Cluster Server Agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access Installation and Configuration Guide

AIX, Linux, Solaris

8.0.2

VERITAS™

Cluster Server Agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access Installation and Configuration Guide

Last updated: 2023-06-30

Legal Notice

Copyright © 2023 Veritas Technologies LLC. All rights reserved.

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

This product may contain third-party software for which Veritas is required to provide attribution to the third party (“Third Party Programs”). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

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

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

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

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

Veritas Technologies LLC

2625 Augustine Drive
Santa Clara, CA 95054

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

infoscaledocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1 Introducing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 7

About the agent for Hitachi TrueCopy / HUR / HP-XP Continuous Access 7
Supported software 8
Supported hardware 8
Typical Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access setup in a VCS cluster 9
Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent functions 10
About the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent's online function 13

Chapter 2 Installing and removing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 15

Before you install the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 15
Installing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 16
Installing the agent IPS package on Oracle Solaris 11 systems 17
Upgrading the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 17
Removing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 19
Configuring LVM on AIX 19
Configuring LVM on HP-UX 20

Chapter 3 Configuring the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 21

Configuration concepts for the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent 21
Resource type definition for the Hitachi TrueCopy agent 21
Attribute definitions for the TrueCopy agent 22
Sample configuration for the TrueCopy agent 33
Before you configure the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 35
About operations on volumes in a CVM environment 36
About cluster heartbeats 37
About configuring system zones in replicated data clusters 37
About preventing split-brain 38
Configuring the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 39
Performing a manual Volume Manager rescan 39
Configuring the agent manually in a global cluster 39
Configuring the agent manually in a replicated data cluster 40
Configuring the agent to compute RPO 41
Considerations for configuring TrueCopy agent in SFRAC or SFCFS environments 42

Chapter 4 Managing and testing clustering support for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 44

How recovers from various disasters in an HA/DR setup with Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access 45
Failure scenarios in global clusters 45
Failure scenarios in replicated data clusters 50
Replication link / Application failure scenarios 54
Testing the global service group migration 54
Testing disaster recovery after host failure 56
Testing disaster recovery after site failure 57
Performing failback after a node failure or an application failure 58
Performing failback after a site failure 59

Chapter 5 Setting up a fire drill 61

About fire drills 61
Fire drill configurations 62

Note on the Gold configuration 63
About the HTCSnap agent 63
HTCSnap agent functions 63
Resource type definition for the HTCSnap agent 64
Attribute definitions for the HTCSnap agent 65
About the Snapshot attributes 66
Before you configure the fire drill service group 66
Configuring the fire drill service group 67
Creating the fire drill service group using Cluster Manager (Java Console) 68
Creating the fire drill service group using the Fire Drill SetUp Wizard 70
Verifying a successful fire drill 71
Sample configuration for a fire drill service group 71

Chapter 1

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

This chapter includes the following topics:

- About the agent for Hitachi TrueCopy / HUR / HP-XP Continuous Access
- Supported software
- Supported hardware
- Typical Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access setup in a VCS cluster
- Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent functions

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

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

The agent monitors and manages the state of replicated Hitachi TrueCopy devices that are attached to VCS nodes. The agent ensures that the system that has the

TrueCopy resource online also has safe and exclusive access to the configured devices.

You can use the agent in replicated data clusters and 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
Hitachi Lightning	data, never, and async
Hitachi Thunder	data and never

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

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

Supported software

For information on the software versions that the (VCS) agent for TrueCopy supports, see the Veritas Services and Operations Readiness Tools (SORT) site:

https://sort.veritas.com/agents

Supported hardware

The agent for Hitachi TrueCopy provides support for the following:

- The agent supports Hitachi TrueCopy replication, provided that the host, HBA, array combination are in Hitachi's hardware compatibility list.
- The agent for Hitachi TrueCopy does not support other Hewlett-Packard replication solutions under the Continuous Access umbrella such as Continuous Access Storage Appliance (CASA). The agent only supports Continuous Access XP.

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

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

The following figure displays a typical cluster setup in a TrueCopy environment.

Figure 1-1 Typical clustering setup for the agent

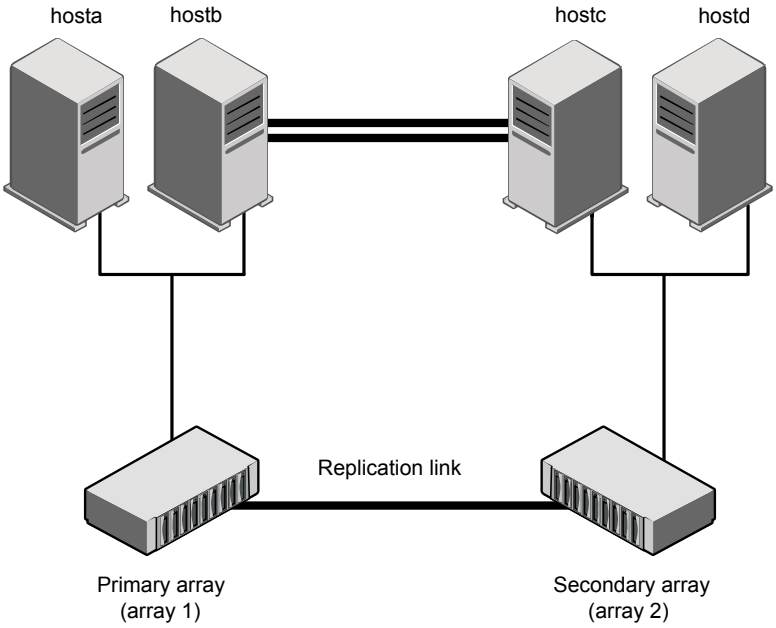

Clustering in a TrueCopy environment typically consists of the following hardware infrastructure:

- The primary array (array1) has one or more P-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to the Hitachi TrueCopy array that contains the TrueCopy P-VOL devices.
- The secondary array (array2) has one or more S-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to a Hitachi TrueCopy array that contains the TrueCopy S-VOL devices. The S-VOL devices are paired with the P-VOL devices in the P-VOL array. The S-VOL hosts and arrays must be at a significant distance to survive a disaster that may occur at the P-VOL side.
- Network heartbeating between the two data centers to determine their health; this network heartbeating could be LLT or TCP/IP.

- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT.
  In a global cluster environment, you must attach all hosts in a cluster to the same Hitachi TrueCopy array.
- In parallel applications like Storage Foundation for Oracle RAC, all hosts that are attached to the same array must be part of the same GAB membership. Storage Foundation for Oracle RAC is supported with TrueCopy only in a global cluster environment and not in a replicated data cluster environment.

Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent functions

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

The agent performs the following functions:

Table 1-2 Agent Functions

Function	Description
online	If the state of all local devices is read-write enabled, the agent makes the devices writable by creating a lock file on the local host. If one or more devices are not in a writable state, the agent runs the `horctakeover` command to enable read-write access to the devices. For S-VOL devices in any state other than SSWS/SSUS/SMPL, the agent runs the `horctakeover` command and makes the devices writable. The time required for failover depends on the following conditions: ■ The health of the original primary. ■ The RAID Manager timeouts as defined in the horcm configuration file for the device group. The agent considers P-VOL devices writable and takes no action other than going online, regardless of their status. If the S-VOL devices are in the COPY state, the agent runs the `horctakeover` command after one of the following: ■ The synchronization from the primary completes. ■ The OnlineTimeout period of the entry point expires, in which case the resource faults. If the S-VOL devices are in the PAIR state, the agent runs the `pairdisplay` command without the `-l` option to retrieve the state of the remote site devices. If it finds that the P-VOL devices are in PAIR state, the agent proceeds with the failover. But if the remote RAID manager is down, then the agent honors the SplitTakeover attribute configuration before performing failover. See "About the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent's online function" on page 13. If the RAID manager is not up for an instance, the agent runs the `horcmstart` command to bring the RAID manager online.
offline	The agent removes the lock file that was created for the resource by the online entry point. The agent does not run any TrueCopy commands because taking the resource offline is not indicative of an intention to give up the devices.

