Veritas™ Cluster Server Agent for Hitachi TrueCopy Configuration Guide

ESX

5.1

symantec™

Veritas Cluster Server Agent for Hitachi TrueCopy

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

Agent version: 5.1

Legal Notice

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

http://www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder.

Technical support

For technical assistance, visit

http://www.symantec.com/enterprise/support/assistance_care.jsp

and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 4 Setting up a fire drill

Index

Chapter 1

Introducing the Veritas agent for Hitachi TrueCopy

This chapter includes the following topics:

About the agent for Hitachi TrueCopy

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

Arrays	Supported fence levels
Lightning	data, never, and async

Thunder data and never

Supported software and hardware

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

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

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

Typical Hitachi TrueCopy setup in a VCS cluster

Figure 1-1 displays a typical cluster setup in a TrueCopy environment.

Figure 1-1 Typical clustering setup for the agent

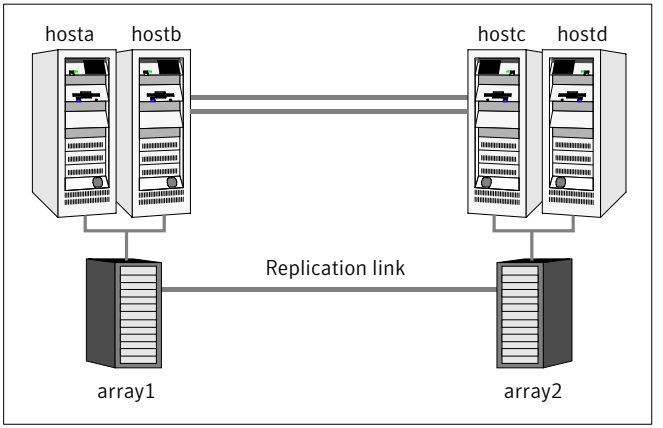

Clustering in a TrueCopy environment typically consists of the following hardware infrastructure:

- The primary array (array1) has one or more P-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to the Hitachi TrueCopy array that contains the TrueCopy P-VOL devices.

- The secondary array (array2) has one or more S-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to a Hitachi TrueCopy array that contains the TrueCopy S-VOL devices. The S-VOL devices are paired with the P-VOL devices in the P-VOL array. The S-VOL hosts and arrays must be at a significant distance to survive a disaster that may occur at the P-VOL side.
- Network heartbeating between the two data centers to determine their health; this network heartbeating could be LLT or TCP/IP.
  See "About cluster heartbeats" on page 14.
- In a global cluster environment, you must attach all hosts in a cluster to the same Hitachi TrueCopy array.

Hitachi TrueCopy agent operations

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 operations:

online	If the state of all local devices is read-write enabled, the agent creates a lock file on the local host to indicate that the resource is online. This action makes the devices writable for the application. If one or more devices are not in a writable state, the agent runs the `horctakeover` command to enable read-write access to the devices.
offline	The agent removes the lock file on the device. The agent does not run any TrueCopy commands because taking the resource offline is not indicative of an intention to give up the devices.
monitor	Verifies the existence of the lock file to determine the resource status. If the lock file exists, the agent reports the status of the resource as online. If the lock file does not exist, the agent reports the status of the resource as offline.
open	Removes the lock file on the system 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 `hastop -force` command.

clean	Determines whether if it is safe to fault the resource if the online entry point fails or times out. The main consideration is whether a management operation was in progress when the online thread timed out and was killed. 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	Resynchronizes the devices from the VCS command line after connectivity failures are detected and corrected. The agent supports the following actions: ■ pairdisplay—Displays information about all devices. ■ pairresync—Resynchronizes the S-VOLs. ■ pairresync-swaps—Promotes the S-VOLs to P-VOLs and resynchronizes the original P-VOLs. ■ localtakeover—Makes the local devices write-enabled.

