

Veritas™ Cluster Server Release Notes

HP-UX 11i v3

5.1 Service Pack 1



Veritas™ Cluster Server Release Notes

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.

Product version: 5.1SP1

Document version: 5.1SP1.0

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate 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 support 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/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

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

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support 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 technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as 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 support 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

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support 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

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec Web site.

<https://sort.symantec.com/documents>

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:

docs@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Release Notes

This document includes the following topics:

- [About this document](#)
- [About Veritas Cluster Server](#)
- [About Symantec Operations Readiness Tools](#)
- [Important release information](#)
- [Changes introduced in 5.1 SP1](#)
- [VCS system requirements](#)
- [Features no longer supported](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation](#)

About this document

This document provides important information about Veritas Cluster Server (VCS) version 5.1 SP1 for HP-UX 11i v3. Review this entire document before you install or upgrade VCS.

The information in the Release Notes supersedes the information provided in the product documents for VCS.

This is Document version: 5.1SP1.0 of the *Veritas Cluster Server Release Notes*. Before you start, ensure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

<http://www.symantec.com/business/support/overview.jsp?pid=15107>

About Veritas Cluster Server

Veritas™ Cluster Server by Symantec (VCS) is a clustering solution that eliminates downtime, facilitates server consolidation and failover, and effectively manages a wide range of applications in heterogeneous environments.

About VCS agents

VCS bundled agents manage a cluster's key resources. The implementation and configuration of bundled agents vary by platform.

For more information about bundled agents, refer to the *Veritas Cluster Server Bundled Agents Reference Guide*.

The Veritas High Availability Agent Pack gives you access to agents that provide high availability for various applications, databases, and third-party storage solutions. The Agent Pack is available through Symantec™ Operations Readiness Tools (SORT).

For more information about SORT, see <https://sort.symantec.com/home>.

For information about agents under development and agents that are available through Symantec consulting services, contact your Symantec sales representative.

VCS provides a framework that allows for the creation of custom agents. Create agents in situations where the Veritas High Availability Agent Pack, the bundled agents, or the enterprise agents do not meet your needs.

For more information about the creation of custom agents, refer to the *Veritas Cluster Server Agent Developer's Guide*. You can also request a custom agent through Symantec consulting services.

About compiling custom agents

Custom agents must be developed using compilers from one of the products listed below:

- HP C/ANSI C Developer's Bundle (S800), part number B3901BA.
- HP aC++ Compiler (S800), part number B3913DB.

These products may be identified by various part numbers in addition to those listed.

Existing custom agents written to run on VCS versions earlier than 1.2 must be recompiled for use with VCS 5.1SP1.

About Symantec Operations Readiness Tools

Symantec™ Operations Readiness Tools (SORT) is a set of Web-based tools and services that lets you proactively manage your Symantec enterprise products. SORT automates and simplifies administration tasks, so you can manage your data center more efficiently and get the most out of your Symantec products. SORT lets you do the following:

- Collect, analyze, and report on server configurations across UNIX or Windows environments. You can use this data to do the following:
 - Assess whether your systems are ready to install or upgrade Symantec enterprise products
 - Tune environmental parameters so you can increase performance, availability, and use
 - Analyze your current deployment and identify the Symantec products and licenses you are using
- Upload configuration data to the SORT Web site, so you can share information with coworkers, managers, and Symantec Technical Support
- Compare your configurations to one another or to a standard build, so you can determine if a configuration has "drifted"
- Search for and download the latest product patches
- Get notifications about the latest updates for:
 - Patches
 - Hardware compatibility lists (HCLs)
 - Array Support Libraries (ASLs)
 - Array Policy Modules (APMs)
 - High availability agents
- Determine whether your Symantec enterprise product configurations conform to best practices
- Search and browse the latest product documentation
- Look up error code descriptions and solutions

Note: Certain features of SORT are not available for all products.

To access SORT, go to:

<http://sort.symantec.com>

Important release information

- The latest product documentation is available on the Symantec Web site at:
<http://www.symantec.com/business/support/overview.jsp?pid=15107>
- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:
<http://www.symantec.com/docs/TECH144835>
- For the latest patches available for this release, go to:
<http://sort.symantec.com/>

Changes introduced in 5.1 SP1

This section lists the changes for Veritas Cluster Server.

Changes related to the installation and upgrades

The product installer includes the following changes.

Using the installer for Veritas Dynamic Multi-pathing (DMP)

You can use the script- or Web-based installer to install, configure, and uninstall Veritas Dynamic Multi-pathing. You can enable DMP using the DMP license or using any Storage Foundation license key.

The VRTScutil and VRTSacclib are no longer in use

For all high availability products, the VRTScutil and VRTSacclib are no longer required.

See the *Veritas Cluster Server Installation and Configuration Guide*.

Installer-related changes to configure LLT private links, detect aggregated links, and configure LLT over UDP

For all high availability products, the installer provides the following new features in this release to configure LLT private links during the VCS configuration:

- The installer detects and lists the aggregated links that you can choose to configure as private heartbeat links.
- The installer provides an option to detect NICs on each system and network links, and sets link priority to configure LLT over Ethernet.
- The installer provides an option to configure LLT over UDP.

- The installer now supports VCS cluster configuration up to 64 nodes.

See the *Veritas Cluster Server Installation and Configuration Guide*.

Installer supports configuration of non-SCSI3 based fencing

You can now configure non-SCSI3 based fencing for VCS cluster using the installer.

See the *Veritas Cluster Server Installation and Configuration Guide*.

Web-based installer supports configuring VCS cluster in secure mode

You can now configure the VCS cluster in secure mode using the Web-based installer.

See the *Veritas Cluster Server Installation and Configuration Guide*.

The installer can copy CPI scripts to any given location using `-copyinstallscripts` option

The installer can copy CPI scripts to given location using `-copyinstallscripts` option. This option is used when customers install SFHA products manually and require CPI scripts stored on the system to perform product configuration, uninstallation, and licensing tasks without the product media.

See the *Veritas Cluster Server Installation and Configuration Guide*.

Web-based installer supports configuring disk-based fencing for VCS

You can now configure disk-based fencing for the VCS cluster using the Web-based installer.

See the *Veritas Cluster Server Installation and Configuration Guide*.

