

Veritas™ Cluster Server One Agent for IBM SVCCopyServices Installation and Configuration Guide

AIX, HP-UX, Linux, Solaris, ESX,
Windows Server 2003/2008

5.1 Service Pack 1

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: 5.0SP1.0

Document version: 5.0SP1.0.0

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas, Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's maintenance offerings include the following:

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

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

www.symantec.com/business/support/index.jsp

Contacting Technical Support

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

www.symantec.com/business/support/contact_techsupp_static.jsp

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

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

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

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

Licensing and registration

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

customercare.symantec.com

Customer service

Customer Care information is available at the following URL:

www.symantec.com/customercare

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

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

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

sfha_docs@symantec.com

Maintenance agreement resources

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

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

Additional enterprise services

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

Enterprise services that are available include the following:

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

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

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	3
Chapter 1	Introducing the Veritas Cluster Server One Agent for IBM SVCCopyServices
	9
	About the agent for IBM SVCCopyServices
	9
	What's new in this release
	10
	Supported software for IBM SVCCopyServices
	10
	Supported hardware for IBM SVCCopyServices
	10
	Typical IBM SVCCopyServices setup in a VCS One cluster
	11
	SVCCopyServices agent functions
	12
	About the SVC Copy Services agent's online function
	13
	About the SVC Copy Services agent's update action function
	16
	About the SVC Copy Services agent's info function
	17
Chapter 2	Installing and removing the agent for IBM SVCCopyServices
	19
	Installing the Veritas Cluster Server One agent for SVCCopyServices
	19
	Installing the agent packages using the installer
	20
	Installing the agent package using the CLI
	20
	Adding the agent resource type definitions
	21
	Removing the Veritas Cluster Server One agent for SVCCopyServices
	23
	Removing the agent packages using the installer
	23
	Removing the agent package using CLI
	24
	Removing the agent type definition from the Policy Master system
	25
Chapter 3	Configuring the agent for IBM SVCCopyServices
	27
	Configuration concepts for the SVCCopy Services agent
	27
	Resource type definition for the IBM SVCCopyServices agent
	27

	Attribute definitions for the SVCCopyServices agent	28
	Sample configuration for the SVCCopyServices agent	30
	Before you configure the agent for IBM SVCCopyServices	31
	Configuring the agent for IBM SVCCopyServices	31
Chapter 4	Managing and testing clustering support for IBM SVCCopyServices	33
	Failure scenarios in global clusters	33
	Testing the global composite service group migration	36
	Testing disaster recovery after site failure	36
	Testing disaster recovery after client system failure	37
	Performing failback after a client system failure or an application failure	38
	Performing failback after a site failure	38
Index		41

Introducing the Veritas Cluster Server One Agent for IBM SVCCopyServices

This chapter includes the following topics:

- [About the agent for IBM SVCCopyServices](#)
- [What's new in this release](#)
- [Supported software for IBM SVCCopyServices](#)
- [Supported hardware for IBM SVCCopyServices](#)
- [Typical IBM SVCCopyServices setup in a VCS One cluster](#)
- [SVCCopyServices agent functions](#)

About the agent for IBM SVCCopyServices

The VCS One enterprise agent for IBM SVCCopyServices 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 SVCCopyServices 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 SVCCopyServices agent supports MetroMirror (i.e. synchronous replication) and Global Mirror (i.e. asynchronous replication).

What's new in this release

The VCSOne agent for IBM SVCCopyServices includes the following new or enhanced features:

- The agent for SVCCopyServices is supported on ESX 3.5/4.0/4i.
- The agent for SVCCopyServices provides support for Windows Server 2003 and 2008.

Supported software for IBM SVCCopyServices

The agent for IBM SVCCopyServices supports the following software versions:

- Veritas Cluster Server One
- VCS One 5.0 SP1 on AIX
 - VCS One 5.0 SP1 on HP-UX 11i v2
 - VCS One 5.0 SP1 on Red Hat Enterprise Linux
 - VCS One 5.0 SP1 on SUSE Linux Enterprise Server
 - VCS One 5.0 SP1 on Solaris SPARC
 - VCS One 5.0 SP1 on Solaris x64 (OPTERON)
 - VCS One 5.0 SP1 on ESX 3.5/4.0/4i
 - VCS One 5.0 SP1 on Windows Server 2003, 2008.