Table 1-2 Agent Functions *(continued)*

Function	Description
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. Based on other attribute values, the monitor entry point examines the state of the devices or the state of the replication link between the arrays.
open	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. Note that the agent does not remove the lock file if the agent starts after the following command: `hastop -force`
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 it was killed. If a management operation was in progress, it could potentially leave the devices in an unusable state.
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	The agent supports the following actions using the `hares -action` command from the command line: ■ pairdisplay—Displays information about all devices. ■ pairresync—Resynchronizes the S-VOL devices from the VCS command line after connectivity failures are detected and corrected. ■ pairresync-swaps—Promotes the S-VOLs to P-VOLs and resynchronizes the original P-VOLs. ■ localtakeover—Makes the local devices write-enabled.

Table 1-2 Agent Functions *(continued)*

Function	Description
action\PreSwitch	Ensures that the remote site cluster can come online during a planned failover within a GCO configuration without data loss. The VCS engine on the remote cluster invokes the PreSwitch action on all the resources of the remote Service Group during a planned failover using the `hagrp -switch` command. For this, the PreSwitch attribute must be set to 1. The option `-nopre` indicates that the VCS engine must switch the servicegroup regardless of the value of the PreSwitch service group attribute. If running the PreSwitch action fails, the failover should not occur. This minimizes the application downtime and data loss . For more information on the PreSwitch action and the PreSwitch feature in the VCS engine, refer to the *Cluster Server Administrator's Guide*.
action\vxdiske	Reports the mapping between the physical disk name and the volume manager disk name for all connected disks.
action\GetCurrentRPO	Fetches the current point in time RPO. The agent performs this action function on the disaster recovery (DR) system where the ComputeDRSLA attribute is set to 1. The RPO is computed in seconds. **Note:** The agent does not compute the RPO when the group is frozen. The agent does not store the computed RPO; make a note of the RPO for future reference.

Note: The agent uses the following internal action functions to compute the RPO: StartRPOComputation, StopRPOComputation, StartWriter, and ReportRPOData.

About the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent's online function

If the state of all local devices is read-write enabled, the agent makes the devices writable by creating a lock file on the local host. The agent considers the P-VOL devices writable and takes no action other than going online, regardless of their status.

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

If one or more devices are not in a writable state, the agent runs the `horctakeover` command to enable read-write access to the devices. If the `horctakeover` command exits with an error (exit code > 5), for example, due to a timeout, then the agent flushes and freezes the group to indicate that user-intervention is required to identify the cause of the error.

For S-VOL devices in any state other than SSWS and SSUS, the agent honors the SplitTakeover attribute and runs the `horctakeover` command to make the devices writable.

The time required for failover depends on the following conditions:

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

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

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

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

- The synchronization from the primary completes.
- When the OnlineTimeout period of the entry point expires, the `horctakeover` command will not be executed, in which case the resource faults.

If S-VOL devices are in the PAIR state, the agent issues the `pairdisplay` command without the `-l` option to get the replication link state. If it finds that the P-VOL devices are in the PAIR state, the agent proceeds with failover. If remote horcm is down, the SplitTakeover attribute is honored before issuing the `horctakeover` command. The agent validates the value of OnlineTimeout for the HTC type is sufficient to run the `horctakeover` command. If the agent finds this value of OnlineTimeout is insufficient, the agent logs an appropriate error message.

Chapter 2

Installing and removing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access

This chapter includes the following topics:

- Before you install the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access
- Installing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access
- Upgrading the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access
- Removing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access

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

Before you install the VCS agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access, ensure that you install and configure VCS on all nodes in the cluster.

Set up replication and the required hardware infrastructure.

For information about setting up Oracle RAC environment, refer to the *Storage Foundation for Oracle RAC Configuration and Upgrade Guide*.

See "Typical Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access setup in a VCS cluster" on page 9.

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

You must install the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster.

These instructions assume that you have already installed VCS or SF for Oracle RAC or Storage Foundation Cluster File System (SFCFS).

To install the agent in a VCS environment

1. Download the Agent Pack from the Veritas Services and Operations Readiness Tools (SORT) site: https://sort.veritas.com/agents.
2. Log in as a superuser.
3. Install the package.

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

   Linux
   ```
   # rpm -ihv
   VRTSvcstc-AgentVersion-Linux_GENERIC.noarch.rpm
   ```

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

   For Solaris 11 systems, refer to the agent IPS package installation procedure.

   See "Installing the agent IPS package on Oracle Solaris 11 systems" on page 17.

Note: On successful installation of the agent, if VCS is running, the agent types definition is automatically added to the VCS configuration.

Installing the agent IPS package on Oracle Solaris 11 systems

To install the agent IPS package on an Oracle Solaris 11 system

1. Copy the `VRTSvcstc.p5p` package from the `pkgs` directory to the system in the `/tmp/install` directory.
2. Disable the publishers that are not reachable as package install may fail, if any of the already added repositories are unreachable.

   ```
   # pkg set-publisher --disable <publisher name>
   ```

   where the publisher name is obtained using the `pkg publisher` command.
3. Add a file-based repository in the system.

   ```
   # pkg set-publisher -g /tmp/install/VRTSvcstc.p5p Veritas
   ```
4. Install the package.

   ```
   # pkg install --accept VRTSvcstc
   ```
5. Remove the publisher from the system.

   ```
   # pkg unset-publisher Veritas
   ```
6. Enable the publishers that were disabled earlier.

   ```
   # pkg set-publisher --enable <publisher name>
   ```

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

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

To upgrade the agent software

1. Save the VCS configuration.

   ```
   # haconf -dump -makero
   ```
2. Stop the agent if it is running.
   - If the HTC agent is configured and running:

     ```
     # haagent -stop HTC -force -sys system
     ```
   - If the HTCSnap agent is configured and running:

     ```
     # haagent -stop HTCSnap -force -sys system
     ```

3 Verify the status of the agent and ensure that it is not running.

- For HTC agent:

  ```
  # haagent -display HTC | grep Running
  ```

  The command output resembles the following:

  ```
  HTC Running No
  ```

- For HTCSnap agent:

  ```
  # haagent -display HTCSnap | grep Running
  ```

  The command output resembles the following:

  ```
  HTCSnap Running No
  ```

4 Upgrade the agent.

- For Linux, run the following command:

  ```
  # rpm -Uvh VRTSvcstc-AgentVersion-Linux_GENERIC.noarch.rpm
  ```

- For AIX and Solaris, perform the following steps:
  - Remove the previous version of the agent from the node.
    See "Removing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access" on page 19.
  - Install the latest version of the agent.
    See "Installing the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access" on page 16.
  - If the agent types file was not added automatically on successful installation of the agent, add the agent types file.

    ```
    # /etc/VRTSvcs/conf/sample_htc/addHTCType.sh 1
    ```

5 Start the agent.

- For HTC agent:

  ```
  # haagent -start HTC -sys system
  ```

- For HTCSnap agent:

```
# haagent -start HTCSnap -sys system
```

6 Verify the status of the agent and ensure that it is running.

- For HTC agent:

  ```
  # haagent -display HTC | grep Running
  ```

  The command output resembles the following:

  ```
  HTC Running Yes
  ```

- For HTCSnap agent:

  ```
  # haagent -display HTCSnap | grep Running
  ```

  The command output resembles the following:

  ```
  HTCSnap Running Yes
  ```

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

Before you attempt to remove the agent, make sure the application service group is not online.

You must remove the TrueCopy agent from each node in the cluster.

To uninstall the agent in a VCS environment

- To remove the agent, type the following command on each node. Answer prompts accordingly:

AIX	`# installp -u VRTSvcstc.rte`
Linux	`# rpm -e VRTSvcstc`
Solaris	`# pkgrm VRTSvcstc`

Configuring LVM on AIX

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

Configuring LVM on HP-UX

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

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

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

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

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

To configure LVM on HP-UX

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

Chapter 3

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

This chapter includes the following topics:

- Configuration concepts for the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent
- Before you configure the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access
- Configuring the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access

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

Review the resource type definition and attribute definitions for the agent.