The installer can automatically detect and configure LLT links

The installer detects link connection status among all cluster nodes and chooses the most suitable links for LLT communication. It then can set the priority of the LLT private heartbeat links based on their media speed. Aggregated and bonded NICs are supported.

See the *Veritas Cluster Server Installation and Configuration Guide*.

The Web-based installer supports adding nodes

The Web-based installer has increased parity with the script-based installer. It now supports the ability to add nodes to a cluster. It also supports configuring secure clusters and fencing configuration.

The installer provides automated, password-less SSH configuration

When you use the installer, it enables SSH or RSH communication among nodes. It creates SSH keys and adds them to the authorization files. After a successful completion, the installer removes the keys and system names from the appropriate files.

When you use the installer for SSH communications, meet the following prerequisites:

- The SSH (or RSH) daemon must be running for auto-detection.
- You need the superuser passwords for the systems where you plan to install VCS.

The installer can check product versions

You can use the installer to identify the version (to the MP/RP/SP level depending on the product) on all platforms. Activate the version checker with `./installer -version system_name`.

Depending on the product, the version checker can identify versions from 3.5 onward.

The `installsfha` and `uninstallsfha` scripts are now available

The `installsfha` and `uninstallsfha` scripts are now available in the `storage_foundation_high_availability` directory to directly install, uninstall, or configure the Storage Foundation and High Availability product.

Upgrade changes

The following lists the upgrade changes in this release.

Supported paths for VCS upgrades that do not require a node reboot

When using the installer program to perform any of the typical upgrade listed in the following upgrade matrix, a node reboot is not required.

Upgrade matrix:

- 5.0 to 5.1SP1
- 5.0.1 to 5.1 SP1
- 5.0MP1 to 5.1SP1

For supported upgrade matrix, refer to the *Veritas Cluster Server Installation Guide*.

Upgrades that follow any other upgrade paths require a reboot.

Packaging updates

The following lists package changes in this release.

New VRTSamf package

VRTSamf is a new package introduced in this release. The Asynchronous Monitoring Framework (AMF) module, along with the VCS Agent Framework (AGFW) and resource agents provides a way to avoid polling for resource state changes. The AMF module allows the agent to register which resources to monitor and when to wait. The module provides the agent with immediate notification so that action can be taken at the time of the event. AMF enables the VCS agents to monitor a large number of resources with a minimal effect on performance.

VRTSacclib package is no longer shipped with VCS 5.1SP1

The VRTSacclib package was available with VCS 5.1. The package is not shipped with VCS 5.1SP1. The latest VRTSacclib package can be accessed from the Agent Pack release.

New VRTSvcsea package

The VCS agent binaries for Oracle, Sybase, and DB2 are now a part of VRTSvcscor package. The following individual agent packages for these agents are not shipped:

- VRTSvcscor
- VRTSvcssy
- VRTSvcscdb

Changes related to the VCS engine

This section lists the new features related to the VCS engine.

The HAD can exchange messages up to 64KB size

The size of the messages that HAD supports is increased from 16KB to 64KB. The messages can be exchanged between different HAD processes (running on different systems) or between CLI and HAD processes.

Refer to the following list for the message, object, attribute, and attribute values:

1. Maximum message size = 64KB
2. Maximum object name size = 1KB
3. Maximum attribute name size = 1KB
4. Maximum scalar attribute value size = 4KB
5. Maximum single key (of key-value pair) size = 4KB
6. Maximum single value (of key-value pair) size = 4KB
7. Maximum size of single element of vector or keylist pair = 4KB
8. Maximum user Name size = 1KB
9. Maximum password size = 255b
10. Maximum password encrypted size = 512b

Note: Points 2 through 10 were already supported in 5.0.1 release.

VCS engine allows deletion of individual value from a vector-type attribute

If there are multiple occurrences of the same value in the vector, then all instances of that value will be deleted.

VCS support for IPv6

VCS now supports IPv6 protocol.

VCS components that support IPv6 are as follows:

- VCS engine:
 - Supports IPv6 and IPv4 in a dual stack configuration and in a pure stack configuration (either IPv4 or IPv6).
 - You can use an IPv6 address as the value for the ClusterAddress attribute in the "Cluster" object.

Note: Simulator on Windows only supports IPv4.

- **Wide-Area Connector (WAC):**
 - You can use an IPv6 address as the value for the ClusterAddress attribute for the Cluster resource.
 - The ClusterAddress of all participating clusters in a global cluster option configuration should be from the same family (either IPv6 or IPv4).
- **Heartbeat agents**—You can use IPv6 addresses as the value of the Arguments attribute for the Icmp and IcmpS agents.
- **Steward**—You can use a list of IPv6 addresses as the value for the Steward attribute in the cluster resource.
- The following HP-UX networking agents now support IPv6 protocol:
 - IP agent
 - NIC agent
 - IPMultiNIC agent
 - MultiNICA agent
 - IPMultiNICB agent
 - MultiNICB agent

Support for a universally unique ID (UUID) for each cluster

This release introduces a universally unique ID for each cluster.

The VCS installer configures a UUID value for each cluster at the end of the configuration. If you manually configure a cluster, you must use the `uuidconfig.pl` utility to create a cluster UUID.

Changes related to the VCS commands

- The folder `/opt/VRTS/bin` includes links to commonly used VCS commands along with other SFHA product commands. Symantec recommends that you add this directory to your PATH environment variable.
For the commands that do not reside in the common folder, the VCS user documents specify the complete path for the command.
- VCS 5.1 SP1 includes the following new options and changes for the `ha` commands:
 - The VCS engine allows deleting an individual element from a vector-type attribute. If the vector list has multiple occurrences of the same value, then the VCS engine deletes all the occurrences of the value.
 - The `hagrpl -resources` command supports `-clus` | `-localclus` options.

```
hagrp -resources group [-clus cluster | -localclus]
```

The command displays the resource of a global group on a remote *cluster*. The option `-clus` displays information for the cluster designated by the variable *cluster*. The option `-localclus` specifies the local cluster.

- The `hastatus` command supports `-time` option.

```
hastatus [-sound] [-time] -sys sys [ -sys sys ... ]
```

The `-time` option prints the system time at which the status was received.

- The `hares` command supports the `-parentprop` option for taking a resource offline.

```
hares -offline [-ignoreparent | -parentprop] res -sys system
```

The `-parentprop` option stops all the parent resources in order before VCS takes the specific resource offline.