See the product's Release Notes for more details on the supported architectures and the operating system versions.

Firmware levels for SVC 4.1.1.2 or higher

Supported hardware for IBM SVCCopyServices


The agent also supports SSH access to SVC.

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 SVCCopyServices setup in a VCS One cluster

Figure 1-1 displays a typical setup in a SVCCopyServices environment.

Figure 1-1 Typical clustering setup for the agent



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

- 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 SVCCopyServices 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.

SVCCopyServices agent functions

The VCS One agent for IBM SVCCopyServices 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 SVC Copy Services agent’s online function” on page 13.</p>
offline	<p>The agent removes the lock file that was created for the resource by the online 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 SVCCopyServices 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"> ■ starttrcrelationship ■ starttrcconsistgrp <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 SVC Copy Services agent’s update action function” on page 16.</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 VCS One was stopped by force and then restarted on the client system.</p>
info	<p>Reports the detailed status of the replication relationship or consistency group being monitored by the agent.</p> <p>See “About the SVC Copy Services agent’s info function” on page 17.</p>

About the SVC Copy Services 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.

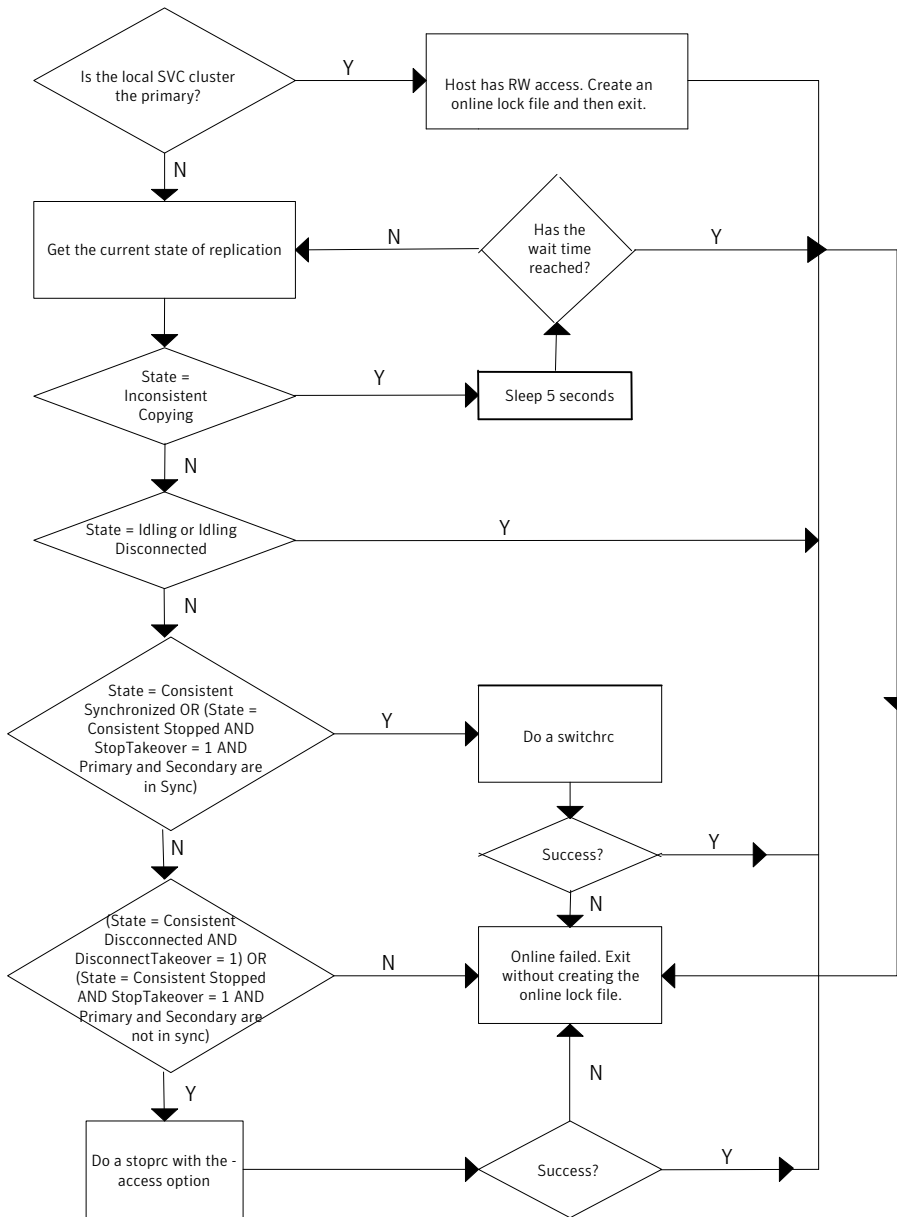
- **Consistent Disconnected** - The online entry point takes action only if the attribute `DisconnectTakeover = 1`. If `DisConnectTakeover = 1`, the online entry point runs the `stoprc` command with the `-access` option.