Resource type definition for the Hitachi TrueCopy agent

The resource type definition defines the agent in VCS.

```
type HTC (
        static keylist RegList = { ComputeDRSLA, SplitTakeover,
        LinkMonitor, RoleMonitor, FreezeSecondaryOnSplit,
        AllowOnlineOnSimplex }
        static keylist SupportedActions = { localtakeover,
        pairresync, pairresync-swaps, pairdisplay, vxdiske,
        PreSwitch, ReportRPOData, StartWriter, GetCurrentRPO,
        StartRPOComputation, StopRPOComputation }
        static int InfoInterval = 300
        static keylist LogDbg = { DBG_1, DBG_2, DBG_3 }
        static int OpenTimeout = 180
        static str ArgList[] = { BaseDir, GroupName, Instance,
        SplitTakeover, LinkMonitor, RoleMonitor, FreezeSecondaryOnSplit,
        AllowOnlineOnSimplex, ComputeDRSLA, AdvancedOpts }
        str BaseDir = "\"/HORCM/usr/bin\""
        str GroupName
        int Instance
        int SplitTakeover
        int LinkMonitor
        int RoleMonitor
        int FreezeSecondaryOnSplit
        boolean AllowOnlineOnSimplex = 0
        temp str VCSResLock
        temp str TargetFrozen
        int ComputeDRSLA
        temp boolean Tagging = 0
        str AdvancedOpts{} = { AllowAutoFailoverInterval="-1" }
        temp str PVOLStateTime
)
```

Attribute definitions for the TrueCopy agent

Table 3-1 lists the attributes associated with the agent:

Table 3-1 Attributes for the Hitachi TrueCopy agent

Attribute	Description
BaseDir	Path to the RAID Manager Command Line Interface. Type-Dimension: string-scalar Default: /HORCM/usr/bin

Table 3-1 Attributes for the Hitachi TrueCopy agent *(continued)*

Attribute	Description
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: integer
SplitTakeover	A flag that determines the following: ■ Whether the agent permits a failover to S-VOL devices if the replication link is disconnected (that is, when P-VOL devices are in the PSUE state) ■ Whether the agent cannot connect to the remote site RAID manager ■ Whether the replication link is manually suspended (that is when P-VOL devices are in the PSUS state) See "About the SplitTakeover attribute for the Hitachi TrueCopy agent" on page 27. Type-Dimension: integer-scalar Default: 0

Table 3-1 Attributes for the Hitachi TrueCopy agent *(continued)*

Attribute	Description
LinkMonitor	An integer that determines the action the agent takes when the replication link is disconnected. Depending on the value of this attribute, the agent takes the following action: ■ The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the `pairresync` command. ■ The value 2 indicates that when the replication link is disconnected, the agent generates SNMP traps or email alerts. If the status of the configured HTC device changes to PSUE, the agent generates an SNMP trap of severity Error or an email alert indicating that the resource health has gone down. For all other types of status changes of the configured HTC devices, the agent generates an SNMP trap of severity Information indicating that the resource health has improved. For information about the VCS severity levels, refer to the *Cluster Server Administrator's Guide*. The agent logs a message in the VCS engine log: `The state of P-VOL/S-VOL devices in device group device group name has changed from previous state to current state.` Type-Dimension: integer-scalar Default: 0 For information about the NotifierMngr agent that starts, stops, and monitors a notifier process, refer to the *Cluster Server Bundled Agents Reference Guide*. The notifier process manages the reception of messages from VCS and the delivery of those messages to SNMP consoles and SMTP servers.

Table 3-1 Attributes for the Hitachi TrueCopy agent *(continued)*

Attribute	Description
RoleMonitor	Determines if the agent must perform detailed monitoring of HTC volumes. If this attribute is set to 0, the agent does not perform detailed monitoring. This attribute is disabled by default. If this attribute is set to 1, the agent monitors the status of the HTC volumes everytime a monitor cycle runs. In addition, the HTC resource comes online only when any of the following conditions are met: ■ When the volume is P-VOL ■ When the volume is S-VOL and the status is SSWS ■ When the volume is S-VOL, the status is SSUS, and the M flag of the corresponding P-VOL is set to W. Type-Dimension: integer-scalar Default: 0
FreezeSecondaryOnSplit	A flag that determines if the agent must freeze the service group in the remote cluster when the TrueCopy replication link is either split or suspended. The value 1 indicates that the agent must freeze the service group in the remote cluster when the replication link is split or suspended. Type-Dimension: integer-scalar Default: 0
AllowOnlineOnSimplex	A flag that determines if the agent must allow a resource to come online when the TrueCopy devices are in SMPL (Simplex) state. This attribute is honored only when the agent attempts to bring a resource online. The value false indicates that the agent must not allow a resource to come online when TrueCopy devices are in SMPL (Simplex) state. Type-Dimension: boolean-scalar Default: false
TargetFrozen	For internal use. Do not modify.
VCSResLock	The agent uses the VCSResLock attribute to guarantee serialized management in case of a parallel application. Type-Dimension: temporary string

Table 3-1 Attributes for the Hitachi TrueCopy agent *(continued)*

Attribute	Description
ComputeDRSLA	Used to enable or disable Recovery Point Objective (RPO) computation. Set this attribute on any one node in the disaster recovery (DR) cluster. Setting this attribute to 1 starts the RPO computation process. Ensure that you reset this attribute to 0 after you use the GetCurrentRPO action function to check the RPO. Type-Dimension: integer-scalar Default: 0
Tagging	This internal attribute is used for maintaining the process of computing RPO.

AdvancedOpts	Used at the time of monitoring. This attribute enables the agent to execute a custom script during the monitor cycle of the resource.
	
	Use the AllowAutoFailoverInterval attribute with this attribute. The agent automatically fails over if certain conditions are met and AllowAutoFailoverInterval is set to 0 (zero) or a positive integer.
	
	To disable the execution of the custom script, set AllowAutoFailoverInterval to `-1` or remove it from the AdvancedOpts attribute. For example:
	
	`AdvancedOpts{} =`
	`{ AllowAutoFailoverInterval="-1" }`
	
	Type-Dimension: string-association

AllowAutoFailoverInterval	The agent uses this attribute to fail over automatically only if all the following conditions are met: ■ This attribute is set to 0 (zero) or a positive integer. ■ The fence level is NEVER. ■ The remote RAID manager is not reachable. The failover takes place only if the value of this attribute is greater than the last registered PAIR state time difference. **Note:** This attribute is applicable only when the fence level is NEVER. See "Special consideration for fence level NEVER" on page 29. Default: `-1` Type-Dimension: string-association

Table 3-1 Attributes for the Hitachi TrueCopy agent *(continued)*

Attribute	Description
PVOLStateTime	This is an internal attribute that is used to maintain the P-VOL state and the timestamp when the instance was last registered as PAIR. **Note:** Do not modify this attribute.

About the SplitTakeover attribute for the Hitachi TrueCopy agent

The SplitTakeover attribute determines the following:

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

SplitTakeover attribute = 0

The default value of the SplitTakeover attribute is 0.

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

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

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

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

SplitTakeover attribute = 1

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

About the FreezeSecondaryOnSplit attribute for the Hitachi TrueCopy agent

In a global cluster environment, if the agent at the P-VOL side detects the PSUE or PSUS state locally and FreezeSecondaryOnSplit is set to 1, then the agent freezes the service group at the S-VOL side to prevent a failover. The agent unfreezes the service group after the link is restored and the devices are resynchronized.

FreezeSecondaryOnSplit attribute = 0

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

About the HTC configuration parameters

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

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

Update the HORCM_INST section of the configuration file.

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

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

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

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

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

- The ActionTimeout value of HTC type should be more than twice the value of remote RAID manager timeout.

Special consideration for fence level NEVER

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

Consider the following failure scenario:

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

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

In this scenario, the agent allows a failover to happen only if `AllowAutoFailoverInterval < (Event B - Event A)`, where:

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

The AllowAutoFailoverInterval value is passed to the AdvancedOpts attribute.