- The `switch group` command supports the `-any` option.

```
hagrp -switch group -any [-clus cluster | -localclus]
```

This option allows the switching of parallel global groups across a cluster. If you run this command to switch a parallel global service group across clusters, VCS brings the parallel service group online on all possible nodes in the remote cluster.

- The `ha` commands with `-modify` option now support `-insert` option. It enables you to add one or more values in the vector/keylist attribute at a given index.

```
hares -modify resource attr -insert index value ...
```

See the *Veritas Cluster Server Administrator's Guide* for more information.

First Failure Data Capture (FFDC) logs for support analysis

If VCS encounters some problem, then First Failure Data Capture (FFDC) logs are generated and dumped along with other core dumps and stack traces. If the debug logging is not turned on, these FFDC logs are useful to analyze the issues that require professional support.

See the *Veritas Cluster Server Administrator's Guide*.

New UUIDCONFIG(1M) man page

The new man page for UUIDCONFIG(1M) describes how to manage the cluster UUID (universally unique id) on the VCS nodes.

Support for intelligent monitoring of VCS resources using IMF

VCS now supports intelligent resource monitoring in addition to poll-based monitoring. Intelligent Monitoring Framework (IMF) is an extension to the VCS agent framework. You can enable or disable the intelligent monitoring functionality of VCS agents as needed.

The benefits of intelligent monitoring over poll-based monitoring are as follows:

- Faster notification of resource state changes.
- Reduction in VCS system utilization which enables VCS to effectively monitor a large number of resources.

See the *Veritas Cluster Server Administrator's Guide* for more information.

The following agents are IMF-aware in VCS 5.1 SP1:

- Mount
- Process
- Application
- Oracle
- Netlsnr
- CFMount
- CVMVxconfigd
- CFSfsckd

Changes related to managing clusters

VCS Single Cluster Manager web console is no longer available

VCS Single Cluster Manager web console is no longer available. For Web-based administration, Symantec recommends that you use Veritas Operations Manager (VOM).

To download the most current version of VCS Management Console, go to <http://go.symantec.com/vom/> and click **Download Now**.

Upgrading removes Cluster Connector component if configured. You need to upgrade VCS Management Console (formerly CMC) to version 5.5 to manage this

version of VCS. After you upgrade, you need to use Cluster Connector to Direct Connection conversion wizard in VCS Management Console.

Changes to Symantec Java Runtime Environment Redistribution

Symantec Java Runtime Environment Redistribution (VRTSjre15) is no longer packaged with VCS. Symantec recommends users to install native JRE 1.5 for any Symantec components that require it.

Make sure that you meet at least one of the following requirements for the Symantec components to run successfully:

- JAVA_HOME is specified and it points to Java v1.5+ installation
- /opt/VRTSjre/jre1.5/bin/java exists
- /usr/bin/java is at least v1.5
- \$PATH has java and is at least v1.5

Changes to VCS Java Console, VCS Simulator and VCS wizards

Following are the changes to the VCS Java Console, VCS Simulator and VCS wizards.

- Cluster Manager (Java Console) is no longer packaged with VCS. Symantec recommends using Veritas Operations Manager (VOM) to manage, monitor and report on multi-cluster environments. You can download this utility at no charge from <http://go.symantec.com/vom/> . If you wish to manage a single cluster using Cluster Manager (Java Console), a version is available for download from http://go.symantec.com/vcsm_download/. You can download VCS Java Console from http://go.symantec.com/vcsm_download/
- The Java-based configuration wizards (hawizards) for Oracle, NFS, and application agents are not supported for this release. Use VOM, the command line, or Cluster Manager (Java Console) to configure service groups for these applications.
- VCS Simulator is no longer packaged with VCS. You can download VCS Simulator from http://go.symantec.com/vcsm_download/.

New attributes

The following sections describes the attributes introduced in VCS 5.1SP1, VCS 5.0.1, and VCS 5.0MP3.

Attributes introduced in VCS 5.1SP1

Application Agent attributes

- **EnvFile**: This attribute specifies the environment file that must be sourced before running `StartProgram`, `StopProgram`, `MonitorProgram` or `CleanProgram`.
- **UseSUDash**: This attribute specifies that the agent must run `su - user -c <program>` or `su user -c <program>` while running `StartProgram`, `StopProgram`, `MonitorProgram` or `CleanProgram`.

RemoteGroup agent attribute

- **ReturnIntOffline**: This attribute can take one of the following three values. These values are not mutually exclusive and can be used in combination with one another. You must set `IntentionalOffline` attribute to 1 for the `ReturnIntOffline` attribute to work.
 - **RemotePartial**: Makes `RemoteGroup` resource to return `IntentionalOffline` when the remote service group is in `ONLINE|PARTIAL` state.
 - **RemoteOffline**: Makes `RemoteGroup` resource to return `IntentionalOffline` when the remote service group is in `OFFLINE` state.
 - **RemoteFaulted**: Makes `RemoteGroup` resource to return `IntentionalOffline` when the remote service group is in `OFFLINE|FAULTED` state.

DiskGroup agent attribute

- **Reservation**: Determines if you want to enable SCSI-3 reservation. For more information, please refer to *Bundled Agents Reference Guide*. In order to support SCSI-3 disk reservation, you must be sure that the disks are SCSI-3 compliant. Since all the disks are not SCSI-3 compliant, reservation commands fail on such disk groups. The `Reservation` attribute helps in resolving this issue. The `Reservation` attribute can have one of the following three values:
 - **ClusterDefault**: The disk group is imported with or without SCSI-3 reservation, based on the cluster-level `UseFence` attribute.
 - **SCSI3**: The disk group is imported with SCSI-3 reservation.
 - **NONE**: The disk group is imported without SCSI-3 reservation. The agent does not care about the cluster-level `UseFence` attribute.

Note: This attribute must be set to `NONE` for all resources of type `DiskGroup` in case of non-SCSI-3 fencing.

NFSRestart agent attribute

- **Lower:** Defines the position of the NFSRestart resource in the service group. The NFSRestart resource below the Share resource needs a value of 1. The NFSRestart resource on the top of the resource dependency tree has a Lower attribute value of 0.

RVGPrimary agent attribute

- **BunkerSyncTimeOut:** The timeout value in seconds that signifies the amount of time that a Secondary RVG can wait for the synchronization from the bunker host to complete before taking over the Primary role.