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.

See [Figure 1-2](#) on page 15.

Figure 1-2 The algorithm for the online entry point



About the SVC Copy Services agent's update action function

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

Warning: Symantec 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 composite service group, `csg_oradb` and the replication resource in this global composite service group is defined by the SVCCopyServices 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 composite service group `csg_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 -action_args master -sys hosta.`
 - **Consistent Stopped / Consistent Disconnected** - 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 composite service group, `csg_oradb` and the replication resource in this global composite service group is defined by the SVCCopyServices resource `svc_mmrel`. If the replication relationship managed by the resource `svc_mmrel` is in the `consistent_stopped` or `consistent_disconnected` state:

- Ensure that the application global composite service group `csg_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.

Symantec 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 SVC Copy Services agent's info function

The info entry point captures the entire output from `svcinfolsrcrelationship` or the `svcinfolsrcconsistgrp` for the replication relationship or consistency group monitored by the SVCCopyServices resource.

The output updates as the value of the ReplicationStatus key in the ResourceInfo attribute for the SVCCopyServices 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 SVCCopyServices agent. The agent always allows the action and info entry points to run.

Installing and removing the agent for IBM SVCCopyServices

This chapter includes the following topics:

- [Installing the Veritas Cluster Server One agent for SVCCopyServices](#)
- [Removing the Veritas Cluster Server One agent for SVCCopyServices](#)

Installing the Veritas Cluster Server One agent for SVCCopyServices

You must install the agent for SVCCopyServices on all the client systems of the VCS One cluster that will host the SVCCopyServices service group. You can install the agent for SVCCopyServices using the installagpack program or using the command line interface (CLI).

The installation of the agent packs involves the following phases:

Installing the agent packages

See [“Installing the agent packages using the installer”](#) on page 20.

Adding the agent resource type definitions

See [“Adding the agent resource type definitions”](#) on page 21.

Note: The installagpack program supports only the -addtypes, -rmtypes, -responsefile, and -rsh options. Symantec recommends that you do not use any of the other options from the `installagpack` command help output.

Installing the agent packages using the installer

You can install the agent packages on one or more client systems of a specific platform type.

Perform the following steps to install the agent packages using the installer

- 1 Mount the VCS One Agent Pack software disc on the client system where you plan to run the installation.
- 2 Depending on the platform type, navigate to the directory containing the agent installer:

AIX `cdl/aix/vcsone/vcsone_version`

HP-UX `cdl/hpux/hpuxos_version/vcsone/vcsone_version`

Where *os_version* is the HP-UX version.

Linux `cdl/linux/dist_arch/vcsone/vcsone_version`

Where *dist* is the Linux distribution and *arch* is the architecture.

Solaris `cdl/solaris/dist_arch/vcsone/vcsone_version`

Where, *dist_arch* is 'sol_sparc' or 'sol_x64'.