Table 3-2 Failover scenarios for the various AllowAutoFailoverInterval values

Value	Conditions	Actions
0	■ The fence level is NEVER. ■ The remote RAID manager is not reachable.	The failover is triggered. The last known PAIR status of P-VOL is recorded but not used.
Greater than 0	■ The fence level is NEVER. ■ The remote RAID manager is not reachable. ■ The last known remote state within `T` seconds is PAIR, where `T` is the AllowAutofailoverInterval value.	The failover is triggered. The last known PAIR status of P-VOL is recorded and is used to determine the failover action.

Table 3-2 Failover scenarios for the various AllowAutoFailoverInterval values *(continued)*

Value	Conditions	Actions
Greater than 0	■ The fence level is NEVER. ■ The remote RAID manager is not reachable. ■ The last known remote state is PAIR and is greater than `T` seconds, where `T` is the AllowAutofailoverInterval value.	The failover is not triggered. The last known PAIR status of P-VOL is recorded and is used to determine the failover action.
Greater than 0	■ The fence level is NEVER. ■ The remote RAID manager is not reachable. ■ The last known remote state is not PAIR.	The failover is not triggered. The last known PAIR status of P-VOL is recorded and is used to determine the failover action.
`-1` or The value is not passed to the AdvancedOpts attribute.	Any	The automatic failover is not enabled. The last known PAIR status of P-VOL is not recorded. Manual intervention is required to restore the operations in this scenario.

Consider the following before using the AllowAutoFailoverInterval attribute:

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

See "Considerations for calculating the AllowAutoFailoverInterval attribute value" on page 31.

Considerations for calculating the AllowAutoFailoverInterval attribute value

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

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

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

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

Variable	Source	Default value	Usage
A = MonitorInterval attribute value	HTC agent	60 seconds	Specifies how frequently the agent polls and records the PAIR status for P-VOL.
B = AYAInterval attribute value	Heartbeat agent	60 seconds	The interval between two heartbeats in the global cluster. You can modify this value using the `hahb` command, for example: `hahb -modify ICMP AYAInterval 45`

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

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

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

```
A + (B x C) + D + BufferTime
= 60 + (60 x 3) + 120 + 40
= 400 seconds
```

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

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

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

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

Sample configuration for the TrueCopy agent

Figure 3-1 shows a dependency graph of a VCS service group that has a resource of type HTC.

Figure 3-1 VCS service group with resource type HTC

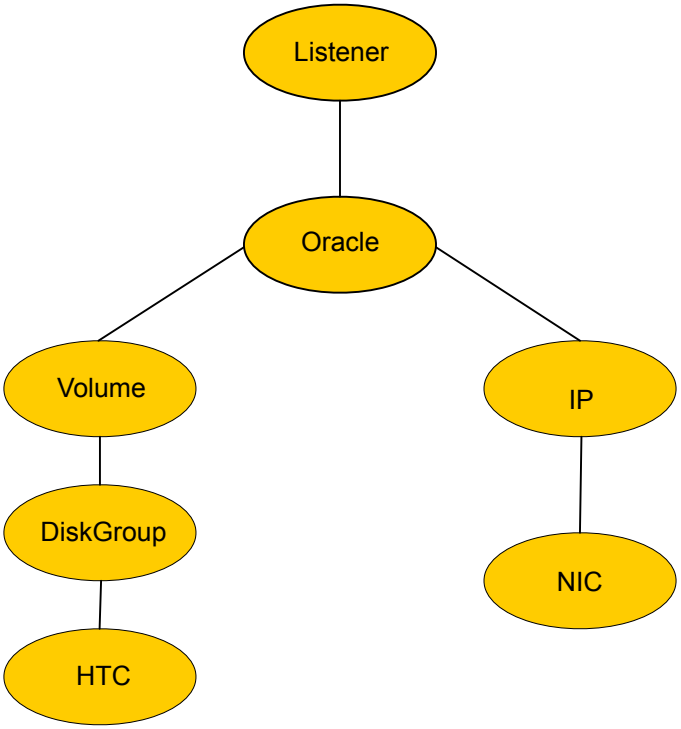

You can configure a resource of type HTC in the main.cf file as:

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

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

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

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

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

        HTC rep_htc (
                GroupName = vg1
                Instance = 1
                )

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

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

```
CFSfsckd vxfsckd (
        )

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

CVMVxconfigd cvm_vxconfigd (
        Critical = 0
        CVMVxconfigdArgs = { syslog }
        )

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus
```

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

Before you configure the agent, review the following information:

- Verify that you have installed the agent on all the cluster nodes.
- Verify the hardware setup for the agent.
  See "Typical Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access setup in a VCS cluster" on page 9.
- Make sure that the cluster has an effective heartbeat mechanism in place.
  See "About cluster heartbeats" on page 37.
  See "About preventing split-brain" on page 38.
- Set up system zones in replicated data clusters.
  See "About configuring system zones in replicated data clusters" on page 37.
- Verify that the clustering infrastructure is in place.
  - If you plan to configure the agent in a global cluster, make sure the global service group for the application is configured.
    For more information, refer to the *Cluster Server Administrator's Guide*.
  - If you want to configure the agent in an SF Oracle RAC environment, verify that the SF Oracle RAC global cluster infrastructure is in place.

- If you plan to configure the agent in a replicated data cluster, make sure the required replication infrastructure is in place and that the application is configured.

- Ensure that the HORC manager is configured to access the device groups.
  - Verify that the HTC instance is configured appropriately and is in a running state.
  - Verify that the HORC manager CLIs execute successfully. This is essential for the HTC and the HTCSnap agents to be able to fetch HTC-related data and to succesfully perform failover, switchover, and other operations.

About operations on volumes in a CVM environment

In a Cluster Volume Manager (CVM) environment, the HTC agent may import and deport the VxVM-managed hardware-replicated (HTC) disk groups that are defined for the corresponding CVMVolDg resources. If you do not want the HTC agent to control these operations, remove the SupportedActions that are defined for the CVMVolDg-related resources.

To remove the SupportedActions definition for all CVMVolDg-related resources

1. View the SupportedActions definition:

   ```
   # hatype -display CVMVolDg | grep -i SupportedActions
   ```

   Sample output:

   ```
   CVMVolDg SupportedActions import deport vxdctlenable
   ```

2. Make the VCS configuration writable:

   ```
   # haconf -makerw
   ```

3. Update the CVMVolDg configuration:

   ```
   # hatype -modify CVMVolDg SupportedActions ""
   ```

4. Verify that the SupportedActions definition has been removed:

   ```
   # hatype -display CVMVolDg | grep -i SupportedActions
   ```

   Sample output:

   ```
   CVMVolDg SupportedActions
   ```

5. Make the VCS configuration read-only:

   ```
   # haconf -dump -makero
   ```

About cluster heartbeats

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

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

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

Use the SystemZones attribute to create these zones.

Figure 3-2 Example system zone configuration

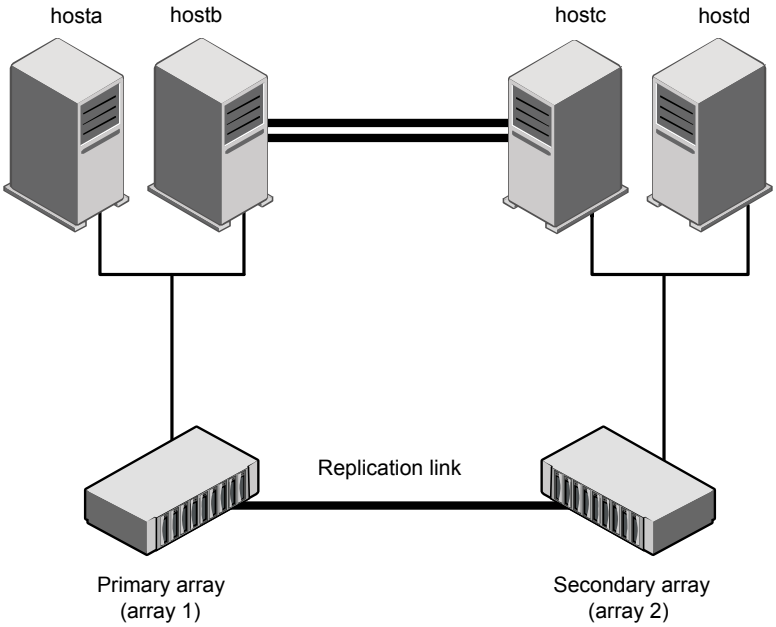

Global clusters do not require system zones because failover occurs on a remote cluster if all local targets have been exhausted.