NotifierSourceIP agent attribute

- **NotifierSourceIP:** Lets you specify the interface that the notifier must use to send packets. This attribute is string/scalar. You must specify an IP address that is either DNS resolvable or appears in the `/etc/hosts` file.

SambaServer agent attributes

- **PidFile:** The absolute path to the Samba daemon (smbd) Pid file. This attribute is mandatory if you are using Samba configuration file with non-default name or path.
- **SocketAddress:** The IPv4 address where the Samba daemon (smbd) listens for connections. This attribute is mandatory if you are configuring multiple SambaServer resources on a node.
- **SambaTopDir:** Parent path of Samba daemon and binaries.

ASMIInst agent attributes

- **MonitorOption:** Enables or disables health check monitoring.

NetBios agent attribute

- **PidFile:** The absolute path to the Samba daemon (nmbd) PidFile. This attribute is mandatory if you are using Samba configuration file with non-default name or path.

Sybase agent attribute

- **Run_ServerFile:** The attribute specifies the location of the RUN_SERVER file for a Sybase instance. If this attribute is not specified, the default location of this file is accessed while starting Sybase server instances.
- **WaitForRecovery:** If this attribute is enabled during the online function, the agent waits until recovery is complete and all databases that can be made online are brought online.

Cluster-level attributes

- **AutoAddSystemToCSG:** Indicates whether the newly joined or added systems in the cluster become a part of the SystemList of the ClusterService service

group if the service group is confirmed. The value 1 (default) indicates that the new systems are added to SystemList of ClusterService. The value 0 indicates that the new systems are not added to SystemList of ClusterService.

- **CounterMissTolerance:** If GlobalCounter does not update in CounterMissTolerance intervals of CounterInterval, then VCS reports about this issue depending on the CounterMissAction (that is, CounterMissTolerance * CounterInterval) time has elapsed since last update of GlobalCounter then CounterMissAction is performed. The default value of CounterMissTolerance is 20.
- **CounterMissAction:** The action mentioned in CounterMissAction is performed whenever the GlobalCounter is not updated for CounterMissTolerance intervals of CounterInterval.
The two possible values of CounterMissAction are LogOnly and Trigger. LogOnly logs the message in Engine Log and SysLog. Trigger invokes a trigger which has a default action of collecting the comms tar file. The Default value of Trigger is LogOnly.
- **PreferredFencingPolicy:** The I/O fencing race policy to determine the surviving subcluster in the event of a network partition. Valid values are Disabled, System, or Group.
Disabled: Preferred fencing is disabled. The fencing driver favors the subcluster with maximum number of nodes during the race for coordination points.
System: The fencing driver gives preference to the system that is more powerful than others in terms of architecture, number of CPUs, or memory during the race for coordination points. VCS uses the system-level attribute FencingWeight to calculate the node weight.
Group: The fencing driver gives preference to the node with higher priority service groups during the race for coordination points. VCS uses the group-level attribute Priority to determine the node weight.

Resource type attributes

- **IMF:** Determines whether the IMF-aware agent must perform intelligent resource monitoring.
It is an association attribute with three keys Mode, MonitorFreq, and RegisterRetryLimit.
 - **Mode:** Defines whether to perform IMF monitoring based on the state of the resource. Mode can take values 0, 1, 2, or 3. Default is 0.
 - **MonitorFreq:** Specifies the frequency at which the agent invokes the monitor agent function. Default is 1.
 - **RegisterRetryLimit:** Defines the maximum number of times the agent attempts to register a resource. Default is 3.

- **IMFRegList:** Contains a list of attributes. The values of these attributes are registered with the IMF module for notification. If an attribute defined in IMFRegList attribute is changed then the resource, if already registered, is unregistered from IMF. If IMFRegList is not defined and if any attribute defined in ArgList is changed the resource is unregistered from IMF.
- **AlertOnMonitorTimeouts:** Indicates the number of consecutive monitor failures after which VCS sends an SNMP notification to the user.
- **ResourceSet:** A resource set is used to define a subset of processors in the system. If a resource set is specified for a workload partition, it can use the processors within the specified resource set only. The value of the ResourceSet attribute is the name of the resource set created using the `mkrset` command. If set, the agent configures the WPAR to use only the resource set specified by this attribute.
- **WorkLoad:** Allows modification of resource control attributes of WPAR - `shares_CPU` and `shares_memory`. This attribute has two keys, CPU and MEM. The key CPU is used to specify the number of processor shares that are available to the workload partition. The key MEM is used to specify the number of memory shares that are available to the workload partition.

Attributes introduced in VCS 5.0.1

VCS 5.0MP3 introduced the following attributes.

Resource type attributes:

- **FaultPropagation:** Specifies if VCS should propagate the fault up to parent resources and take the entire service group offline when a resource faults.
- **AgentFile:** Complete name and path of the binary for an agent. Use when the agent binaries are not installed at their default locations.
- **AgentDirectory:** Complete path of the directory in which the agent binary and scripts are located. Use when the agent binaries are not installed at their default locations.

Cluster level attributes:

- **DeleteOnlineResource:** Defines whether you can delete online resources.
- **HostMonLogLvl:** Controls the behavior of the HostMonitor daemon. Configure this attribute when you start the cluster. You cannot modify this attribute in a running cluster.
- **EngineShutdown:** Provides finer control over the `hastop` command.
- **BackupInterval:** Time period in minutes after which VCS backs up configuration files.

- **OperatorGroups:** List of operating system user account groups that have Operator privileges on the cluster.
- **AdministratorGroups:** List of operating system user account groups that have administrative privileges on the cluster.
- **Guests:** List of users that have Guest privileges on the cluster.

System level attributes:

- **EngineVersion:** Specifies the major, minor, maintenance-patch, and point-patch version of VCS.

Service group level attributes:

- **TriggerResFault:** Defines whether VCS invokes the resfault trigger when a resource faults.
- **AdministratorGroups:** List of operating system user account groups that have administrative privileges on the service group.
- **OperatorGroups:** List of operating system user account groups that have Operator privileges on the service group.
- **Guests:** List of users that have Guest privileges on the service group.

Changes to bundled agents

This section describes changes to the bundled agents for VCS.

New bundled agents

VCS has the following new agents:

- **CoordPoint**—Monitors coordination points in I/O fencing configurations.

