

Cluster Server Agent for IBM SVC Copy Services Installation and Configuration Guide

AIX, HP-UX, Linux, Solaris

7.0

Veritas InfoScale™ Availability Agents

Last updated: 2021-07-07

Legal Notice

Copyright © 2021 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 IBM SVC Copy Services	6
	6
	About the agent for IBM SVC Copy Services	6
	Supported software	7
	Supported hardware for IBM SVC Copy Services	7
	Typical IBM SVC Copy Services setup in a VCS cluster	7
	SVCCopyServices agent functions	9
	About the SVCCopyServices agent's online function	10
	About the SVCCopyServices agent's update action function	13
	About the SVCCopyServices agent's info function	14
 Chapter 2	 Installing and removing the agent for IBM SVC Copy Services	 15
	15
	Before you install the agent for IBM SVC Copy Services	15
	Installing the agent for IBM SVC Copy Services	16
	Installing the agent IPS package on Oracle Solaris 11 systems	17
	Upgrading the agent for IBM SVC Copy Services	18
	Removing the agent for IBM SVC Copy Services	19
 Chapter 3	 Configuring the agent for IBM SVC Copy Services	 20
	20
	Before you configure the agent for IBM SVC Copy Services	20
	About cluster heartbeats	20
	About configuring system zones in replicated data clusters	21
	About preventing split-brain	21
	Configuration concepts for the SVCCopyServices agent	22
	Resource type definition for the SVCCopyServices agent	22
	Attribute definitions for the SVCCopyServices agent	23
	Sample configuration for the IBM SVCCopyServices agent	25
	Configuring the agent for IBM SVC Copy Services	26
	Configuring the agent manually in a global cluster	26
	Configuring the agent in an SF for Oracle RAC environment or Storage Foundation Cluster File System (SFCFS) environment	27

Chapter 4	Managing and testing clustering support for IBM SVC Copy Services	28
	Typical test setup for the IBM SVC Copy Services agent	28
	Testing service group migration	29
	Testing host failure	30
	Testing cluster failure	31
	Performing a disaster test	31
	Performing the failback test	32
	Failure scenarios for IBM SVC Copy Services	33
	Site disaster	33
	All host or all application failure	33
	Cluster failure	33
	Replication link failure	34
	Split-brain in a SVC Copy Services environment	34
Chapter 5	Setting up a fire drill	35
	About fire drills	35
	Fire drill configurations	36
	About the SVCCopyServicesSnap agent	37
	Configuring the agent for LVM support	37
	SVCCopyServicesSnap agent functions	37
	Resource type definition for the SVCCopyServicesSnap agent	38
	Attribute definitions for the SVCCopyServicesSnap agent	39
	About the Snapshot attributes	41
	Before you configure the fire drill service group	41
	Configuring the fire drill service group	41
	Creating the fire drill service group using Cluster Manager (Java Console)	42
	Verifying a successful fire drill	44
	Sample configuration for a fire drill service group	45

Introducing the agent for IBM SVC Copy Services

This chapter includes the following topics:

- [About the agent for IBM SVC Copy Services](#)
- [Supported software](#)
- [Supported hardware for IBM SVC Copy Services](#)
- [Typical IBM SVC Copy Services setup in a VCS cluster](#)
- [SVCCopyServices agent functions](#)

About the agent for IBM SVC Copy Services

The Cluster Server agent for IBM SVC Copy Services manages replication relationships and consistency groups that are defined on SVC clusters. An SVC cluster brings storage devices together in a virtual pool to make all storage appear as one logical device to centrally manage and to allocate capacity as needed.

Each resource managed by the agent manages one replication relationship or one consistency group defined on a specific SVC cluster. The agent supports inter-cluster replication relationships. The agent does not support intra-cluster replication relationships or intra-cluster consistency groups.

The attributes of the resource managed by the SVC Copy Services agent contain the necessary information about the replication relationship or consistency group managed by the resource. For example, the SVC cluster IP address that is used to communicate with the SVC cluster; the absolute path of the SSH identity key file location, and the absolute path of the SSH binary required for communicating with the SVC cluster.

The SVC Copy Services agent supports MetroMirror (synchronous replication) and Global Mirror (asynchronous replication).

Note: The SVC Copy Services agent also supports IBM Storwize V7000.

Supported software

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

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

Supported hardware for IBM SVC Copy Services

The agent also supports SSH access to SVC.


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

Note: Refer to the IBM SVC documentation for more information on the distance limitations of the SVC metro mirror and global mirror configuration.

Typical IBM SVC Copy Services setup in a VCS cluster

[Figure 1-1](#) displays a typical setup in a SVC Copy Services environment.

Figure 1-1 Typical clustering setup for the agent



Clustering in a SVC Copy Services environment typically consists of the following hardware infrastructure:

- 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 array.
- In the SVC cluster, one node acts as the point of contact or the preferred node. Each SVC cluster has a preferred node, through which replication occurs across the fiber.
- To access an SVC cluster from a host, you need an SSH identity file on that host. The SVC Copy Services agent assumes that the SVC cluster has the information about the host and the public key (generated as a counterpart of the private key on the host) has been uploaded to the SVC cluster. This ensures that the SVC cluster can identify the host from which the agent invokes the SVC commands via SSH.
- 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 SVC only in a global cluster environment and not in a replicated data cluster environment.

SVCCopyServices agent functions

The VCS agent for IBM SVC Copy Services manages the replication relationships or consistency groups that are defined on an SVC cluster.