About preventing split-brain

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

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

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

You can configure clustered application in a disaster recovery environment by:

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

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

See "Sample configuration for the TrueCopy agent" on page 33.

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

Performing a manual Volume Manager rescan

If you configure Volume Manager diskgroups on the disks that are replicated, the diskgroups do not come online the first time after failover on the secondary node. You must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the diskgroups online. This rescans all Volume Manager objects and must be performed only once after which the failover works uninterrupted.

To perform a manual Volume Manager rescan

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

Configuring the agent manually in a global cluster

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

To configure the agent in a global cluster

1. Start Cluster Manager (Java Console) and log on to the cluster.
2. If the agent resource type (HTC) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types**, and select:

   /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 service group, configure the service group using the Global Group Configuration Wizard.

   Refer to the *Cluster Server Administrator's Guide* for more information.
8. Change the ClusterFailOverPolicy attribute from the default, if necessary. Veritas recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.
9. Repeat step 5 through step 8 for each service group in each cluster that uses replicated data.
10. The configuration must be identical on all cluster nodes, both primary and disaster recovery.

Configuring the agent manually in a replicated data cluster

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

To configure the agent in a replicated data cluster

1. Start Cluster Manager and log on to the cluster.
2. If the agent resource type (HTC) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types** and select:

   /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 bottom of the service group.

Link the VMDg and HTC resources so that VMDg resources depend on Hitachi Truecopy.

6 Configure the attributes of the HTC resource.

7 Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array.

Configuring the agent to compute RPO

In a global cluster environment, the agent for Hitachi TrueCopy / HUR / HP-XP Continuous Access can compute the recovery point objective (RPO), which is a disaster recovery (DR) SLA. In a DR configuration where data is replicated asynchronously to the DR site, the DR site data is not always as current as the primary site data.

RPO is the maximum acceptable amount of data loss in case of a disaster at the primary site. The agent computes RPO in terms of time, that is, in seconds.

Before you configure the agent to compute the RPO, ensure that the following pre-requisites are met:

- The service group containing the HTC resource and the VxVM disk group resource are online at the production site.
- The disk group resource is dependent on the HTC resource.

To configure the agent to compute the RPO:

1 In the DR cluster, on any one of the nodes where devices are asynchronously replicated and where the service group is configured, run the following command to start the RPO computation:

```
hares -modify HTC_resource_name ComputeDRSLA 1 -sys system_name.
```

2 Run the following command on the same node in the DR cluster:

```
hares -action HTC_resource_name GetCurrentRPO -sys system_name
```

The action entry point displays the RPO. The agent does not store the computed RPO; make a note of the RPO for future reference.

If the RPO is not reported, it indicates that the agent needs more time to finish computing the RPO. Wait for some more time before you run the GetCurrentRPO action function again.

3 To stop RPO computation, run the following command:

```
hares -modify HTC_resource_name ComputeDRSLA 0 -sys system_name
```

Considerations for configuring TrueCopy agent in SFRAC or SFCFS environments

Consider the following attribute definitions and usage when configuring the TrueCopy agent in the SFRAC or SFCFS environments with Cluster Volume Manager (CVM).

CVMDeportOnOffline

The CVMVolDg agent uses the CVMDeportOnOffline attribute to determine whether or not to deport a shared disk group when the corresponding CVMVolDg resource is taken offline.

CVMDeportOnOffline attribute value	CVMVolDg agent behavior
0 (Default)	Does not deport the disk group when the CVMVolDg resource is taken offline.
1	Deports the disk group when the CVMVolDg resource is taken offline.

Note: You must set the CVMDeportOnOffline attribute to 1 for all the CVMVolDg resources that depend on the VCS hardware replicated managed devices, such as TrueCopy.

Run the following commands to set this attribute to 1:

1. `# haconf -makerw`
2. `# hares -modify cvmvoldg_res CVMDeportOnOffline 1`
3. `# haconf -dump -makero`

Run the following command to verify that the attribute value is set as expected:

```
# hares -display cvmvoldg_res | grep CVMDeportOnOffline
```

ClearClone

The TrueCopy agent uses the ClearClone attribute to update the on-disk UDID content for TrueCopy hardware-replicated devices at a disk group level.

Note: Do not use the ClearClone attribute with hardware clone devices like TrueCopy.

When the DiskGroup or CVMVolDg resources are defined with the ClearClone attribute set to 1, VCS calls the underlying VxVM command to import the disk group.

VxVM provides the VxDg import option (`-c`) to update the UDID-related content. When the `-c` option is used, the `udid_mismatch` and the subsequent `clone_disk` flags are cleared in a single operation from the disks in the specified disk group.

```
# grep -i "clear clone" /opt/VRTSvcs/bin/CVMVolDg/actions/import
# '-c' option to clear clone flag and import clone dg as standard dg.
VCSAG_LOG_MSG "I" "Importing $cvmvoldg_dgname DG with -c option
to clear clone flag on disk." 1111 "$cvmvoldg_dgname"
```

For shared disk groups, ensure that the `-c` option is specified in the CVMVolDg import actions script.

VxVM performs additional checks when using DMP to determine whether the device is a hardware replicated device, or a hardware clone. This additional safeguard is not available when using third-party drivers such as MPxIO, MPIO, and EMC PowerPath.

Note: MPIO and EMC PowerPath are not supported with TrueCopy when it is used in combination with VxVM or CVM and therefore with the VCS agent for TrueCopy.

The `/etc/VRTSvcs/conf/config/CVMTypes.cf` file contains the ClearClone definition. The ClearClone attribute of a CVMVolDg resource only takes an integer value. You can verify this as follows:

```
# grep -w ClearClone /etc/VRTSvcs/conf/config/CVMTypes.cf
static str ArgList[] = { CVMDiskGroup, CVMVolume, CVMActivation,
    CVMVolumeIoTest, CVMDGAction, CVMDeportOnOffline,
    CVMDeactivateOnOffline, State, ClearClone }
int ClearClone
```

Note: The `VRTScavf` package contains the CVMVolDg action scripts.

See “Sample configuration for the TrueCopy agent” on page 33.

Chapter 4

Managing and testing clustering support for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access

This chapter includes the following topics:

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

How recovers from various disasters in an HA/DR setup with Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access

This topic lists various failure scenarios and describes how responds to the failures in the following DR cluster configurations.

Global clusters

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

See "Failure scenarios in global clusters" on page 45.

Replicated data clusters

When service group faults or system faults occur, the failover behavior depends on the value of the AutoFailOver attribute of the faulted service group. The agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access ensures safe and exclusive access to the configured Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access devices.

See "Failure scenarios in replicated data clusters" on page 50.

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

Failure scenarios in global clusters

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

Table 4-1 Failure scenarios in a global cluster configuration with the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access

Failure	Description and response

Application failure	Application cannot start successfully on any hosts at the primary site.
	response at the secondary site:
	■ Causes global service group at the primary site to fault and displays an alert to indicate the fault.
	■ Does the following based on the ClusterFailOverPolicy global service group attribute:
	  ■ Auto or Connected— automatically brings the faulted global group online at the secondary site.
	  ■ Manual—No action. You must bring the global group online at the secondary site.
	Agent response:
	■ Write enables the devices at the secondary site, except when the link is manually suspended with the read-only option.
	■ Swaps the P-VOL/S-VOL role of each device in the device group.
	■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site.
	See "Performing failback after a node failure or an application failure" on page 58.
	See "Replication link / Application failure scenarios" on page 54.

Host failure	All hosts at the primary site fail.
	response at the secondary site:
	■ Displays an alert to indicate the primary cluster fault.
	■ Does the following based on the ClusterFailOverPolicy global service group attribute:
	  ■ Auto— automatically brings the faulted global group online at the secondary site.
	  ■ Manual or Connected—No action. You must bring the global group online at the secondary site.
	The agent does the following:
	■ Write enables the devices at the secondary site, except when the link is manually suspended with the read-only option.
	■ Swaps the P-VOL/S-VOL role of each device in the device group.
	■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site.
	See "Performing failback after a node failure or an application failure" on page 58.