The following Veritas Volume Replicator agents are now bundled as well:

- **RVG**—Brings the RVG online, monitors read and write access to the RVG, and takes the RVG offline.
- **RVGPrimary**—Attempts to migrate or takeover a Secondary to a Primary upon an application failover.
- **RVGSnapshot**—Creates and destroys a transactionally consistent space-optimized snapshot of all volumes in a VVR secondary replicated data set.
- **RVGShared**—Monitors the RVG in a shared environment.
- **RVGLogowner**—Assigns and unassigns a node as the logowner in the CVMcluster.

- RVGSharedPri—Attempts to migrate or takeover a Secondary to a Primary when a parallel service group fails over.

See the *Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide*.

Support for Veritas dynamic multi-pathing

The following agent supports Veritas Dynamic Multi-Pathing (DMP):

- LVMVolumeGroup agent
- LVMCombo agent

Support for IMF

The following agents support IMF:

- Process agent
- Mount agent
- Application agent

Note: Intelligent Monitoring Framework for mounts is supported only for the VxFS, CFS, and NFS mount types.

Changes to database agents

Updates to the VCS agent for Oracle

The Veritas Cluster Server agent for Oracle includes the following new or enhanced features:

- The VCS agent binaries for Oracle are now part of VRTSvcsea package. This package also includes the VCS agent binaries for DB2 and Sybase.
- If you installed the VCS agent binaries using the installer program, the program updates the main.cf file to include the appropriate agent types.cf files.
- The Oracle ASMInst agent has two new attributes: StartUpOpt and ShutDownOpt.

Changes to the Oracle agent

- Oracle agent now supports IMF monitoring.

Changes to the Sybase agent

- The Sybase agent supports a new optional attribute `Run_ServerFile`. The attribute specifies the location of the `RUN_SERVER` file for a Sybase instance. If this attribute is not specified, the default location of the `RUN_SERVER` file is accessed while starting Sybase server instances.
- The VCS agent binaries for Sybase are now a part of `VRTSvcsea` package. This package also includes the VCS agent binaries for DB2 and Oracle.
- The agent supports a new attribute `WaitForRecovery`. If this attribute is enabled, during the online function, the agent waits until recovery is completed and all databases that can be made online are brought online.

Changes to I/O fencing

This section covers the new features and changes related to I/O fencing in this release.

- Support for Coordination Point server (CP server)
You can use CP server as a coordination point with server-based I/O fencing. The Coordination Point server is a software solution based on the customized fencing mechanism, running on a remote system or cluster that provides arbitration functionality by allowing client cluster nodes to perform the fencing tasks.
- Unique I/O fencing keys for coordinator disks
The `vxfen` driver now encodes the LLT cluster ID in the SCSI3 keys registered on the coordinator disks. If the disk is zoned to multiple clusters, the I/O fencing key allows you to identify which cluster a coordinator disk belongs to. VCS 5.1 SP1 does not support sharing of coordinator disks across multiple clusters.
- New command options for `vxfenclearpre`
The `vxfenclearpre` command now includes the following options:
 - A coordinator-only disk option
 - An option to clear all keys from coordinator disks
 - An option to clear all keys with the VF prefix from the coordinator disks
 - An option to clear only the keys from the coordinator disks you specify in the `clusterid`
- New `-W` option for `vxfenconfig` command
The `vxfenconfig` command now has a `-W` option. You can use this option to display the supported and the current I/O fencing protocol versions.

- New `vxfen_vxfnd_tmt` tunable parameter
I/O fencing introduces a new tunable parameter `vxfen_vxfnd_tmt`. You can use this parameter to tune the time in seconds that the I/O fencing driver VxFEN must wait for the I/O fencing daemon VXFEND to return after completing a given task.

Support for preferred fencing

Traditional fencing prevents a split-brain condition by allowing only one of multiple sub-clusters to continue its operation in case a network partition disrupts regular communication between nodes. The preferred fencing feature gives preference to one sub-cluster over other sub-clusters in determining the surviving sub-cluster. This preference is based on factors such as which of the sub-clusters is running higher priority applications or the total importance of nodes which form that sub-cluster or both.

See the *Veritas Cluster Server Installation and Configuration Guide* and the *Veritas Cluster Server Administrator's Guide* for more details.

Support for Non-SCSI3 fencing

In environments that do not support SCSI-3, non-SCSI-3 fencing provides reasonable data protection by causing the winning side to delay by a configurable amount (`loser_exit_delay`, default 55). Additionally, Symantec has enhanced the fencing component to help panic the losing side quickly. Together, these enhancements help narrow down the window of potential data corruption drastically.

See the *Veritas Cluster Server Installation and Configuration Guide* and the *Veritas Cluster Server Administrator's Guide* for more details.

Changes to LLT

This release includes the following new features and changes to LLT:

- LLT startup time through the LLT init script is now optimized to use a constant time. LLT takes less than 16 seconds to start irrespective of the number of links specified in `/etc/llttab` file.
In the previous releases, LLT took around $(5 * \text{number_of_links_specified_in_the_}/etc/llttab_file)$ seconds to start.
- The `lltstat` command includes the following new options:
 - `lltstat -nv active`
This command filters the output of `lltstat -nv` to display the status of only the active nodes in the cluster.

- **lltstat -nv configured**

This command filters the output of `lltstat -nv` to display the status of only the configured nodes in the cluster. Configured nodes include active nodes and any additional nodes which are listed in the `/etc/llthosts` file.

See the `lltstat` manual page for more information.

- **Support for different link speeds for LLT links**

LLT now removes the restriction to use private NICs with same media speed. You can now use different media speed for the private NICs and configure the NICs with lesser speed as low-priority links to enhance LLT performance.

- **Support for destination-based load balancing**

LLT now also provides destination-based load balancing where the LLT link is chosen based on the destination node id and the port. With destination-based load balancing, LLT sends all the packets of a particular destination on a link.

- **The llconfig command includes the following new options:**

- **-W**, to print the LLT supported, broadcast, and current protocol version information.

- **-P**, to make some of the LLT parameters configurable.

- **Added a mechanism inside LLT to track the operating system timeouts registered by LLT.**

- **Added a separate tunable "peertroublelo" for specifying the trouble time for lo-pri links.**

- **The default heartbeating mechanism in LLT is now point-to-point unicast and not broadcast heartbeating.**

See the *Veritas Cluster Server Installation and Configuration Guide* and the *Veritas Cluster Server Administrator's Guide* for more details.