Additional considerations in an ESX environment

The agent performs the following actions before coming online:

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

Chapter 2

Configuring the agent for Hitachi TrueCopy

This chapter includes the following topics:

Configuration concepts for the Hitachi TrueCopy agent

Review the configuration concepts and failure scenarios for the agent.

Resource type definition for the Hitachi TrueCopy agent

The resource type definition defines the agent in VCS.

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

```
        temp str TargetFrozen
        )
```

Attribute definitions for the Hitachi TrueCopy agent

The descriptions of the agent attributes are as follows:

BaseDir	Path to the RAID Manager Command Line interface. Type-dimension: string-scalar
GroupName	Name of the device group that is managed by the agent. Type-dimension: string-scalar
Instance	The Instance number of the device that the agent manages. Multiple device groups may have the same instance number. Do not define the attribute if the instance number is zero. Type-dimension: string-scalar
SplitTakeover	A flag that determines whether the agent permits a failover to S-VOL devices if the if the replication link is disconnected, that is, if P-VOL devices are in the PSUE state. See "About the SplitTakeover attribute for the Hitachi TrueCopy agent" on page 13. Type-dimension: integer-scalar Default: 0
LinkMonitor	A flag that defines whether the agent periodically attempts to resynchronize the S-VOL side if the replication link is disconnected. The agent uses the `pairresync` command to resynchronize arrays. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the `pairresync` command. Type-dimension: integer-scalar Default: 0
TargetFrozen	For internal use. Do not modify.

About the SplitTakeover attribute for the Hitachi TrueCopy agent

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

The default value for this attribute is 0. The value 0 indicates that the agent does not permit a failover to S-VOL devices if the P-VOL devices are in the PSUE state. If a failover occurs when the replication link is disconnected, data loss may occur because the S-VOL devices may not be in sync.

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

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

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

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

Sample configuration for the Hitachi TrueCopy agent

Figure 2-1 shows a dependency graph of a VCS service group that has a resource of type HTC.

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

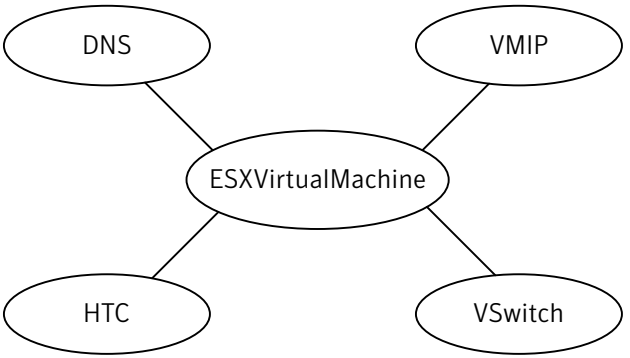

You can configure a resource of type HTC in the main.cf file as:

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

Before you configure the agent for TrueCopy

Before you configure the agent, review the following information:

- Review the configuration concepts, which describe the agent's type definition and attributes.
  See "Configuration concepts for the Hitachi TrueCopy agent" on page 11.
- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
- Make sure that the cluster has an effective heartbeat mechanism in place.
  See "About cluster heartbeats" on page 14.
  See "About preventing split-brain" on page 15.

About cluster heartbeats

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

Hitachi arrays do not support a native heartbeating mechanism between the arrays. The arrays send a support message on detecting replication link failure.

You can take appropriate action to recover from the failure and to keep the devices in a synchronized state. The TrueCopy agent supports those actions that can automate the resynchronization of devices after a replication link outage is corrected.

About preventing split-brain

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

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

Configuring the agent for TrueCopy

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

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

Configuring the agent manually in a global cluster

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

To configure the agent in a global cluster

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

4. Save the configuration.
5. Add a resource of type HTC at the bottom of the service group.
6. Configure the attributes of the HTC resource.
7. If the service group is not configured as a global group, configure the service group using the Global Group Configuration Wizard. See the *Veritas Cluster Server User's Guide* for more information.
8. Change the ClusterFailOverPolicy from the default, if necessary. Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.