Table 4-1 Failure scenarios in a global cluster configuration with the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access *(continued)*

Failure	Description and response

Site failure	All hosts and the storage at the primary site fail.
	
	response at the secondary site:
	
	■ Displays an alert to indicate the cluster fault.
	■ Does the following based on the ClusterFailOverPolicy global service group attribute:
	  ■ Auto— automatically brings the faulted global group online at the secondary site.
	  ■ Manual or Connected—No action. You must bring the global group online at the secondary site.
	
	Agent response: The agent does the following on the secondary site in case of a manual failover based on the value of the SplitTakeover attribute of the HTC resource:
	
	■ 1—The agent issues the `horctakeover` command to make the HTC devices write-enabled. The HTC devices go into the SSWS (Suspend for Swapping with S-VOL side only) state. If the original primary site is restored, you must execute the pairresync-swaps action on the secondary site to establish reverse replication.
	■ 0—Agent does not perform failover to the secondary site.
	
	See "Performing failback after a site failure" on page 59.

Table 4-1 Failure scenarios in a global cluster configuration with the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access *(continued)*

Failure	Description and response
Replication link failure	Replication link between the arrays at the two sites fails. The volume state on the primary site becomes PSUE. response: No action. Agent response: When the replication link is disconnected, the agent does the following based on the value of LinkMonitor attribute of the HTC resource: ■ 0—No action. ■ 1—The agent periodically attempts to resynchronize the S-VOL side using the `pairresync` command. The agent also logs a warning message to indicate that the replication link is broken. ■ 2—The agent periodically attempts to resynchronize the S-VOL side and also sends notifications about the disconnected link. Notifications are sent in the form of either SNMP traps or emails. For information about the NotifierMngr agent, refer to the *Symantec Cluster Server Bundled Agents Reference Guide*. If the value of the LinkMonitor attribute is not set to 1 or 2, you must manually resynchronize the HTC devices after the link is restored. To manually resynchronize the HTC devices after the link is restored: ■ Before you resync the S-VOL device, you must split off the Shadow Image device from the S-VOL device at the secondary site. ■ You must initiate resync of the S-VOL device using the agent's pairresync action. ■ After P-VOL and S-VOL devices are in sync, re-establish the mirror relationship between the Shadow Copy and the S-VOL devices. If you initiate a failover to the secondary site when resync is in progress, the online function of the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent waits for the resync to complete and then initiates a takeover of the S-VOL devices. **Note:** If you did not configure Shadow Copy devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Veritas recommends configuring Shadow Copy devices at both the sites. See "Replication link / Application failure scenarios" on page 54.

Table 4-1 Failure scenarios in a global cluster configuration with the agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access *(continued)*

Failure	Description and response

Network failure	The network connectivity and the replication link between the sites fail.
	response at the secondary site:
	■ at each site concludes that the remote cluster has faulted.
	■ Does the following based on the ClusterFailOverPolicy global service group attribute:
	  ■ Manual or Connected—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue.
	  ■ Auto— brings the global group online at the secondary site which may lead to a site-wide split brain. This causes data divergence between the devices on the primary and the secondary arrays.
	  When the network (WAC and replication) connectivity is restored, you must manually resync the data.
	  Note: Veritas recommends that the value of the ClusterFailOverPolicy attribute is set to Manual for all global groups to prevent unintended failovers due to transient network failures.
	To resynchronize the data after the network link is restored:
	■ Take the global service group offline at both the sites.
	■ Manually resynchronize the data.
	Use the `pairresync-swap` command to resynchronize from the secondary.
	■ Bring the global service group online on the secondary site.
	Agent response: Similar to the site failure.

Storage failure	The array at the primary site fails.
	response at the secondary site:
	■ Causes the global service group at the primary site to fault and displays an alert to indicate the fault.
	■ Does the following based on the ClusterFailOverPolicy global service group attribute:
	  ■ Auto or Connected— automatically brings the faulted global service group online at the secondary site.
	  ■ Manual—No action. You must bring the global group online at the secondary site.
	Agent response: The agent does the following based on the SplitTakeover attribute of the HTC resource:
	■ 1—The agent issues the `horctakeover` command to make the HTC devices write-enabled. The S-VOL devices go into the SSWS state.
	■ 0—The agent faults the HTC resource.

Failure scenarios in replicated data clusters

The following table lists the failure scenarios in a replicated data cluster configuration, and describes the behavior of and the agent in response to the failure.

Table 4-2 Failure scenarios in a replicated data cluster configuration with agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access

Failure	Description and response

Application failure	Application cannot start successfully on any hosts at the primary site.
	response:
	■ Causes the service group at the primary site to fault.
	■ Does the following based on the AutoFailOver attribute for the faulted service group:
	  ■ 1— automatically brings the faulted service group online at the secondary site.
	  ■ 2—You must bring the service group online at the secondary site.
	The agent does the following:
	■ Write enables the devices at the secondary site, except when the link is manually suspended with the read-only option.
	■ Swaps the P-VOL/S-VOL role of each device in the device group.
	■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site.
	See "Performing failback after a node failure or an application failure" on page 58.
	See "Replication link / Application failure scenarios" on page 54.

Host failure	All hosts at the primary site fail.
	response:
	■ Causes the service group at the primary site to fault.
	■ Does the following based on the AutoFailOver attribute for the faulted service group:
	  ■ 1— automatically brings the faulted service group online at the secondary site.
	  ■ 2—You must bring the service group online at the secondary site.
	The agent does the following:
	■ Write enables the devices at the secondary site, except when the link is manually suspended with the read-only option.
	■ Swaps the P-VOL/S-VOL role of each device in the device group.
	■ Restarts replication from P-VOL devices on the secondary site to the S-VOL devices at the primary site.
	See "Performing failback after a node failure or an application failure" on page 58.

Table 4-2 Failure scenarios in a replicated data cluster configuration with agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access *(continued)*

Failure	Description and response

Site failure	All hosts and the storage at the primary site fail.
	response:
	■ Causes the service group at the primary site to fault.
	■ Does the following based on the AutoFailOver attribute for the faulted service group:
	  ■ 1— automatically brings the faulted service group online at the secondary site.
	  ■ 2—You must bring the service group online at the secondary site.
	Agent response: The agent does the following based on the SplitTakeover attribute of the HTC resource:
	■ 1— The agent issues the `horctakeover` command to make the HTC devices write-enabled. The HTC devices go into the SSWS (Suspend for Swapping with S-VOL side only) state. If the original primary site is restored, you must execute the pairresync-swaps action on the secondary site to establish reverse replication.
	■ 0 — Agent does not perform failover to the secondary site.
	See "Performing failback after a site failure" on page 59.

Table 4-2 Failure scenarios in a replicated data cluster configuration with agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access *(continued)*

Failure	Description and response
Replication link failure	Replication link between the arrays at the two sites fails. response: No action. Agent response: When the replication link is disconnected, the agent does the following based on the LinkMonitor attribute of the HTC resource: ■ 0—No action. ■ 1—The agent periodically attempts to resynchronize the S-VOL side using the `pairresync` command. The agent also logs a warning message to indicate that the replication link is broken. ■ 2—The agent periodically attempts to resynchronize the S-VOL side and also sends notifications about the disconnected link. Notifications are sent in the form of either SNMP traps or emails. For information about the NotifierMngr agent, refer to the *Symantec Cluster Server Bundled Agents Reference Guide*. If the value of the LinkMonitor attribute is not set to 1 or 2, you must manually resynchronize the HTC devices after the link is restored. To manually resynchronize the HTC devices after the link is restored: **1** Before you resync the S-VOL device, you must split off the Shadow Image device from the S-VOL device at the secondary site. **2** You must initiate resync of S-VOL device using the agent's pairresync action. **3** After P-VOL and S-VOL devices are in sync, reestablish the mirror relationship between the Shadow Copy and the S-VOL devices. If you initiate a failover to the secondary site when resync is in progress, the online function of the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent waits for the resync to complete and then initiates a takeover of the S-VOL devices. **Note:** If you did not configure Shadow Copy devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Veritas recommends configuring Shadow Copy devices at both the sites. See “Replication link / Application failure scenarios” on page 54.