Changes to GAB

This section lists the new features and changes related to GAB in this release.

- **GAB logging daemon**

GAB implements a distributed network protocol. For situations when GAB decides to take the drastic action of killing its userland client process or panicking a node to resolve an issue, data from the affected node alone may not suffice for a meaningful support analysis. The new `gablogd` daemon attempts to address this issue. GAB starts this daemon by default at GAB configuration time.

See the *Veritas Cluster Server Administrator's Guide* for more information.

- **Registration monitoring**

The registration monitoring feature lets you configure GAB behavior when the VCS engine (HAD) is killed and does not reconnect after a specified time interval. This feature uses the settings in the environment variables `VCS_GAB_RMTIMEOUT` and `VCS_GAB_RMACTION` that are defined in the `vcse` file.

The `hashadow` process is now a real-time process.

See the *Veritas Cluster Server Administrator's Guide* for more information.
- **New `-W` option for `gabconfig` command**

The `gabconfig` command now has a `-W` option. You can use this option to display the supported and the current `gab` protocol versions.

Changes to VCS clusters running in secure mode

This section lists the changes in VCS 5.1 SP1 for clusters running in secure mode.

Support for passwordless login for non-root users

Support is added for passwordless login for non-root users to run HA commands on secure clusters.

See the *Veritas Cluster Server Administrator's Guide* for more information.

Support to enable LDAP authentication in secure clusters using AT CLIs

You can now use the `addldapdomain` and the `atldapconf` commands to enable LDAP authentication in secure clusters.

See the *Veritas Cluster Server Installation and Configuration Guide* for more details.

About the `ReturnIntOffline` attribute

The `ReturnIntOffline` attribute can take one of three values: `RemotePartial`, `RemoteOffline`, and `RemoteFaulted`.

These values are not mutually exclusive and can be used in combination with one another. You must set the `IntentionalOffline` attribute of `RemoteGroup` resource to 1 for the `ReturnIntOffline` attribute to work.

About the `RemotePartial` option

Select the `RemotePartial` value of this attribute when you want the `RemoteGroup` resource to return an `IntentionalOffline` when the remote service group is in an `ONLINE | PARTIAL` state.

About the RemoteOffline option

Select the RemoteOffline value of this attribute when you want the RemoteGroup resource to return an IntentionalOffline when the remote service group is in an OFFLINE state.

About the RemoteFaulted option

Select the RemoteFaulted value of this attribute when you want the RemoteGroup resource to return an IntentionalOffline when the remote service group is in an OFFLINE | FAULTED state.

Configuring RemoteGroup resources in parallel service groups

When a RemoteGroup resource is configured inside parallel service groups, it can come online on all the cluster nodes, including the offline nodes. Multiple instances of the RemoteGroup resource on cluster nodes can probe the state of a remote service group.

Note: The RemoteGroup resource automatically detects whether it is configured for a parallel service group or for a failover service group. No additional configuration is required to enable the RemoteGroup resource for parallel service groups.

A RemoteGroup resource in parallel service groups has the following characteristics:

- The RemoteGroup resource continues to monitor the remote service group even when the resource is offline.
- The RemoteGroup resource does not take the remote service group offline if the resource is online anywhere in the cluster.
- After an agent restarts, the RemoteGroup resource does not return offline if the resource is online on another cluster node.
- The RemoteGroup resource takes the remote service group offline if it is the only instance of RemoteGroup resource online in the cluster.
- An attempt to bring a RemoteGroup resource online has no effect if the same resource instance is online on another node in the cluster.

Changes to VCS logs

VCS 5.1 SP1 prints warning messages to STDERR. In earlier releases, VCS sent warning messages to STDOUT.

VCS system requirements

This section describes system requirements for VCS.

The following information applies to VCS clusters. The information does not apply to SF Oracle RAC installations.

VCS requires that all nodes in the cluster use the same processor architecture and run the same operating system version. All cluster nodes must be at the same patch level.

See “[Hardware compatibility list](#)” on page 30.

Hardware compatibility list

VCS supports vpar/npar from release VCS 5.0.1 and later.

The compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware go to the following URL:

<http://entsupport.symantec.com/docs/330441>

Before installing or upgrading Veritas Cluster Server, review the current compatibility list to confirm the compatibility of your hardware and software.

Supported software for VCS

VCS supports the following volume managers and file systems:

- Logical Volume Manager (LVM)
- HP File System (HFS)
- Veritas Storage Foundation (SF): Veritas Volume Manager (VxVM) with Veritas File System (VxFS)

VCS 5.1 SP1 supports the following versions of SF:

- SF 5.1 SP1
 - VxVM 5.1 SP1 with VxFS 5.1 SP1
- SF 5.0.1
 - VxVM 5.0.1 with VxFS 5.0.1

Note: VCS supports the previous version of SF and the next version of SF to facilitate product upgrades.

Supported VCS agents

[Table 1-1](#) lists the agents for enterprise applications and the software that the agents support.

Table 1-1 Supported software for the VCS agents for enterprise applications

Agent	Application	Application version	HP-UX 11i v3 version
DB2	DB2 Enterprise Server Edition	8.1, 8.2, 9.1, 9.5, 9.7	HP-UX 11i v3
Oracle	Oracle	11gR1, 10gR2, 11gR2	HP-UX 11i v3
Sybase	Sybase Adaptive Server Enterprise	12.5.x, 15.x	HP-UX 11i v3

See the *Veritas Cluster Server Installation Guide* for the agent for more details.

For a list of the VCS application agents and the software that the agents support, see the [Veritas Cluster Server Agents Support Matrix](#) at Symantec website.

Features no longer supported

No longer supported agents and components

VCS no longer supports the following:

- Configuration wizards
- Disk agent
- CampusCluster agent
- ServiceGroupHB agent.
This release does not support disk heartbeats. Symantec recommends using I/O fencing.
- VRTSWebApp
- Oracle 8.0.x, Oracle 8.1.x, and Oracle 9i - not supported by the updated Oracle agent.
- VCS documentation package (VRTSvcsdc)
The VCS documentation package (VRTSvcsdc) is deprecated. The software disc contains the documentation for VCS in Portable Document Format (PDF) in the *cluster_server/docs* directory.

Symantec recommends copying pertinent documents from the disc to your system directory /opt/VRTS/docs for reference.