The agent performs the following functions:

online	<p>If the state of all local devices is read-write enabled, the agent creates a lock file on the local host to indicate that the resource is online.</p> <p>See “About the SVCCopyServices agent's online function” on page 10.</p>
offline	<p>The agent removes the lock file that was created for the resource by the online function. The agent does not run any SVCCopyServices commands because taking the resource offline does not indicate that the direction of the replication needs to be reversed or even that the replication should be stopped.</p>
monitor	<p>Verifies the existence of the lock file to determine the resource status. If the lock file exists, the agent reports the status of the resource as online. If the lock file does not exist, the agent reports the status of the resource as offline.</p>
clean	<p>The agent removes the lock file from the local host. The agent does not run any SVCCopyServices commands because taking the resource offline does not indicate that the direction of the replication needs to be reversed or even that the replication should be stopped.</p>
action / update	<p>The update action is invoked to resync the SVC disks at the primary site with the up to date data from the secondary site. A need to resync the data arises when the agent brings the SVC Copy Services resource online on the secondary and the replication is in a stopped or disconnected state.</p> <p>The action updates and resynchronizes the differences in the data between the primary and secondary sites.</p> <p>It allows the user to determine which site should be maintained as the primary site in the replication and updates it with the data from the secondary site.</p> <p>The update action supports the following commands:</p> <ul style="list-style-type: none"> ■ <code>starttrcrelationship</code> ■ <code>starttrcconsistgrp</code> <p>The action needs one argument i.e. the SVC cluster (master or auxiliary) that the user wants to retain as the Primary.</p> <p>See “About the SVCCopyServices agent's update action function” on page 13.</p>

action/PreSwitch	<p>Ensures that the remote site cluster can come online during a planned failover within a GCO configuration. The VCS engine on the remote cluster invokes the PreSwitch action on all the resources of the remote site during a planned failover using the <code>hagrp -switch</code> command. For this, the PreSwitch attribute must be set to 1. The option <code>-nopre</code> indicates that the VCS engine must switch the servicegroup regardless of the value of the PreSwitch service group attribute.</p> <p>If running the PreSwitch action fails, the failover should not occur. This minimizes the application downtime and data loss .</p> <p>For more information on the PreSwitch action and the PreSwitch feature in the VCS engine, refer to the <i>Component User's Guide</i>.</p>
open	<p>Removes the online lock file on the host where the entry point is called. This function prevents potential concurrency violation if the service group fails over to another node while this host is down.</p> <p>The agent does not remove the lock file if the agent was started after running the <code>hastop -force</code> command.</p>
info	<p>Reports the detailed status of the replication relationship or consistency group being monitored by the agent.</p> <p>See “About the SVCCopyServices agent's info function” on page 14.</p>

About the SVCCopyServices agent's online function

If the local SVC cluster is the primary, the host has read/write access to the disks and the online entry point creates a lock file and exits.

If the local SVC cluster is not the primary, it may takeover the role of the primary in the replication. The takeover depends on the state of the relationship.

- **Inconsistent Copying** - In this state, the agent waits for the remainder of the online entry point time (before the online entry point times out) for the replication to move out of this state. If it does not move out of the `inconsistent_copying` state, the online entry point times out. Else, the online entry point takes an appropriate action based on the current state of the replication.
- **Idling / Idling Disconnected** - The online entry point exits without taking any action since neither of the SVC clusters is the primary and read-write access to the disks is enabled from both SVC clusters.
- **Consistent Stopped** - There are two possible scenarios in this state:
When the primary and secondary are in sync and if `StopTakeOver = 1`: The online entry point runs the `switchrc` command to switch roles of the primary and secondary sites.

When the primary and secondary are not in sync and if StopTakeOver = 1: The online entry point runs the `stoprc` command with the `-access` option.

If the switch between the primary and secondary, or the `stoprc` command with the `-access` option is successful, the online entry point creates the lock file and exits.

If StopTakeover = 0, or if the `switchrc` command fails, or if the `stoprc` command with the `-access` option fails, the online entry point exits without creating the online lock file.

- Consistent Synchronized - The online entry point runs the `switchrc` command to switch the roles of replication.

If the switch is successful, the online entry point creates the lock file, else the online entry point exits without creating the lock file.

- Consistent Disconnected - The online entry point takes action only if the attribute DisconnectTakeover = 1. If DisconnectTakeover = 1, the online entry point runs the `stoprc` command with the `-access` option.

If the `stoprc` command is successful, then the online entry point creates the lock file and exits, else the online entry point exits without creating the lock file.


- Inconsistent Stopped / Inconsistent Disconnected - The online entry point does not do anything to enable read/write access to the local SVC cluster. The online entry point simply exits in these states.

- Consistent Copying - If the value of the CopyTakeover attribute is set to 1, the online entry point does the following:

- Runs the `stoprc` command.
- Changes the cycling mode to none.
- Runs the `startrc` command with the `-force` option.
- Waits for the state to change to `consistent_synchronized`.
- Runs the `switchrc` command to change the direction of replication.
- Runs the `stoprc` command.
- Changes the cycling mode back to multi.
- Runs the `startrc` command.

If these actions fail, the online entry point runs the `stoprc` command with the `-access` option. If the `stoprc` command is successful, the online entry point creates the lock file and exits. If the `stoprc` command fails, the online entry point exits without creating the lock file.