Table 4-2 Failure scenarios in a replicated data cluster configuration with agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access *(continued)*

Failure	Description and response
Network failure	The LLT and the replication links between the sites fail. response: ■ at each site concludes that the nodes at the other site have faulted. ■ Does the following based on the AutoFailOver attribute for the faulted service group: ■ 2—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue. ■ 1— brings the service group online at the secondary site which leads to a cluster-wide split brain. This causes data divergence between the devices on the arrays at the two sites. When the network (LLT and replication) connectivity is restored, takes all the service groups offline on one of the sites and restarts itself. This action eliminates concurrency violation where in the same group is online at both the sites. After taking the service group offline, you must manually resynchronize the data. **Note:** Veritas recommends that the value of the AutoFailOver attribute is set to 2 for all service groups to prevent unintended failovers due to transient network failures. To resynchronize the data after the network link is restored: **1** Take the service groups offline at both the sites. **2** Manually resynchronize the data. Depending on the site whose data you want to retain run the `pairresync` or the `pairresync-swap` command. **3** Bring the service group online on one of the sites. Agent response: Similar to the site failure.

Table 4-2 Failure scenarios in a replicated data cluster configuration with agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access *(continued)*

Failure	Description and response

Storage failure	The array at the primary site fails.
	response:
	■ Causes the service group at the primary site to fault and displays an alert to indicate the fault.
	■ Does the following based on the AutoFailOver attribute for the faulted service group:
	  ■ 1— automatically brings the faulted service group online at the secondary site.
	  ■ 2—You must bring the service group online at the secondary site.
	Agent response: The agent does the following based on the SplitTakeover attribute of the HTC resource:
	■ 1—The agent issues the `horctakeover` command to make the HTC devices write-enabled. The S-VOL devices go into the SSWS state.
	■ 0—The agent does not perform failover to the secondary site.

Replication link / Application failure scenarios

The following table lists the link failure scenarios and recommended actions.

Table 4-3 Replication link / Application failure scenarios

Event	Fence level	Recommended action
Link fails and is restored, but application does not fail over.	never, async	Run the `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.
Action faults due to I/O errors.	data	Run the `localtakeover` action to write enable the local devices. Clear faults and restart service group.

Testing the global service group migration

After you configure the Cluster Server agent for Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access, verify that the global service group can

migrate to hosts across the sites. Depending on your DR configuration, perform one of the following procedures.

To test the global service group migration in global cluster setup

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

   Perform the following steps:

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

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

     VCS brings the global service group online on a node at the secondary site.
   - Verify that the HTC devices at the secondary site are write-enabled and the device state is PAIR.

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

   Perform the following steps:

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

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

     VCS brings the global service group online at the primary site.
   - Verify that the HTC devices at the secondary site are write-enabled and the device state is PAIR.

To test service group migration in replicated data cluster setup

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

   Perform the following steps:

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

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

     VCS brings the service group online on a node at the secondary site.
   - Verify that the HTC devices at the secondary site are write-enabled, and the device state is PAIR.

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

   Perform the following steps:

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

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

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

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

Testing disaster recovery after host failure

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

See "Failure scenarios in global clusters" on page 45.

See "Failure scenarios in replicated data clusters" on page 50.

Perform the procedure that is applicable to your DR configuration to test how VCS recovers after all hosts at the primary site fail.

To test disaster recovery for host failure in global cluster setup

1. Halt the hosts at the primary site.

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

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

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

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

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

   ```
   hagrp -state global_group
   ```

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

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

1. Halt the hosts at the primary site.

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

   - 1—VCS brings the faulted service group online at the secondary site.

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

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

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

```
hagrp -state global_group
```

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

Testing disaster recovery after site failure

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

See "Failure scenarios in global clusters" on page 45.

See "Failure scenarios in replicated data clusters" on page 50.

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

To test disaster recovery for site failure in global cluster setup

1 Halt all nodes and the arrays at the primary site.

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

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

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

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

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

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

```
hagrp -state global_group
```

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

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

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

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

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

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

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

   ```
   hagrp -state global_group
   ```

Performing failback after a node failure or an application failure

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

See “Failure scenarios in global clusters” on page 45.

See “Failure scenarios in replicated data clusters” on page 50.

After the nodes at the primary site are restarted, you can perform a failback of the global service group to the primary site.

Perform the procedure that is applicable to your DR configuration.

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

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

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

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

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

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

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

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

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

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

Performing failback after a site failure

After a site failure at the primary site, the hosts and the storage at the primary site are down. VCS brings the global service group online at the secondary site and the Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access agent write enables the S-VOL devices.

The device state is SSWS.

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

See “Failure scenarios in global clusters” on page 45.

See “Failure scenarios in replicated data clusters” on page 50.

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

To perform failback after a site failure in global cluster

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

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

2. Since the application has made writes on the secondary due to a failover, resynchronize the primary from the secondary site and reverse the P-VOL/S-VOL roles with the `pairresync-swaps` action on the secondary site.

   After the resync is complete, the devices in the secondary are P-VOL and the devices in the primary are S-VOL. The device state is PAIR at both the sites.

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

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

   This again swaps the role of P-VOL and S-VOL.

To perform failback after a site failure in replicated data cluster

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

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

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

   After the resync is complete, the devices in the secondary are P-VOL and the devices in the primary are S-VOL. The device state is PAIR at both the sites.

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

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

   This again swaps the roles of P-VOL and S-VOL.

Chapter 5

Setting up a fire drill

This chapter includes the following topics:

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

About fire drills

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

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

In clusters employing Hitachi TrueCopy / HUR / Hewlett-Packard XP Continuous Access, the HTCSnap resource manages the replication relationship during a fire drill.

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

The HTCSnap agent supports fire drill for storage devices that are managed using Veritas Volume Manager.

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

Fire drill configurations

VCS supports the following fire drill configurations for the agent:

Gold

Runs the fire drill on a snapshot of the target array. The replicated device keeps receiving writes from the primary.

Veritas recommends this configuration because it does not affect production recovery.

In the Gold configuration, VCS does the following:

- Suspends replication to get a consistent snapshot.
- Takes a snapshot of the target array on a ShadowImage device.
- Resumes replication.
- Modifies the disk group name in the snapshot.
- Brings the fire drill service group online using the snapshot data.

For Gold configurations, you must use Veritas Volume Manager to import and deport the storage.

You can use the Gold configuration only with ShadowImage pairs created without the `-m noread` flag to the `paircreate` command.

Silver

VCS takes a snapshot, but does not run the fire drill on the snapshot data. VCS breaks replication and runs the fire drill on the replicated target device.

If a disaster occurs while resynchronizing data after running the fire drill, you must switch to the snapshot for recovery.

In the Silver configuration, VCS does the following:

- Suspends replication to get a consistent snapshot.
- Takes a snapshot of the target array on a ShadowImage device.
- Resumes replication
- Modifies the disk name and the disk group name in the snapshot.
- Brings the fire drill service group online using the data on the target array; the agent does not use the snapshot data for the fire drill.

You can use the Silver configuration only with ShadowImage pairs created with the `-m noread` flag to the `paircreate` command.

Bronze	VCS breaks replication and runs the fire drill test on the replicated target. VCS does not take a snapshot in this configuration. If a disaster occurs while resynchronizing data after the test, it may result in inconsistent data as there is no snapshot data. In the Bronze configuration, VCS does the following: ■ Suspends replication. ■ Brings the fire drill service group online using the data on the target array.

Note on the Gold configuration

Perform the following steps for a successful Gold configuration fire drill.