- hahbsetup tool. This tool is removed as no supported feature requires this tool.
- VRTScutil package. This package is no longer supported.

Fixed issues

This section covers the incidents that are fixed in this release.

See the corresponding Release Notes for a complete list of fixed incidents related to that product.

See [“Documentation”](#) on page 69.

Fixed issues for VCS

[Table 1-2](#) lists the fixed issues for VCS .

Table 1-2 VCS fixed issues

Incident	Description
1142970	VCS logs an error "SSL Handshake failed" if the client creates a channel and then disconnects.
1179782	The haclus -state command displays inconsistent output in four-node cluster configurations
1214140	HAD reports incorrect CPU utilization
1403471 1397692	When in secure mode, GCO takes a long time to detect a cluster fault on a remote cluster
1404384	HAD crashes while switching over Global group and PreSwitch is set to TRUE.
1456724	Group switch/failover logic does not complete if the parent group gets autodisabled in between.

Table 1-2 VCS fixed issues (*continued*)

Incident	Description
1458338	<p>Even though /var/VRTSvcs/diag/had directory was empty, VCS was moving it to /var/VRTSvcs/diag/had.<timestamp>, and printed the following kind of a message on the console.</p> <p>VCS NOTICE V-16-1-53021 Diagnostics directory moved to /var/VRTSvcs/diag/had.1227075736, please check its contents and contact VERITAS Support.</p> <p>This was due to incorrect handling of directories.</p>
1479349	<p>On disabling the array side switch port, the system panics intermittently.</p>
1531512	<p>The Oracle agent picks up only the last corresponding action from oraerror.dat ignoring the previous actions. Even though the log shows the errors, the resource does not move to FAULT state.</p>
1531720	<p>Seeding of a port a does not seed other ports.</p>
1537433 1403471	<p>In secure global cluster environment, VCS took a long time to detect a cluster fault of a remote cluster. This was because in secure connection, a socket was opened as a blocking socket.</p>
1539089	<p>The agent framework leaked memory if there is continuous logging into agent's log file.</p>
1587173	<p>The LC_ALL value was set to empty string by the hastart script even though it was not present in the environment.</p>
1588784	<p>VCS engine does not support system names starting with numbers.</p>
1634031	<p>If a resource faults in a planned offline of a global group on node1 in a primary cluster, followed by planned online of the group in a remote cluster, the group can go online in the primary cluster if node1 gets rebooted, resulting in global concurrency violation. This may lead to data corruption.</p>
1710470	<p>Had may crash in a global cluster environment. In a global cluster environment, if the SystemList of a global group is modified to add new system in C1 then ResourceInfo attribute for all remote resources of this group, should get set to default value in C2 . This was not happening and hence hares -display for remote resources in C2 was causing _had to get SEGV.</p>
1780722	<p>For a SAMBA GROUP, "netbios" resource, with CIDR address for interface, fails to come ONLINE.</p>

Table 1-2 VCS fixed issues (*continued*)

Incident	Description
1789808	Cluster does not accept HA commands without reboot of whole cluster.
1795151	Global group fails to come online on the DR site with a message that it is in the middle of a group operation
1829180	The VCSAG_SU() function from vcsag_i18n_inc.sh file, has incorrect options to execute the su command.
1830978	HAD may crash while sending notifications to notifier if the send fails. This happens due to incorrect data-access.
1851078	RemoteGroup agent faults when set up for monitor only the local service group is taken offline
1852521	DiskGroupSnap agent assumes all nodes are part of a campus cluster configuration
1862229	<p>Symptom: Resource coordpoint keeps fluctuating between online/faulted status without actual failure.</p> <p>Description: The coordpoint resource reports a faulted status only when the number of coordination points with missing keys (of the local node) exceeds a predefined fault tolerance value. Though no coordpoint with missing keys was found, the agent reports as one or more of the coordination points are not accessible/missing out registrations. The fault is not with the cpagent functionality, but is observed that the <code>cpsat listpdcommand</code> was failing sometimes when more than the predefined processes execute the command simultaneously (using threads).</p> <p>Resolution: Moved out the code to execute the command only once much before forking the threads and store the value to be used by all them later. Thus, preventing each thread to simultaneously execute it which is leading to a failure sometimes.</p>
2097935	Need strict host name matching in coordination point installer.
2130967	Health check monitoring is performed in ASMinst agent, if newly added MonitorOption attribute is set to 0.

Bundled agents fixed issues

[Table 1-3](#) lists the fixed issues for bundled agents.

Table 1-3 Bundled agents fixed issues

Incident	Description
251704	The index number of the virtual interfaces changes after a failover. As a result applications using the interfaces may face communication failure.
252354	Application agent does not pass the value of CleanReason to the script that the Agent executes as per the value of the CleanProgram attribute.
2001039	In a frozen local service group, if a resource goes offline outside VCS control, a RemoteGroup resource also goes offline.
2019904	The OS determines the interface that the notifier uses to send packets. The user must be optionally able to specify the interface. This is an enhancement.
2045972	If you upgrade to VCS 5.1, Application resources may go to the FAULTED state.

VCS engine fixed issues

[Table 1-4](#) lists the fixed issues for VCS engine.

Table 1-4 VCS engine fixed issues

Incident	Description
254693	If the OnlineRetryLimit and PreOnline attributes are set for a group, then the group does not fail over in case of a fault.
970971	At unfreeze of a failover group, VCS does not evaluate the group for Concurrency violation.
1074707	A global group goes online on a remote site despite a concurrency violation. This behavior is observed if a related group resource goes intentionally online on the remote site.
1472734	VCS must log an alert message to increase the value of the ShutdownTimeout attribute on a multi-CPU computer.
1634494	If the GlobalCounter attribute does not increase at the configured interval, VCS must report the failure.