- 3 Enter the following command to start the agent pack installation:

```
# ./installagpack [-rsh]
```

You can use the `-rsh` option if `rsh` and `rcp` are used for communication between systems instead of the default `ssh` and `scp`. This option requires that systems be preconfigured such that the `rsh` commands between systems execute without prompting for passwords or confirmations.

- 4 Enter the name of the client systems where you want to install the agents.
- 5 Choose whether to install all the agents or any specific agent. Follow the installer prompt to specify your option.
- 6 Review the output as the installation program installs the agent packages.

You can view installation logs in the `/var/VRTS/install/logs` directory.

Installing the agent package using the CLI

You can install the desired agent package using the CLI, on one or more client systems of a specific platform type.

Perform the following steps to install the agent packages using CLI

- 1 Mount the VCS One Agent Pack software disc on the client system where you plan to run the installation.
- 2 Depending on the platform type, navigate to the directory containing the agent installer:

AIX # `cd1/aix/vcsone/vcsone_version/pkgs`

HP-UX # `cd1/hpux/hpuxos_version/vcsone/vcsone_version/depot`

Linux # `cd1/linux/dist_arch/vcsone/vcsone_version/rpms`

Where, *dist* is the Linux distribution and *arch* is the architecture

Solaris # `cd1/solaris/dist_arch/vcsone/vcsone_version/pkgs`

Where *dist_arch* is 'sol_sparc' or 'sol_x64'

- 3 Type the following command on each client system to install the agent. Answer the prompt accordingly:

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

HP-UX # `swinstall -s `pwd` VRTSvcssvc`

Linux # `rpm -ivh VRTSvcssvc_rpm_filename`

Solaris # `pkgadd -d . VRTSvcssvc`

Adding the agent resource type definitions

You must add the agent resource type definitions to the Policy Master database configuration. You can perform this task from any client system in the VCS One cluster.

Note: You must add the agent resource type definitions only one time per platform type.

To add the agent resource types to the policy master database configuration

- 1** Set up RSH or SSH communications between the client system and the policy master system.

For information on configuring SSH for remote communication, refer to the *Veritas Cluster Server One Installation Guide*.

- 2** Make sure that the PM daemon is running.

```
# /opt/VRTSvcsone/bin/haclus -display
```

The output should show ClusterState is RUNNING.

- 3** If you have just installed the agents on VCS One client systems and still have the VCS One Agent Pack software disc mounted, skip to step [6](#).
- 4** Mount the VCS One Agent Pack software disc.
- 5** Depending on the platform type, navigate to the directory containing the agent installer:

AIX `cdl/aix/vcsone/vcsone_version`

HP-UX `cdl/hpux/hpuxos_version/vcsone/vcsone_version`

Where `os_version` is the HP-UX version.

Linux `cdl/linux/dist_arch/vcsone/vcsone_version`

Where `dist` is the Linux distribution and `arch` is the architecture.

Solaris `cdl/solaris/dist_arch/vcsone/vcsone_version`

Where `dist_arch` is the `sol_sparc` or `sol_x64`.

- 6** Enter the command to start the agent pack installer for adding resource types to the Policy Master configuration database. Use the `-addtypes` option:

```
# ./installagpack -addtypes
```

- 7** When the installer prompts, enter the virtual IP address of the Policy Master.
- 8** Review the output as the installer verifies communication with the Policy Master system.

- 9 Choose whether to add the type definitions for all the agents or for specific agents. Follow the installer prompts to add the type definitions.
- 10 Review the output as the installer adds the agent types to the PM database configuration and copies the appropriate types.xml files to the PM system.
You can view installation logs in the /var/VRTS/install/logs directory.

Removing the Veritas Cluster Server One agent for SVCCopyServices

Removing the agent package involves removing the agent files from each client system where it was installed.

You can remove the packages using the agent pack installer or the command line.

See [“Removing the agent packages using the installer”](#) on page 23.

See [“Removing the agent package using CLI”](#) on page 24.

After removing the agent packages you can remove the agent type definition from the Policy Master system.

See [“Removing the agent type definition from the Policy Master system”](#) on page 25.