   Repeat step 5 through step 8 for each service group in each cluster that uses replicated data.

Chapter 3

Managing and testing clustering support for Hitachi TrueCopy

This chapter includes the following topics:

Typical test setup for the Hitachi TrueCopy agent

A typical test environment includes the following characteristics:

- Two hosts (hosta and hostb) are attached to the P-VOL array.
- Two hosts (hostc and hostd) are attached to the S-VOL array.
- The application runs on hosta and devices in the local array are P-VOLs in the PAIR state.

Figure 3-1 depicts a typical test environment.

Figure 3-1 Typical test setup

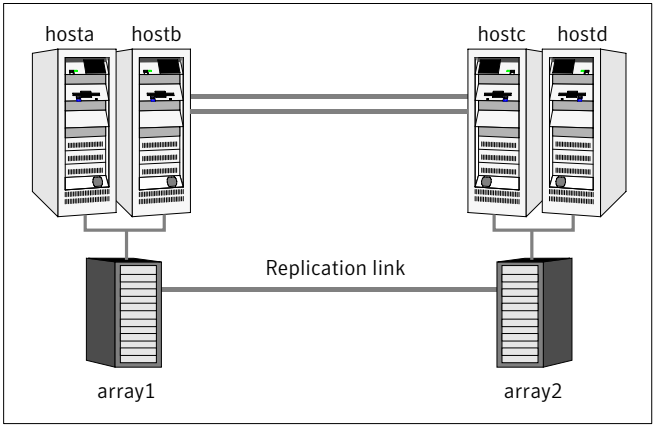

Testing service group migration

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

To perform the service group migration test

1. Migrate the service group to a host that is attached to the same array. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
2. Click **Switch To**, and click the system that is attached to the same array (hostb) from the menu.

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

   The service group comes online on hostc and volumes there transition to the P-VOL/PAIR state, changing the original P-VOLs to S-VOLs.

5 Migrate the service group back to its original host. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.

6 Click **Switch To**, and click the system on which the group was initially online (hosta).

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

Testing host failure

In this scenario, the host where the application runs is lost.

To perform the host failure test

1 Halt or shut down the host where the application runs (hosta).

2 Halt or shut down hostb.

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

3 Reboot the two hosts that were shut down. A swap resynchronization is required to demote the original P-VOLs:

```
hares -action HTCRes pairresync-swaps -sys system
```

4 Switch the service group to its original host when VCS starts. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.

5 Click **Switch To**, and click the system on which the service group was initially online (hosta).

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

Performing a disaster test

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

To perform a disaster test

1. Shut down all hosts on the source side and shut down the source array.

   If you can not shut down the primary array, disconnect the replication link between the two arrays and simultaneously shut down the hosts. This action mimics a disaster scenario to the secondary side.
2. In a global cluster, the administrator is notified of the failure. The administrator can then initiate the failover.

Performing the failback test

You can set up your cluster for a failback test.

To perform a failback test

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

Failure scenarios for Hitachi TrueCopy

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

Site disaster

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

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

All host or all application failure

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

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

In both environments, multiple service groups can fail over in parallel.

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

Replication link failure

Hitachi arrays send an alert in the following situations:

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

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

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

Table 3-1 shows the situations that require administrative action after you repair a link failure.

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

Table 3-1 Replication link failure scenarios

Event	Fence Level	Recommended Action
Link fails and is restored, but application does not fail over.	never, async	Run the `pairresync` action to resynchronize the S-VOLs.

Table 3-1 Replication link failure scenarios *(continued)*

Event	Fence Level	Recommended Action
Link fails and application fails to the S-VOL side.	never, async, or data	Run the `pairresync-swaps` action to promote the S-VOLs to P-VOLs and resynchronize the original P-VOLs.
Application faults due to I/O errors.	data	Run the `localtakeover` action to write-enable the local devices. Clear faults and restart service group.