Figure 1-2 The algorithm for the online entry point



About the SVCCopyServices agent's update action function

The update action is invoked when the agent brings the SVC Copy Services resource online on the secondary site. It resynchronizes the data between the primary and secondary sites.

Warning: Veritas recommends that the update action is run only when the resource is offline on both the primary and secondary sites.

The action requires exactly one argument i.e. the SVC cluster (master or auxiliary), which the user wants to retain as the primary. Alternately, the action requires no arguments, in which case the existing primary is retained. However, if you do not specify the direction of replication, there is a possibility of data corruption.

- **Idling / Idling Disconnected** - In these states, there is no primary defined for the replication relationship or the consistency group. The update action expects exactly one argument i.e. the SVC cluster that the user wants to retain as the primary (one of master or auxiliary) for the replication relationship or the consistency group. The action fails with an appropriate error message if any other value is used. The update action starts the replication by specifying the primary.

For example: Assume that the application is defined by the global service group, `grp_oradb` and the replication resource in this global service group is defined by the SVC Copy Services resource `svc_mmrel`. If the replication relationship managed by the resource `svc_mmrel` is in the `idling` or `idling_disconnected` state:

- Ensure that the application global service group `grp_oradb` is offline on all the clusters (assuming that the user wants to make the master SVC cluster the primary for the relationship).
- Then, start the update action for `svc_mmrel` on system `hosta` using the following command: `hares -action svc_mmrel update -actionargs master -sys hosta`
- **Consistent Stopped / Consistent Disconnected / Consistent Copying**- There is a primary already defined for the replication or the consistency group. The update action expects no arguments. In this state, the action simply starts the replication and ignores any argument that may be used. The direction of the replication remains unchanged.

For example: Assume that the application is defined by the global service group, `grp_oradb` and the replication resource in this global service group is defined by the SVC Copy Services resource `svc_mmrel`. If the replication relationship managed by the resource `svc_mmrel` is in the `consistent_stopped`, `consistent_disconnected`, or `consistent_copying` state:

- Ensure that the application global service group grp_oradb is offline on all the clusters.
- Then, start the update action for svc_mmrel on system hosta using the following command: `hares -action svc_mmrel update -sys hosta.`

When the primary and secondary are not in sync, the update action uses the `-force` flag to start the replication.

Veritas recommends that the update action is run only after the status of the replication changes to online. If the update action is run as soon as the link is restored, it may fail because the status of the replication is `io_channel_offline`.

About the SVCCopyServices agent's info function

The info entry point captures the entire output from `svcinfolsrcrelationship` or the `svcinfolsrconsistgrp` for the replication relationship or consistency group monitored by the SVC Copy Services resource.

The output updates as the value of the ReplicationStatus key in the ResourceInfo attribute for the SVC Copy Services resource.

To view the current replication status as stored in the ResourceInfo attribute on system hosta, use the command: `hares -value svc_mmrel ResourceInfo -sys hosta.`

When the resource faults or goes offline, the ResourceInfo attribute for that resource is marked as stale. This indicates that the value in the ResourceInfo attribute is not the latest information but from some time in the past. The TS key in the attribute has the timestamp of when the ResourceInfo was last modified. Therefore, it is likely that the command `svcinfolsrcrelationship` runs on the SVC cluster and results in a different output from what is stored in the ResourceInfo attribute.

The info entry point for a resource does not get invoked on a system where the resource is not currently online.

Note: The attributes ActionTimeout and InfoTimeout for the action and info entry points do not influence the SVC Copy Services agent. The agent always allows the action and info entry points to run.

Installing and removing the agent for IBM SVC Copy Services

This chapter includes the following topics:

- [Before you install the agent for IBM SVC Copy Services](#)
- [Installing the agent for IBM SVC Copy Services](#)
- [Installing the agent IPS package on Oracle Solaris 11 systems](#)
- [Upgrading the agent for IBM SVC Copy Services](#)
- [Removing the agent for IBM SVC Copy Services](#)

Before you install the agent for IBM SVC Copy Services

Before you install the VCS agent for IBM SVC Copy Services, 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 IBM SVC Copy Services setup in a VCS cluster”](#) on page 7.

Installing the agent for IBM SVC Copy Services

You must install the IBM SVC Copy Services 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.

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>.

You can download the complete Agent Pack `tar` file or the individual agent `tar` file.

- 2 Uncompress the file to a temporary location, say `/tmp`.
- 3 If you downloaded the complete Agent Pack `tar` file, navigate to the directory that contains the package for the platform that is applicable to your environment.

```
AIX      cd1/aix/vcs/replication/svccopyservices_agent/  
agent_version/pkgs/
```

```
HP-UX    cd1/hpux/generic/vcs/replication/svccopyservices_agent/  
(PA)    agent_version/PA/depot/
```

```
HP-UX    cd1/hpux/generic/vcs/replication/svccopyservices_agent/  
(IA)    agent_version/IA/depot
```

```
Linux    cd1/linux/generic/vcs/replication/svccopyservices_agent/  
agent_version/rpms/
```

```
Solaris  cd1/solaris/dist_arch/vcs/replication/svccopyservices_agent/  
agent_version/pkgs/
```

(AIX, HP-UX, or Solaris) If you downloaded the individual agent `tar` file, navigate to the `pkgs` directory.

(Linux) If you downloaded the individual agent `tar` file, navigate to the `rpms` directory.