Removing the agent packages using the installer

You can remove all the agent packages or the desired agent package using the `uninstallagpack` program.

Note: The `uninstallagpack` program supports only the `-responsefile` and `-rsh` options. Symantec recommends that you do not use any of the other options from the `uninstallagpack` command help output.

To remove the agent packages from the client systems

- 1 Freeze the service groups that hosts the application, on the system from which you want to remove the agent package.

```
# hagr -freeze <groupname>
```

- 2 Stop the agent on all client systems before you remove the agent package from the system.

```
# haagent -stop -notransition <AgentName> -sys <system_name>
```

- 3 Ensure that the agent operations are stopped on all the cluster systems.

```
# haagent -display <AgentName>
```

- 4 Mount the VCS One Agent Pack software disc on the client system where you plan to run the uninstallagpack program.
- 5 Depending on the platform type, navigate to the directory containing the agent uninstaller:

AIX `cd1/aix/vcsone/vcsone_version`

HP-UX `cd1/hpux/hpuxos_version/vcsone/vcsone_version`

Where `os_version` is the HP-UX version.

Linux `cd1/linux/dist_arch/vcsone/vcsone_version`

Where `dist` is the Linux distribution and `arch` is the architecture.

Solaris `cd1/solaris/dist_arch/vcsone/vcsone_version`

Where `dist_arch` is the `sol_sparc` or `sol_x64`.

- 6 Start the uninstallagpack program.

```
# ./uninstallagpack [-rsh]
```

- 7 Enter the name of the client systems on which you want to uninstall the agent pack. The names must be separated by spaces.
- 8 Choose whether to remove all the agent packages or a specific agent package. Follow the installer prompt to remove the agent package.
- 9 Review the output as the program verifies the agent pack that you installed and removes the agent packages.

You can view logs in the `/var/VRTS/install/logs` directory.

Removing the agent package using CLI

You can remove a desired agent package using the CLI.

Note: You must remove this agent package from each client system in the cluster.

To remove the agent for SVCCopyServices from a client system

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

AIX # **installp -u VRTSvcssvc**

HP-UX # **swremove VRTSvcssvc**

Linux # **rpm -e VRTSvcssvc**

Solaris # **pkgrm VRTSvcssvc**

Removing the agent type definition from the Policy Master system

After you remove the agent packages, you can remove the agent type definitions for all the agents for specific agents from the Policy Master system.

To remove the agent type definition from the Policy Master system

- 1 Navigate to the following directory on the client system.

```
# cd /opt/VRTS/install
```

- 2 Run the following command to remove the agent type definition from the Policy Master system:

```
# ./installagpack -rmtypes
```

- 3 When the installer prompts, enter the virtual IP address of the Policy Master.
- 4 Choose whether to remove the type definitions for all the agents or for specific agents. Follow the installer prompts to remove the type definitions.

You can view logs in the /var/VRTS/install/logs directory.

Configuring the agent for IBM SVCCopyServices

This chapter includes the following topics:

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

Configuration concepts for the SVCCopy Services agent

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

Resource type definition for the IBM SVCCopyServices agent

The resource type represents the VCS One configuration of the agent and specifies how the agent is defined in the configuration file `main.xml`. For more information, refer to the sample `SVCCopyServicesTypes.platform.xml` files in the `/etc/VRTSagents/ha/conf/SVCCopyServices` directory on the primary client system.

Attribute	Default value
SSHBinary	<code>/usr/bin/ssh</code>
SSHPATHToIDFile	<No default value>
GroupName	<No default value>

Attribute	Default value
IsConsistencyGroup	1
SVCClusterIP	<No default value>
SVCUserName	admin
DisconnectTakeover	0
StopTakeover	0

See “[Attribute definitions for the SVCCopyServices agent](#)” on page 28.