Split-brain in a TrueCopy environment

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

Rescanning Host Bus Adapters (HBAs) on VCS nodes

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

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

Chapter 4

Setting up a fire drill

This chapter includes the following topics:

About fire drills

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

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

In clusters employing Hitachi TrueCopy, the HTCSnap resource manages the replication relationship during a fire drill.

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

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

About the HTCSnap agent

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

Refer to the Figure 4-1 for the HTCSnap agent dependency tree.

Figure 4-1 HTCSnap agent dependency tree

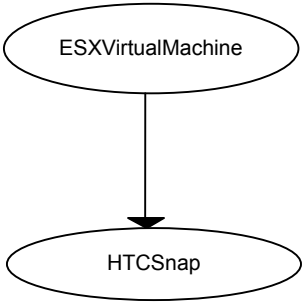

HTCSnap agent operations

TheHTCSnap agent performs the following operations:

Online	Destroys any existing snapshot and takes a new snapshot. See "About the agent's online operation" on page 24.
Offline	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 operation.
Action	For internal use.

About the agent's online operation

The agent's online function performs the following actions:

- Resynchronizes snapshots LUNs with source LUNs.
- If TargetResName is of type HTC (remote fire drill), the agent suspends TrueCopy replication to get a consistent snapshot.

- Splits the ShadowImage replication into PSUS-SSUS state in order to obtain a new writeable snapshot.
- If TargetResName is of type HTC (remote fire drill), the agent resumes Truecopy replication between the arrays.
- Scans for new datastores with following VMware LVM settings.
  EnableResignature = 1;
  DisAllowSnapshotLun = 0;
- Makes the snapshot of the virtual machine ready for the fire drill. Updates the dependent ESXVirtualMachine resource with the full path to the snapshot of the .vmx file.
  See "Processing the snapshot" on page 25.
  See "Configuring the fire drill service group" on page 29.
- Creates a lock file to indicate that the resource is online.

Processing the snapshot

- Temporarily sets following VMware LVM settings and scans for new datastores:
  Refer to the VMware documentation for more information on these settings.
- Registers the new virtual machine with the ESX host using the new snapshot datastore that was created.
- Populates the CfgFile attribute of the ESXVirtualMachine resource (parent resource) in the fire drill service group. The display name of the new virtual machine is the value of the vmname attribute, suffixed by _snapshot. The suffix clearly identifies the virtual machine as the snapshot virtual machine.
- Generates new MAC addresses for each vmnic and patches the configuration files of the new virtual machine with the new MAC addresses.
- Disconnects the new virtual machine from the public switches by setting the ethernet[n].networkName attribute in the vmx file to None. This action enables you to configure IP address and other parameters in the guest OS of the virtual machine after it boots up.

Resource type definition for the HTCSnapagent

Following are the resource type definitions for the HTCSnapSnap agent:

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

```
int ShadowInstance
int MountSnapshot = 1
int UseSnapshot = 1
int RequireSnapshot = 1
temp str Responsibility
temp str FDFile
)
```

HTCSnap agent attributes

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

ShadowInstance	The instance number of the ShadowInstance P-VOL group. Type-dimension: integer-scalar The P-VOL group must include one of the following: ■ The same LUNs as in the TrueCopy S-VOL group (for remote fire drill) or the TrueCopy P-VOL group (for local fire drill).
TargetResName	Name of the resource managing the LUNs that you want to take snapshot of. Set this attribute to the name of HTC resource for running the remote firedrill. Set this attribute to the name of the ESXVirtualMachine resource for running the local fire drill. The local fire drill runs on the primary site and the remote fire drill runs on the secondary site. Type-dimension: string-scalar
UseSnapshot	Specifies whether the HTCSnap resource takes a local snapshot of the target array. Default Value : 1 Set this attribute to 1. The current release of the agent supports only the default value. Type-Dimension: integer-scalar
RequireSnapshot	Specifies whether the HTCSnap resource must take a snapshot before coming online. Default Value : 1 Set this attribute to 1. The current release of the agent supports only the default value. Type-Dimension: integer-scalar