4 Log in as a superuser.

5 Install the package.

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

```
HP-UX        # swinstall -s 'pwd' VRTSvcssvc  
(IA/PA)
```

```
Linux        # rpm -ihv \  
VRTSvcssvc-AgentVersion-Linux_GENERIC.noarch.rpm
```

```
Solaris      # pkgadd -d . VRTSvcssvc
```

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 `VRTSvcssvc.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/VRTSvcssvc.p5p Veritas
```

4 Install the package.

```
# pkg install --accept VRTSvcssvc
```

- 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 IBM SVC Copy Services

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

To upgrade the agent software

- 1 Save the VCS configuration and stop the VCS engine.

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

```
# hstop -all -force
```

- 2 Remove the previous version of the agent from the node.

See [“Removing the agent for IBM SVC Copy Services”](#) on page 19.

- 3 Delete the `/etc/VRTSvcs/conf/config/SVCCopyServicesTypes.cf` file.

- 4 Install the latest version of the agent.

See [“Installing the agent for IBM SVC Copy Services”](#) on page 16.

- 5 Copy the file `SVCCopyServicesTypes.cf` from the `/etc/VRTSvcs/conf/` to the `/etc/VRTSvcs/conf/config` directory.

- 6 Repeat step 2 through step 5 on each node.

- 7 From a node in the cluster, edit the `/etc/VRTSvcs/conf/config/main.cf` configuration file.

Configure the new attributes, if applicable.

- 8 Verify the configuration by running the following command:

```
# hacf -verify config
```

- 9 Start VCS on the local node first.

```
# hstart - local
```

- 10 Start VCS on other nodes.

Removing the agent for IBM SVC Copy Services

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

You must remove the SVC Copy Services agent from each node in the cluster.

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

```
AIX          # installp -u VRTSvcssvc.rte
```

```
HP-UX        # swremove VRTSvcssvc
```

```
Linux        # rpm -e VRTSvcssvc
```

```
Solaris      # pkgrm VRTSvcssvc
```

Note: To uninstall the agent IPS package on a Solaris 11 system:

```
# pkg uninstall VRTSvcssvc
```

Configuring the agent for IBM SVC Copy Services

This chapter includes the following topics:

- [Before you configure the agent for IBM SVC Copy Services](#)
- [Configuration concepts for the SVCCopyServices agent](#)
- [Configuring the agent for IBM SVC Copy Services](#)

Before you configure the agent for IBM SVC Copy Services

Before you configure the agent, review the following information:

- Set up the SSH identity file on the VCS hosts prior to configuring the service group. Use the SSH keygen, if required. Generate a public and private key pair using the ssh-keygen utility. Copy the public key on the SVC cluster.
- Verify that you have installed the agent on all the cluster nodes.
- Verify the hardware setup for the agent.
See [“Typical IBM SVC Copy Services setup in a VCS cluster”](#) on page 7.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 20.

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.

About configuring system zones in replicated data clusters


In a replicated data cluster, you can prevent unnecessary SVC Copy Services 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 3-1](#) 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 3-1 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 Master to Auxiliary and Auxiliary to Master. 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.

Configuration concepts for the SVCCopyServices agent

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

Resource type definition for the SVCCopyServices agent

Following is the resource type definition for the SVC Copy Services agent:

```
type SVCCopyServices (
    static keylist SupportedActions = { update, PreSwitch }
    static int OfflineMonitorInterval = 0
    static int OnlineTimeout = 600
    static int OpenTimeout = 180
    static int RestartLimit = 1
    static str ArgList[] = { SSHBinary, SSHPathToIDFile,
        GroupName, IsConsistencyGroup, SVCClusterIP, SVCUserName,
        DisconnectTakeover, StopTakeover, CopyTakeover }
    str SSHBinary = "/usr/bin/ssh"
    str SSHPathToIDFile
    str GroupName
    int IsConsistencyGroup = 1
    str SVCClusterIP
    str SVCUserName = "admin"
    int DisconnectTakeover = 0
    int StopTakeover = 0
    int CopyTakeover = 0
    temp str VCSResLock
)
```

Attribute definitions for the SVCCopyServices agent

The descriptions of the agent attributes are as follows:

GroupName	<p>Name of the replication relationship or consistency group that is managed by the agent.</p> <p>Type-dimension: string-scalar</p>
IsConsistency Group	<p>Indicates whether the value specified in the GroupName attribute is the name of a single replication relationship or of a consistency group consisting of several replication relationships.</p> <p>Attribute value is either 0 or 1. Default is 1.</p> <p>Type-dimension: integer-scalar</p>
SSHBinary	<p>Contains the absolute path to the SSH binary. SSH is the mode of communication with the SVC cluster that is connected to the node.</p> <p>Default is "/usr/bin/ssh".</p> <p>Type-dimension: string-scalar</p>
SSHPathToID File	<p>Contains the absolute path to the identity file used for authenticating the host with the SVC cluster. The corresponding public key must be uploaded on the SVC cluster so that the SVC cluster can correctly authenticate the host.</p> <p>Type-dimension: string-scalar</p>
SVCClusterIP	<p>The IP address or name that is used to manage the SVC cluster in the dot notation. The agent uses this IP address to communicate with the SVC cluster.</p> <p>Type-dimension: string-scalar</p>
SVCUserName	<p>User name that authenticates the SSH connection with the SVC cluster.</p> <p>Default is admin.</p> <p>Type-dimension: string-scalar</p>

StopTakeover

Determines whether the agent makes read-write access available to the host when the replication is in a stopped state (i.e. `consistent_stopped`).

The status of the replication goes into a stopped state when the user fires the `stopprrelationship` or the `stopprconsistgrp` command. Thus, no replication occurs between the primary and secondary SVC clusters.

Attribute value is either 0 or 1. Default value is 0. If it is set to 1, there is a possibility for data loss, if after the replication was stopped, the application continues to write data on the Primary. Thus, when the agent enables read/write access on the secondary SVC cluster, the secondary SVC cluster does not have up-to-date data on it.

The possible stopped states are:

`inconsistent_stopped`

`consistent_stopped`

When the state of the replication is `consistent_stopped` and `StopTakeover = 1`, the agent enables read-write access for the SVC cluster.

When the state of the replication is `inconsistent_stopped`, the agent does not enable read-write access for the SVC cluster.

Type-dimension: integer-scalar

Disconnect Takeover

Determines whether the agent makes read-write access available to the host when the replication is in a disconnected state (i.e. `consistent_disconnected`).

The status of the replication goes into a disconnected state when the primary and secondary SVC clusters lose communication with each other. Thus, no replication occurs between the primary and secondary SVC clusters.

Attribute value is either 0 or 1. Default is 0.

The possible disconnected states are:

`idling_disconnected`

`inconsistent_disconnected`

`consistent_disconnected`

When the state of the replication is `consistent_disconnected` and `DisconnectTakeover = 1`, then the agent enables read/write access for the SVC cluster. When the state of the replication is `idling_disconnected`, the agent does not enable read-write access for the SVC cluster.


Type-dimension: integer-scalar

CopyTakeover	<p>Determines whether the agent makes read-write access available to the host when the replication is in a consistent_copying state.</p> <p>The agent attempts to synchronize the data before making read/write access available. The status of the replication is in a consistent_copying state when the multi cycling mode is used. In this mode, the data is always consistent, but can lag behind.</p> <p>The value of this attribute can be set to 0 or 1. If it is set to 1, and the application continues to write data on the primary even after replication was stopped, there is a possibility of data loss. In such a case, when the agent enables read/write access on the secondary SVC cluster, the secondary SVC cluster does not have the most current data.</p> <p>Type-dimension: integer-scalar</p> <p>Default value: 0</p>
VCSResLock	<p>This attribute is used for internal purpose.</p> <p>Type-dimension: temporary string</p>

Sample configuration for the IBM SVCCopyServices agent

The following figure shows a dependency graph of a VCS service group that has a resource of type SVCCopyServices.

Figure 3-2 VCS service group with resource type SVCCopyServices



The sample SVCCopyServices resource configuration to manage a relationship is:

```

SVCCopyServices oradata_svc_relationship
(
    GroupName = oradata_svc_relationship

```

```
SVCClusterIP = "<IP address>"
IsConsistencyGroup = 0
)
```

The sample SVCCopyServices resource configuration to manage a consistency group is:

```
SVCCopyServices oradata_svc_consistency_grp
(
  GroupName = oradata_svc_consistency_grp
  SVCClusterIP = "<IP address>"
)
```

Configuring the agent for IBM SVC Copy Services

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

- Converting their devices to SVC Copy Services devices
- Synchronizing the devices
- Adding the IBM SVC Copy Services agent to the service group

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

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 IBM SVC Copy Services fails to detect a change in the replication state if the role reversal is done externally and RoleMonitor is disabled.

Configuring the agent manually in a global cluster

The following procedure describes the tasks involved in configuring the agent manually in a global cluster.

To configure the agent in a global cluster

- 1 Start Cluster Manager (Java Console) and log on to the cluster.
- 2 Add a resource of type SVCCopyServices at the bottom of the service group.
- 3 Configure the attributes of the SVCCopyServices resource.
- 4 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.

- 5 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.
- 6 Repeat step 2 through step 5 for each service group in each cluster that uses replicated data.

Configuring the agent in an SF for Oracle RAC environment or Storage Foundation Cluster File System (SFCFS) environment

To configure the agent to manage the volumes that Storage Foundation for Oracle RAC uses, do the following:

To configure the agent in a Storage Foundation for Oracle RAC or SFCFS environment

- 1 Configure the SupportedActions attribute for the CVMVolDg resource.
- 2 Add the following keys to the list: import, deport, and vxdctlenable.
- 3 Run the following commands sequentially to add the entry points to the CVMVolDg resource:

```
# haconf -makerw

# hatype -modify CVMVolDg SupportedActions import deport
vxdctlenable

# haconf -dump -makero
```

Note that SupportedActions is a resource type attribute, and it defines a list of action tokens for the resource.

Managing and testing clustering support for IBM SVC Copy Services

This chapter includes the following topics:

- [Typical test setup for the IBM SVC Copy Services agent](#)
- [Testing service group migration](#)
- [Testing host failure](#)
- [Testing cluster failure](#)
- [Performing a disaster test](#)
- [Performing the failback test](#)
- [Failure scenarios for IBM SVC Copy Services](#)

Typical test setup for the IBM SVC Copy Services agent

A typical test environment includes the following characteristics:

- A master SVC cluster is attached to the storage at the back end and the application hosts at the front end.
- An auxiliary SVC cluster is attached to the storage at the back end and the application hosts at the front end.
- Two hosts (hosta and hostb) are attached to the master SVC cluster.


- Two hosts (hostc and hostd) are attached to the auxiliary SVC cluster.
- SSH configuration is established on hosts (hosta and hostb) and the master SVC cluster so that the hosts (hosta and hostb) can communicate via SSH with the master SVC cluster.
- SSH configuration on hosts (hostc and hostd) and the auxiliary SVC cluster so that the hosts (hostc and hostd) can communicate via SSH with the auxiliary SVC cluster.
- The application runs on hosta which is connected to the master SVC cluster, also the primary for the replication relationship being managed by the IBM SVC Copy Services agent.

To verify that the host can communicate with the SVC cluster via SSH, run the following command from the host:

```
ssh -l<username> -i <path_to_ssh_id_file_on_host> <SVC ClusterName  
or IP address> <SVC command>
```

Figure 4-1 depicts a typical test environment.

Figure 4-1 Typical test setup



Testing service group migration

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

To perform the service group migration test

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

Migrate the service group to a host (hostb) that is attached to the primary SVC cluster.
- 2** In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

Migrate the service group to a host (hostc) that is attached to a different SVC cluster (auxiliary SVC cluster).
- 3** Click **Switch To**, and click the system (hostc) that is connected to the auxiliary SVC cluster.

The service group comes online on hostc. The state of the replication stays online but the auxiliary SVC cluster now becomes the primary.
- 4** In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

Migrate the service group back to its original host.
- 5** Click **Switch To** and click the system on which the group was initially online (hosta).

The service group comes online on hosta that is attached to the master SVC cluster. The state of the replication is online and the master SVC cluster now becomes the primary again.

Note: If the replication stops or the replication link is disconnected, set the DisconnectTakeover attribute to 1 for a successful failback between the primary and secondary.

Testing host failure

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

To perform the host failure test

- 1** Halt or shut down the host where the application runs (hosta).
 The application fails over to hostb, also located in the primary site. The IBM SVC Copy Services agent performs no operation on the online entry point as hostb is attached to the primary SVC cluster.
- 2** Halt or shut down hostb.
 In a replicated data cluster, the group fails over to hostc or hostd depending on the FailOverPolicy attribute in the cluster.
 The online entry point of the IBM SVC Copy Services agent attempts to switch roles of replication or at least enable read write access to the storage on the secondary so that the application comes up successfully on the secondary site.
 In a global cluster, a cluster down alert appears and gives you the opportunity to fail over the service group manually.

Testing cluster failure

In this scenario, the primary cluster fails i.e. the nodes in the cluster are shut down.

To perform the cluster failure test

- 1** Halt or shut down the master i.e. primary, where the application runs (hosta).
- 2** If you cannot halt hosta, disable all the ports on the switch or switches connected to the SVC cluster.
 Wait for a few minutes for the VCS agents to detect the loss of storage connectivity and for VCS to initiate the service group failover.
 The service group fails to come online in the primary site.
- 3** Set the attribute DisconnectTakeOver = 1 to bring the service group online on the secondary site.
 In a replicated data cluster, the service group automatically fails over to the secondary site (hostc or hostd) depending upon the FailOverPolicy in the cluster.
 In a global cluster, a cluster down alert appears and gives you the opportunity to fail over the service group manually.

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, the SVC cluster, and the storage arrays connected to the backend of the SVC cluster.

If you cannot shut down the storage arrays or the SVC cluster, disconnect the replication link between the SVC clusters on the two sites. The action mimics a disaster scenario to the secondary site.

- 2 VCS fails over to the hosts on the secondary site (hostc or hostd).

Performing the failback test

You can set up your cluster for a failback test.

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

To perform a failback test

- 1 Reconnect the replication link and reboot the hosts on the primary site.
- 2 Take the service group offline.
- 3 Manually resynchronize the data between the primary and secondary sites.

If the current primary is retained as the primary for the replication relationship and the replication is not in a stopped state, stop the replication manually.

- Use the `stopprrelationship` or the `stopprconsistgrp` command, without the `-access` option.
- Use the `starttrrelationship` or the `starttrconsistgrp` command with the `-force` option on the SVC cluster that the user wants to retain as primary.
- Alternatively, use the update action entry point with no arguments specified.

If you change the primary for the replication relationship, stop the replication manually (if it is not stopped already).

- Enable read / write access on both SVC clusters by using the `stopprrelationship` or the `stopprconsistgrp` command with the `-access` option.
- Use the `starttrrelationship` or the `starttrconsistgrp` command with both the `-force` and `-primary` options.

- Alternatively, use the update action entry point with exactly one argument that is the SVC cluster, which the user retains as primary.
- 4 After the resynchronization is complete, migrate the application back to the original primary side.

Failure scenarios for IBM SVC Copy Services

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

Site disaster

In a total site failure, all hosts, the SVC cluster, and storage arrays are completely disabled, either temporarily or permanently.

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

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

All host or all application failure

If all hosts on the primary side are disabled or if the application cannot start successfully on any primary host, the service group fails over.

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

In replicated data cluster environments, site failure is detected the same way as a total host failure, that is, the loss of all LLT heartbeats. This type of failure is communicated by the VCS engine to the other site.

In a global cluster environment, VCS detects the failure by the loss of the ICMP heartbeat between the clusters. This type of failure is communicated by the wac process to the other site.

Cluster failure

If all hosts on the primary side are disabled or if the application cannot start successfully on any primary host, the replication relationship transitions to a disconnected state.

Any VCS agent, DiskGroup, Volume or Mount that monitors the storage infrastructure detects the loss of storage connectivity.

When VCS detects a fault in one of the storage resources (DiskGroup, Volume or Mount), the service group is taken offline and declared faulted on hosta.

VCS initiates a failover to hostb, which is also in the primary site and connected to the same SVC cluster as hosta.

The service group fails to come online on hostb since the SVC cluster is shut down.

The service group is brought online on the secondary site (either on hostc or hostd), only if the DisconnectTakeover attribute is set to 1.

If the attribute DisconnectTakeover = 1, there is a possibility of data loss since the replication is in a disconnected state and indicates that the data on the primary and secondary sites may not be in sync.

Warning: Veritas recommends caution when setting the DisconnectTakeover to 1 or StopTakeover to 1 for the replication resource.

Replication link failure

The IBM SVC Copy Services agent does not monitor the replication link status and cannot detect link failures. When the replication link between the primary and secondary cluster fails, VCS takes no action. If the replication relationship does not stop after the link failure, the two SVC clusters are automatically resynchronized after the replication link is reconnected.

Split-brain in a SVC Copy Services environment

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

The user must resynchronize the data between the two clusters manually.

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

Setting up a fire drill

This chapter includes the following topics:

- [About fire drills](#)
- [Fire drill configurations](#)
- [About the SVCCopyServicesSnap 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 IBM SVC Copy Services, the SVCCopyServicesSnap 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.

Fire drill configurations

VCS supports the following fire drill configurations for the agent:

Gold	<p>Runs the fire drill on a snapshot of the target array. The replicated device keeps receiving writes from the primary.</p> <p>Veritas recommends this configuration because it does not affect production recovery.</p> <p>In the Gold configuration, VCS does the following:</p> <ul style="list-style-type: none">■ Takes a snapshot of the target array on snapshot LUNs using FlashCopy.■ Modifies the disk group name in the snapshot.■ Brings the fire drill service group online using the snapshot data.
Silver	<p>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.</p> <p>In the Silver configuration, VCS does the following:</p> <ul style="list-style-type: none">■ Suspends replication.■ Takes a snapshot of the target array on snapshot LUNs using FlashCopy.■ Modifies the disk group name in the replicated target device.■ 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. <p>Note: After running the silver fire drill for SVCCopyServices at the disaster recovery site, the replication target is resynchronized from the primary. After the resynchronization is complete, you must discard the snapshot to be able to bring the application service group online at the recovery site. If disaster happens before the resynchronization is complete, you must restore the replication target from the snapshot.</p>
Bronze	<p>VCS breaks replication and runs the fire drill test on the replicated target. VCS does not take a snapshot in this configuration.</p> <p>If a disaster occurs while resynchronizing data after the test, it may result in inconsistent data as there is no snapshot data.</p> <p>In the Bronze configuration, VCS does the following:</p> <ul style="list-style-type: none">■ Suspends replication.■ Modifies the disk group name while importing.

About the SVCCopyServicesSnap agent

The SVCCopyServicesSnap agent is the fire drill agent for IBM SVC Copy Services.

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

The SVCCopyServicesSnap agent supports fire drill for storage devices that are managed using Veritas Volume Manager (VxVM). Additionally, on AIX, the SVCCopyServicesSnap agent supports LVM for Gold fire drill configurations. See [“Configuring the agent for LVM support”](#) on page 37.

Configuring the agent for LVM support

To configure the agent for LVM support:

- 1 Run the following command to set the Physical Volume Identifier (PVID) for all the FlashCopy disks:

```
chdev -l <hdisk#> -a pv=yes
```

- 2 Run the `lspv` command to ensure that the output shows the PVID for the FlashCopy disks on all nodes of the cluster.

- 3 Create a Volume Group comprising of FlashCopy disks.

- 4 Vary on the Volume Group on any one node.

- 5 Run the following command:

```
hares -action <SVCCopyServicesSnap resource name> getfcdisks  
-actionargs <FlashCopy Volume Group name> -sys  
<VCS system where the Volume Group is active>
```

- 6 The `getfcdisks` action function exports the FlashCopy Volume Group. The SVCCopyServicesSnap fire drill agent recreates the Volume Group during fire drill.

SVCCopyServicesSnap agent functions

The SVCCopyServicesSnap agent performs the following functions:

Table 5-1 Agent functions

Function	Description
online	<ul style="list-style-type: none"> ■ For Gold and Silver configurations, takes a local snapshot of the target LUN. ■ For Gold configuration,takes the fire drill service group online by mounting the snapshot LUN. ■ For Bronze and Silver configurations, takes the fire drill service group online by mounting the target LUN. ■ For Bronze and Silver configurations, suspends replication between the source and the target arrays. ■ Creates a lock file to indicate that the resource is online.
offline	<ul style="list-style-type: none"> ■ Resumes the replication between the source and the target arrays. ■ Removes the lock file created by the online function.
monitor	Verifies the existence of the lock file to make sure the resource is online.
clean	<ul style="list-style-type: none"> ■ Resumes the replication between the source and the target arrays. ■ Removes the lock file created by the online function.
action/getfcdisks	<p>Gets a list of target FlashCopy disks based on the Volume Group that is imported with FlashCopy target disks. This action entry point exports the Volume Group after successfully getting the (physical volume identifier (PVid) and disks required for fire drill.</p> <p>For more information on the role the getfcdisks action entry point plays when configuring the agent to support fire drill for LVM: See “Configuring the agent for LVM support” on page 37.</p>

Resource type definition for the SVCCopyServicesSnap agent

Following is the resource type definition for the SVCCopyServicesSnap agent:

```
type SVCCopyServicesSnap (
    static keylist SupportedActions = { getfcdisks }
    static int MonitorInterval = 300
    static int OpenTimeout = 180
    static int NumThreads = 1
    static int OfflineMonitorInterval = 0
```

```

static int OnlineTimeout = 6000
static int RestartLimit = 1
static str ArgList[] = { TargetResName, MountSnapshot,
UseSnapshot, RequireSnapshot, FCMapGroupName,
IsSpaceEfficientSnapshot }
str TargetResName
int MountSnapshot
int UseSnapshot
int RequireSnapshot
str FCMapGroupName
int IsSpaceEfficientSnapshot = 0
temp str Responsibility
temp str FDFile
)

```

Attribute definitions for the SVCCopyServicesSnap agent

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

Table 5-2 Agent attributes

Attribute	Description
TargetResName	<p>Name of the resource managing the LUNs that you want to take snapshot of. Set this attribute to the name of the SVCCopyServices resource if you want to take a snapshot of replicated data. Set this attribute to the name of the DiskGroup resource if the data is not replicated.</p> <p>For example, in a typical 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.</p> <p>Type-Dimension: string-scalar</p>
UseSnapshot	<p>Specifies whether the SVCCopyServicesSnap resource takes a local snapshot of the target array. Set this attribute to 1.</p> <p>Type-Dimension: integer-scalar</p> <p>See “About the Snapshot attributes” on page 41.</p>

Table 5-2 Agent attributes (*continued*)

Attribute	Description
RequireSnapshot	<p>Specifies whether the SVCCopyServicesSnap resource must take a snapshot before coming online.</p> <p>Set this attribute to 1 if you want the resource to come online only after it succeeds in taking a snapshot.</p> <p>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if UseSnapshot is set to 1.</p>
MountSnapshot	<p>Specifies whether the resource uses the snapshot to bring the service group online. Set this attribute to 1.</p> <p>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if the UseSnapshot attribute is set to 1.</p>
FCMapGroupName	<p>Name of the FlashCopy mapping or FlashCopy consistency group.</p> <p>If the target SVCCopyServices resource contains a consistency group, set FCMapGroupName to a flashcopy consistency group. If the target SVCCopyServices resource contains a relationship, set FCMapGroupName to a flashcopy mapping. This attribute is optional for Bronze configurations.</p> <p>Type-Dimension: string-scalar</p>
IsSpaceEfficientSnapshot	<p>Specifies whether the SVCCopyServices resource uses space-efficient snapshots.</p> <p>If you want space-efficient snapshots to be used for fire drill operations, set this attribute to 1.</p> <p>The SVCCopyServices resource can use space-efficient snapshots only if the snapshot target is thin-provisioned and copy rate is 0.</p> <p>Note: This attribute is applicable only to storage devices that are managed using Veritas Volume Manager (VxVM).</p> <p>Type-Dimension: integer-scalar</p>

About the Snapshot attributes

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

[Table 5-3](#) lists the snapshot attribute values for fire drill configurations:

Table 5-3 Snapshot attribute values for fire drill configurations

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

Setting the RequireSnapshot attribute to 0 enables a Gold or Silver configuration to run in the Bronze mode if the snapshot operation fails.

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 SVCCopyServices resource.
- Make sure the infrastructure to take snapshots is properly configured between the source and target arrays. This process involves associating FlashCopy mapping or consistency group and quiescing the application.
- If you plan to run a fire drill on SVC Copy Services devices, you must have a FlashCopy license.
- When you use the Gold or Silver configuration, make sure FlashCopy for IBM SVC is installed and configured at the target array.

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 45.

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 42.

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 **SVCCopyServices** resource and click **Delete**.
- 4 Add a resource of type **SVCCopyServicesSnap** 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.

Enabling the FireDrill attribute

You must edit certain resource types so they are FireDrill-enabled. Making a resource type FireDrill-enabled changes the way that VCS checks for concurrency violations. Typically, when FireDrill is not enabled, resources cannot come online

on more than one node in a cluster at a time. This behavior prevents multiple nodes from using a single resource or from answering client requests. Fire drill service groups do not interact with outside clients or with other instances of resources. They can safely come online even when the application service group is online.

Typically, you would enable the FireDrill attribute for the resource type that is used to configure the agent. For example, in a service group monitoring Oracle, enable the FireDrill attribute for the Oracle resource type.

To enable the FireDrill attribute

- 1 In Cluster Explorer, click the **Types** tab in the left pane, right-click the type to be edited, and click **View > Properties View**.
- 2 Click **Show All Attributes**.
- 3 Double-click **FireDrill**.
- 4 In the **Edit Attribute** dialog box, enable **FireDrill** as required, and click **OK**.
- 5 Repeat the process of enabling the FireDrill attribute for all required resource types.

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 SVCCopyServicesSnap resource replaces the SVCCopyServices resource.

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

```
SVCCopyServicesSnap oradg_fd (  
    TargetResName = oradg_cg  
    MountSnapshot = 1  
    UseSnapshot = 1  
    RequireSnapshot = 1  
    IsSpaceEfficientSnapshot = 1  
    FCMapGroupName = fc_cg1  
)
```