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>
IsConsistencyGroup	<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>Default is: "C:\Program files (x86)\PuTTY\Plink.exe"</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 of 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>

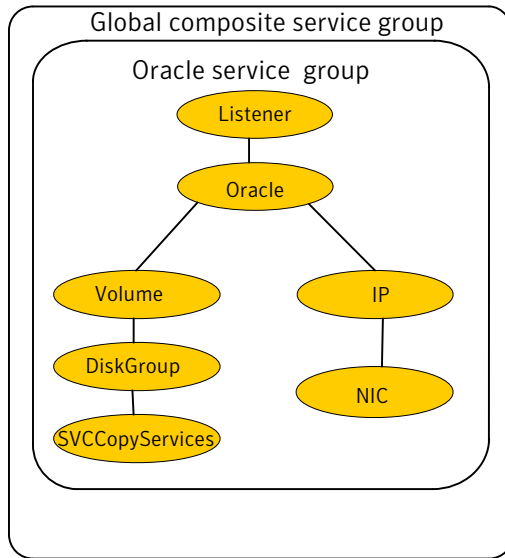
SVCUserName	<p>User name that authenticates the SSH connection with the SVC cluster.</p> <p>Default is admin.</p> <p>Type-dimension: string-scalar</p>
StopTakeover	<p>Determines whether the agent makes read-write access available to the host when the replication is in a stopped state (i.e. <code>consistent_stopped</code>).</p> <p>The status of the replication goes into a stopped state when the user fires the <code>stopprrelationship</code> or the <code>stopprconsistentgrp</code> command. Thus, no replication occurs between the primary and secondary SVC clusters.</p> <p>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.</p> <p>The possible stopped states are:</p> <p><code>inconsistent_stopped</code></p> <p><code>consistent_stopped</code></p> <p>When the state of the replication is <code>consistent_stopped</code> and <code>StopTakeover = 1</code>, the agent enables read-write access for the SVC cluster.</p> <p>When the state of the replication is <code>inconsistent_stopped</code>, the agent does not enable read-write access for the SVC cluster.</p> <p>Type-dimension: integer-scalar</p>

Disconnect Takeover	<p>Determines whether the agent makes read-write access available to the host when the replication is in a disconnected state (i.e. <code>consistent_disconnected</code>).</p> <p>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.</p> <p>Attribute value is either 0 or 1. Default is 0.</p> <p>The possible disconnected states are:</p> <p><code>idling_disconnected</code></p> <p><code>inconsistent_disconnected</code></p> <p><code>consistent_disconnected</code></p> <p>When the state of the replication is <code>consistent_disconnected</code> and <code>DisconnectTakeover = 1</code>, then the agent enables read/write access for the SVC cluster. When the state of the replication is <code>idling_disconnected</code>, the agent does not enable read-write access for the SVC cluster.</p> <p>Type-dimension: integer-scalar</p>
------------------------	--

Sample configuration for the SVCCopyServices agent

Figure 3-1 shows a dependency graph of a VCS One global composite service group that has a resource of type SVCCopyServices.

Figure 3-1 VCS One global composite service group with resource type SVCCopyServices



Before you configure the agent for IBM SVCCopyServices

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.
- Review the configuration concepts, which describe the agent's type definition and attributes.
 See [“Configuration concepts for the SVCCopy Services agent”](#) on page 27.
- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
 See [“Typical IBM SVCCopyServices setup in a VCS One cluster”](#) on page 11.

Configuring the agent for IBM SVCCopyServices

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

- Converting their devices to SVCCopyServices devices

- Synchronizing the devices
- Adding the IBM SVCCopyServices agent to the service group

Note: You must not change the replication state of a cluster from primary to secondary and viceversa, outside of a VCS One setup. The agent for IBM SVCCopyServices fails to detect a change in the replication state if the role reversal is done externally.

Managing and testing clustering support for IBM SVC Copy Services

This chapter includes the following topics:

- [Failure scenarios in global clusters](#)
- [Testing the global composite service group migration](#)
- [Testing disaster recovery after site failure](#)
- [Testing disaster recovery after client system failure](#)
- [Performing failback after a client system failure or an application failure](#)
- [Performing failback after a site failure](#)

Failure scenarios in global clusters

[Table 4-1](#) lists the failure scenarios in a global cluster configuration and describes the behavior of VCS One and the agent in response to the failure.