To create a Gold configuration fire drill

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

About the HTCSnap agent

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

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

HTCSnap agent functions

The HTCSnap agent performs the following functions:

Table 5-1 Agent functions

Function	Description
online	■ Suspends replication between the source and the target arrays. ■ Takes a local snapshot of the target LUN. ■ Resumes the replication between the arrays. ■ Takes the fire drill service group online by mounting the replication target LUN. ■ Creates a lock file to indicate that the resource is online.
offline	■ Destroys the snapshot by synchronizing data between the target array and the device on which snapshot was taken. ■ Removes the lock file created by the online function. ■ Resumes replication between the source and the target arrays. ■ Synchronizes data between the target array and the device on which the snapshot was taken. Destroys the snapshot of the target array after the data is synchronized. ■ Resumes the replication between the source and the target arrays. ■ Removes the lock file created by the online operation.
monitor	Verifies the existence of the lock file to make sure the resource is online.
clean	Restores the state of the LUNs to their original state after a failed online function.

Resource type definition for the HTCSnap agent

Following is the resource type definition for the HTCSnap agent:

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

```
int MountSnapshot
int UseSnapshot
int RequireSnapshot
temp str Responsibility
temp str FDFile
temp str VCSResLock
)
```

Attribute definitions for the HTCSnap agent

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

Table 5-2 Agent attributes

Attribute	Description
ShadowInstance	The instance number of the ShadowInstance P-VOL group. The P-VOL group must include one of the following: ■ The same LUNs as in the TrueCopy S-VOL group (if taking snapshots of replicated data). ■ The same LUNs as in the VxVM disk group (if taking snapshots of non-replicated data). Type-Dimension: integer-scalar
TargetResName	Name of the resource managing the LUNs that you want to take snapshot of. Set this attribute to the name of the HTC resource if you want to take a snapshot of replicated data. Set this attribute to the name of the DiskGroup resource if the data is not replicated. For example, in a typical Oracle setup, you might replicate data files and redo logs, but you may choose to avoid replicating temporary tablespaces. The temporary tablespace must still exist at the DR site and may be part of its own disk group. Type-Dimension: string-scalar
UseSnapshot	Specifies whether the HTCSnap resource takes a local snapshot of the target array. Set this attribute to 1. Type-Dimension: integer-scalar See "About the Snapshot attributes" on page 66.

Table 5-2 Agent attributes *(continued)*

Attribute	Description
RequireSnapshot	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. Type-Dimension: integer-scalar **Note:** Set this attribute to 1 only if UseSnapshot is set to 1.
MountSnapshot	Specifies whether the resource uses the snapshot to bring the service group online. Set this attribute to 1. Type-Dimension: integer-scalar **Note:** Set this attribute to 1 only if the UseSnapshot attribute is set to 1.

About the Snapshot attributes

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

Table 5-3 Snapshot attribute values for fire drill configurations

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

If the snapshot operation fails, set the RequireSnapshot attribute to 0 to enable a Gold or a Silver configuration to run in the Bronze mode.

Before you configure the fire drill service group

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

- Make sure the application service group is configured with a HTC resource.
- Make sure the infrastructure to take snapshots is properly configured. This process involves creating the ShadowImage pairs.

- If you plan to use Gold or Silver configuration, make sure ShadowImage for TrueCopy is installed and configured at the target array.
- For the Gold configuration, you must use Veritas Volume Manager to import and deport the storage.
- You can use the Silver configuration only with ShadowImage pairs that are created with the `-m noread` flag to the `paircreate` command. A fire drill uses the `-E` flag to split the pairs, which requires a 100% resynchronization. The Silver mode that preserves the snapshots as `noread` after a split.
- The name of the ShadowImage device group must be the same as the replicated device group for both replicated and non-replicated LUNs that are to be snapshot. The instance number may be different.
- Make sure the HORCM instance managing the S-VOLs runs continuously; the agent does not start this instance.
- For non-replicated devices:
  - You must use Veritas Volume Manager.
    On HP-UX, you must use Veritas Volume Manager 5.0 MP1.
  - For Gold configuration to run without the Bronze mode, set the RequireSnapshot attribute to 1.
- Add vxdctlenable action in the list of SupportedActions for the CVMVolDg resource in an SF for Oracle RAC or a Storage Foundation Cluster File System (SFCFS) environment.
  Use the following sequence of commands:

```
haconf -makerw
hatype -modify CVMVolDg SupportedActions vxdctlenable
haconf -dump -makero
```

Configuring the fire drill service group

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

See “Sample configuration for a fire drill service group” on page 71.

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

- Cluster Manager (Java Console)

See "Creating the fire drill service group using Cluster Manager (Java Console)" on page 68.

- Fire Drill Setup wizard
  This text-based wizard is available at /opt/VRTSvcs/bin/fdsetup-htc.
  See "Creating the fire drill service group using the Fire Drill SetUp Wizard" on page 70.

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

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

Open the Cluster Manager and perform the following procedure:

To create the fire drill service group

1. Log on to the cluster and click **OK**.
2. Click the **Service Group** tab in the left pane and click the **Resources** tab in the right pane.
3. Right-click the cluster in the left pane and click **Add Service Group**.
4. In the **Add Service Group** dialog box, provide information about the new service group.
   - In Service Group name, enter a name for the fire drill service group, for example: **MyApplication_AppSG_FD**.
   - Select systems from the Available Systems box and click the arrows to add them to the Systems for Service Group box.
   - Click **OK.**
5. After the fire drill service group is successfully created, set the Failover attribute to **false** so that the fire drill service group does not fail over to another node during a test.

To disable the AutoFailOver attribute

1. Click the **Service Group** tab in the left pane and select the fire drill service group.
2. Click the **Properties** tab in the right pane.
3. Click the **Show all attributes** button.

4 Double-click the **AutoFailOver** attribute.

5 In the **Edit Attribute** dialog box, clear the **AutoFailOver** check box.

6 Click **OK** to close the **Edit Attribute** dialog box.

7 Click the **Save and Close Configuration** icon in the toolbar.

Adding resources to the fire drill service group

Add resources to the new fire drill service group to re-create key aspects of the application service group.

To add resources to the service group

1 In Cluster Explorer, click the **Service Group** tab in the left pane, click the application service group and click the **Resources** tab in the right pane.

2 Right-click the resource at the top of the tree, select **Copy** > **Self and Child Nodes**.

3 In the left pane, click the fire drill service group.

4 Right-click the right pane, and click **Paste**.

5 In the **Name Clashes** dialog box, specify a way for the resource names to be modified, for example, insert an **_FD** suffix. Click **Apply**.

6 Click **OK**.

Configuring resources for fire drill service group

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

To configure resources for the fire drill service group

1 In Cluster Explorer, click the **Service Group** tab in the left pane.

2 Click the fire drill service group in the left pane and click the **Resources** tab in the right pane.

3 Right-click the HTC resource and click **Delete**.

4 Add a resource of type HTCSnap and configure its attributes.

5 Right-click the resource to be edited and click **View** > **Properties View**. If a resource to be edited does not appear in the pane, click **Show All Attributes**.

6 Edit the attributes to reflect the configuration at the remote site. For example, change the Mount resources so that they point to the volumes that are used in the fire drill service group.

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

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

See "Fire drill configurations " on page 62.

To create the fire drill service group

1 Start the Fire Drill SetUp Wizard.

```
/opt/VRTSvcs/bin/fdsetup-htc
```

2 Enter the name of the application service group for which you want to configure a fire drill service group.

3 Select the supported snapshot configurations:

Gold, Silver, or Bronze

4 Choose whether to run a Bronze fire drill, if the snapshot fails with Gold or Silver configurations.

```
If snapshot fails, should bronze be used? [y,n,q](n)
```

5 Specify the ShadowImage instance.

6 Press **Return** to verify the snapshot infrastructure.

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

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

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

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

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

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

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

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

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

```
fdsched-htc
```

Verifying a successful fire drill

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

To verify a successful fire drill

1 Bring the fire drill service group online on a node at the secondary site where the application is not running.

If the fire drill service group comes online, it validates your DR configuration. The service group at the primary site can fail over to the secondary site in the event of an actual failure or a 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 service group offline after its functioning has been verified.

Failing to take this service group offline may cause failures in your environment. For example, if the application service group fails over to the node where the fire drill service group is hosted, both the service groups may fault owing to resource conflicts.

Sample configuration for a fire drill service group

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

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

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

```
MountSnapshot = 1
}
```