary

Review the configuration and make any necessary changes to the resource names.

```
appgrp3_fd
    HTCSnap-VG03
    Process-scrib3_fd
    DG-VG03_fd
    Mount-vol3_fd
```

Attributes
Fire Drill Group
SystemList = { thor174 }
AutoStartList = { thor174 }

   a. To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.

   b. Click **Finish**.

   The wizard starts running commands to create the fire drill service group. Various messages indicate the status of these commands.

9. In the Completing the Fire Drill Configuration Wizard dialog box, select the **check box** to bring the service group online on the local system.

10. Click **Close**.

Verifying a Successful Fire Drill

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 that your disaster recovery solution is configured correctly and the production service group will fail over to the secondary site in the event of an actual failure (disaster) at the primary site.

If the fire drill service group does not come online, review the VCS Engine log to troubleshoot the issues so that corrective action can be taken as necessary in the production service group. You can also view the fire drill log, located at `/tmp/fd-`*`servicegroup`*.

Caution Remember to take the fire drill offline once 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 were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

Index