See the *Veritas Cluster Server One User's Guide* for more information on the DR configurations and the global composite service group attributes.

Table 4-1 Failure scenarios in a global cluster configuration with VCS One agent for IBM SVCCopyServices

Failure	Description and VCS One response
Application failure	<p>Application cannot start successfully on any client system at the primary site.</p> <p>VCS One response at the secondary site:</p> <ul style="list-style-type: none"> ■ Causes global composite service group at the primary site to fault and triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ The state of the replication is Consistent Synchronized and the SVC cluster at the secondary site becomes the primary SVC cluster. <p>See “Performing failback after a client system failure or an application failure” on page 38.</p>
Client system failure	<p>All client systems at the primary site fail.</p> <p>VCS One response at the secondary site:</p> <ul style="list-style-type: none"> ■ Triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ The state of the replication is Consistent Synchronized and the SVC cluster at the secondary site becomes the primary SVC cluster.
Site failure	<p>All PMs, client systems, and their storage at the primary site fail.</p> <p>VCS One response at the secondary site:</p> <ul style="list-style-type: none"> ■ Triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ If the state of the replication is Consistent Disconnected and the value of the DisconnectTakeover attribute is 1, the agent runs the <code>stoprc</code> command with the <code>-access</code> option. This makes the SVC devices write-enabled. If the state of the replication is Idling Disconnected, then the devices are already write-enabled. <p>See “Performing failback after a site failure” on page 38.</p>

Table 4-1 Failure scenarios in a global cluster configuration with VCS One agent for IBM SVCCopyServices (*continued*)

Failure	Description and VCS One response
Replication link failure	<p>Replication link between the arrays at the two sites fails.</p> <p>VCS One response: No action.</p> <p>Agent response: No action.</p>
Network failure	<p>The network connectivity and the replication link between the sites fail.</p> <p>VCS One response:</p> <ul style="list-style-type: none"> ■ VCS One at each site concludes that the remote cluster has faulted. ■ The global cluster failover in VCS One is manual. No action. ■ You must confirm the cause of network failure from the cluster administrator at each site and fix the issue. ■ If the availability of the application is unaffected by the network failure, you can leave the global composite service group as is. <p>Agent response: No action.</p> <ul style="list-style-type: none"> ■ If the availability of the application is affected, you must takeover the global composite service group from the remote site. <p>Agent response on the secondary site: Makes the devices write-enabled and the replication state Idling Disconnected.</p> <p>After the network failure is fixed, you can perform a failback.</p> <p>See “Performing failback after a site failure” on page 38.</p>
Storage failure	<p>The array at the primary site fails.</p> <p>A storage failure at the primary site causes the state of the replication on the secondary site to go to a disconnected state.</p> <p>VCS One response at the secondary site:</p> <ul style="list-style-type: none"> ■ Causes the global composite service group at the primary site to fault and triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ If the state of the replication is Consistent Disconnected and the value of the DisconnectTakeover attribute is 1, the agent runs the <code>stoprc</code> command with the <code>-access</code> option. This makes the SVC devices write - enabled.

Testing the global composite service group migration

After you configure the VCS One agent for IBM SVCCopyServices, verify that the global composite service group can migrate to hosts across the sites. Depending on your DR configuration, perform one of the following procedures.

To test the global composite service group migration in global cluster setup

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

Perform the following steps:

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

```
hacsg -switch <global_csg> -clus <secondary_clusname>
```

VCS One brings the global composite service group online at the secondary site.

- Verify that the state of the replication is Consistent Synchronized and the SVC cluster at the secondary site is the primary SVC cluster.

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

Perform the following steps:

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

```
hacsg -switch <global_csg> -clus <primary_clusname>
```

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

- Verify that the state of the replication on the primary site is Consistent Synchronized and the SVC cluster at the primary site is the primary SVC cluster.

Testing disaster recovery after site failure

Review the details on site failure and how VCS One and the agent for IBM SVCCopyServices behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 33.

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

To test disaster recovery for site failure in global cluster setup

- 1 Halt all client systems and the arrays at the primary site.