MountSnapshot	Specifies whether the resource uses the snapshot to bring the service group online. Default Value : 1 Set this attribute to 1. The current release of the agent supports only the default value. Type-Dimension: integer-scalar

Internal attributes

Do not modify internal attributes:

Responsibility	Do not modify. For internal use only. Used by the agent to keep track of resynchonizing snapshots. Type-Dimension: temporary string
FDFile	Do not modify. For internal use only. Used by the agent to locate the latest fire drill report. Type-dimension: temporary string

Sample configuration for a fire drill service group

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

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

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

Before you configure the fire drill service group

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

- Make sure the application service group is configured with a TrueCopy resource.
- Make sure the infrastructure to configure hardware snapshots is properly configured.
- Make sure HTC ShadowImage is installed and configured at the target array.
- The name of the ShadowImage device group must be the same as the Truecopy device group. The instance number may be different.

Additional requirements for running a fire drill in an ESX environment

Follow these guidelines for fire drills in an ESX environment:

- Ensure that each service group contains only one virtual machine resource to be able to support fire drill.
- Symantec recommends dedicating one node in the cluster for fire drills. Configure fire drill service groups on this node. Do not configure replication resources on this node.
- Configure your array such that snapshot LUNs (i.e. target LUNs on which hardware snapshots are taken) are visible to this node only; the other nodes must not see snapshot LUNs. The restriction occurs because other nodes in the cluster may have LVM settings LVM.EnableResignature = 0 and LVM.DisAllowSnapshotLUN = 0 set by the replication agents, in which case VMWare recommends that no snapshot LUNs should be exposed to such ESX hosts.

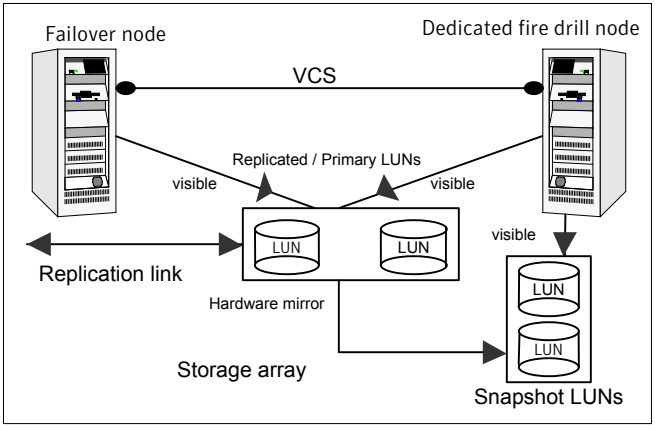

- Install VMware tools on all virtual machines in the fire drill configuration. Otherwise, the missed virtual machine heartbeats cause the virtual machine resource to fault.
- Ensure that no virtual machines are configured with raw device mapping.
- In each fire drill service group, ensure that the CfgFile attribute of the ESXVirtualMachine agent resource is left blank. The fire drill process fills in this value. When the value of the CfgFile attribute is blank, Cluster Explorer displays a red question mark on the ESXVirtualMachine resource.
- Ensure that the resources configured within the fire drill service group are not marked as critical. Not marking the fire drill service group resources as critical prevents service group failover.

Configuring the fire drill service group

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

To create the fire drill service group

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

    See "Sample configuration for the Hitachi TrueCopy agent" on page 13.
11. Link resources such that the ESXVirtualMachine resource depends on the HTCSnap resource.

Running the fire drill

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

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

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

You are now ready to run the fire drill.

To run the fire drill

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

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

   Failing the take the fire drill service group offline could cause failures in your environment. For example, if the application service group fails over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

Index