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

You must bring the global composite service group online at the secondary site. Run the following command:

```
hacsg -online -force global_csg
```

- 2 Verify that the SVC disks at the secondary site are write-enabled and are in the disconnected state.
- 3 Verify that the global composite service group is online at the secondary site.

```
hacsg -state global_csg
```

Testing disaster recovery after client system failure

Review the details on client system failure and how VCS One and the agent for IBM SVCCopyServices behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 33.

Depending on the DR configuration, perform one of the following procedures to test how VCS One recovers after all client systems at the primary site fail.

To test disaster recovery for client system failure in global cluster setup

- 1 Halt the client system at the primary site.

You must bring the global composite service group online at the secondary site. Run the following command:

```
hacsg -online -force global_csg
```

- 2 Verify that the state of the replication is Consistent Synchronized and the SVC cluster at the secondary site is the primary SVC cluster. .
- 3 Verify that the global composite service group is online at the secondary site.

```
hacsg -state global_csg
```

Performing failback after a client system failure or an application failure

Review the details on client system failure and application failure and how VCS One and the agent for IBM SVCCopyServices behave in response to these failures.

See [“Failure scenarios in global clusters”](#) on page 33.

After the client systems at the primary site are restarted, you can perform a failback of the global composite service group to the primary site. Depending on your DR configuration, perform one of the following procedures.

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

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

```
hacsg -switch <global_csg> -clus <remote_clusname>
```

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

Verify that the SVC disks on the primary site are write-enabled and the SVC cluster at the primary site is the primary SVC cluster.

Performing failback after a site failure

After a site failure at the primary site, the hosts and the storage at the primary site are down. The administrator uses VCS One and brings the global composite service group online at the secondary site and the SVCCopyServices agent write enables the disks on the secondary site.

The state of the replication is disconnected. Review the details on site failure and how VCS One and the agent for SVCCopyServices behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 33.

When the hosts and the storage at the primary site are restarted and the replication link is restored, the state of the replication is Idling. The devices are write-enabled at both sites. You can now perform a failback of the global composite service group to the primary site.

To perform failback after a site failure in global cluster

- 1** Take the global composite service group offline at the secondary site. At the secondary site, run the following command:

```
hacsg -offline global_csg
```

- 2** Resynchronize the SVC disks at the primary site. To resync the SVC disks, restart the replication using the `startrc` command with the `-primary` option, specifying the SVC cluster at the secondary site as the primary cluster. Once the replication moves to a Consistent Synchronized state, make the SVC cluster at the primary site the primary SVC cluster using the `switchrc` command.

Alternately you can use the SVCCopyServices agent update action to resync the SVC disks at the primary site.

See [“About the SVC Copy Services agent’s update action function”](#) on page 16.

- 3** Bring the global composite service group online at the primary site. Run the following command:

```
hacsg -online global_csg
```


Index

A

- action/update function 12
- attribute definitions 28
- attributes 28
 - DisconnectTakeover attribute 28
 - GroupName attribute 28
 - IsConsistencyGroup attribute 28
 - SSHBinary attribute 28
 - SSHPATHToIDFile attribute 28
 - StopTakeover attribute 28
 - SVCClusterIP attribute 28
 - SVCUserName attribute 28

C

- clean function 12

D

- DisconnectTakeover attribute 28

E

- enterprise agent
 - removing 23

F

- failure scenarios
 - global clusters 33
 - application failure 33
 - client system failure 33
 - network failure 33
 - replication link failure 33
 - site failure 33
 - storage failure 33
- functions
 - action 12
 - clean 12
 - monitor 12
 - offline 12
 - online 12
 - open 12

G

- global clusters
 - failure scenarios 33
- GroupName attribute 28

I

- info function 12
- IsConsistencyGroup attribute 28

M

- monitor function 12

O

- offline function 12
- online function 12
- open function 12

S

- sample configuration 30
- SSHBinary attribute 28
- SSHPATHToIDFile attribute 28
- StopTakeover attribute 28
- SVC CopyServices agent
 - attribute definitions 28
- SVCClusterIP attribute 28
- SVCUserName attribute 28