Table 1-4 VCS engine fixed issues (*continued*)

Incident	Description
1859387	When a parent group completely faults in a system zone, an online-local-hard (OLH) child group fails over to another system. This behavior is observed only if the child group is marked for manual failover in campus cluster. That is, the value of the AutoFailover attribute for the group is equal to 2.
1861740	The Subject line of the SMTP notification email must contain the Entity name. This is an enhancement.
1866434	When hashadow attempts to restart the High Availability Daemon (HAD) module, if the <code>/var/VRTSvcs/lock/.hadargs</code> file does not exist, hashadow gets the SIGSEGV signal.
1874261	VCS sets the MonitorOnly attribute of the resources in a frozen service group to 0. This behavior occurs even if the ExternalStateChange attribute is not set to OfflineGroup for a resource that goes intentionally offline.
1874533	If a resource, with the ExternalStateChange attribute set to OfflineGroup, goes intentionally offline, VCS sets the MonitorOnly attribute of all resources in a frozen service group to 0. VCS must set the MonitorOnly attribute only when required.
1915908	The <code>hares</code> command lets you have the "." special character in a resource name.
1958245	Notifier Agent is unable to get the local IP address in the linked-based IPMP.
1971256	If a service group faults during failover, then VCS may not honor the Prerequisites/Limits attributes on the target system.
2022123	The <code>notifier</code> command must let you specify the originating source IP address. This is an enhancement.
2061932	If you override a static attribute for a resource with an empty type definition, then VCS dumps duplicate entries for the attribute in the configuration file.
2081725	After a child group autostarts, VCS does not bring a partially-online parent to ONLINE state.

Table 1-4 VCS engine fixed issues (*continued*)

Incident	Description
2083232	If you configure an online-local dependency between two parallel groups, then the parent service group does not AutoStart on all nodes.
2084961	If the VCS configuration includes a file with an absolute pathname, then you cannot load templates from the cluster manager GUI.
2111296	In rare cases, VCS tries to bring online a failover service group that is already online on another node.
2128658	If you add a script-based WAN heartbeat, the heartbeat agent generates a core file and fails to report status.

Enterprise agents fixed issues

[Table 1-5](#) lists the fixed issues for enterprise agents.

Table 1-5 Enterprise agents fixed issues

Incident	Description
776230	The Sybase agent offline script must support the databases that require a longer time to shut down. This is an enhancement.
2117378	When the monitor entry function encounters an invalid process ID (PID) in a PID file, the function fails to detect the exit status of the process. This behavior occurs if the PID has an empty line below it. The function also accordingly fails to detect the correct state of a failed DB2 Connect instance.

LLT, GAB, and I/O fencing fixed issues

[Table 1-6](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-6 LLT, GAB, and I/O fencing fixed issues

Incident	Description
1908938	[GAB] In a large cluster, cascaded lowest node failures result in GAB panic during sequence space recovery.

Table 1-6 LLT, GAB, and I/O fencing fixed issues (*continued*)

Incident	Description
1840826	[GAB] Prevent 'gabconfig -c' while port 'a' is in the middle of iofence processing.
1989645	[GAB] Whenever there is a disparity in connected memberships, GAB master must wait for maximum <code>gab_conn_wait</code> time. However, when the GAB master is a new joinee it does not wait which caused wrong iofence message being sent to healthy nodes in the cluster.
2066020	[LLT] The <code>dlpiping</code> utility exits with an error similar to "dlpiping: send ECHO_REQ failed."
2005045	[LLT] The <code>hastart</code> command fails to start HAD on one of the nodes with message "GabHandle::open failed errno = 16" in syslog after HAD is stopped on all the nodes in the cluster simultaneously.
1859023	[LLT] The <code>lltconfig -T query</code> command displays a partially incorrect output
1846387 2084121	[Fencing] The <code>vxfenswap</code> and the <code>vxfentsthdw</code> utilities fail when rsh or ssh communication is not set to the same node.
1922413	[Fencing] The <code>vxfentsthdw</code> utility should detect storage arrays which interpret NULL keys as valid for registrations/reservations.
1847517	[Fencing] The <code>vxfenswap</code> utility has an incorrect usage message for <code>-n</code> option
1992560	[Fencing] The <code>vxfentsthdw</code> utility uses <code>scp</code> to communicate with the local host.
1512956	[Fencing] The <code>vxfenclearpre</code> utility displays error messages

Known issues

This section covers the known issues in this release.

See the corresponding Release Notes for a complete list of known issues related to that product.

See "[Documentation](#)" on page 69.

Issues related to installation

This section describes the known issues during installation and upgrade.

The Web-based installer does not work from the disc (2321818)

The Web-based installer fails to run.

Workarounds:

For this first workaround, you need to have about 1.7 GB of local storage available. Copy the disc to a local system and start the Web-based installer from the local copy. Symantec recommends that you use `cpio` for these operations.

If you have limited local disk space, use the second workaround.

To start the Web-based installer workaround

- 1 Create a mount point.

```
# mkdir /mnt/dvd
```

- 2 Optionally to find the specific device path (`/dev/dsk/cxtxdx`), run this command:

```
# /usr/sbin/ioscan -fnkC disk
```

- 3 Mount the disc to the mount point.

```
# mount /dev/dsk/cxtxdx /mnt/dvd
```

- 4 Create a temporary installation directory.

```
# mkdir /tmp/HXRT51SP1
```

- 5 Create a symbolic link from the disc to the temporary installation directory.

```
# ln -s /mnt/dvd/* /tmp/HXRT51SP1/
```

- 6 Remove the installer link from the temporary installation directory.

```
# rm -rf /tmp/HXRT51SP1/scripts
```

- 7 Copy the installer scripts from the disc to the temporary installation directory.

```
# cp -rf /mnt/dvd/scripts/ /tmp/HXRT51SP1/
```

- 8 Start the Web-based installer from the temporary installation directory.

```
# /tmp/HXRT51SP1/webinstaller start
```

Installation precheck can cause the installer to throw a license package warning (2320279)

If the installation precheck is attempted after another task completes (for example checking the description or requirements) the installer throws the license package warning. The warning reads:

```
VRTSvlic not installed on system_name
```

Workaround:

The warning is due to a software error and can be safely ignored.

While configuring authentication passwords through the Veritas product installer, the double quote character is not accepted (1245237)

The Veritas product installer prompts you to configure authentication passwords when you configure Veritas Cluster Server (VCS) as a secure cluster, or when you configure Symantec Product Authentication Service (AT) in authentication broker (AB) mode. If you use the Veritas product installer to configure authentication passwords, the double quote character (") is not accepted. Even though this special character is accepted by authentication, the installer does not correctly pass the characters through to the nodes.

Workaround: There is no workaround for this issue. When entering authentication passwords, do not use the double quote character (").

Manual upgrade of VRTSvlic package loses keyless product levels

If you upgrade the `VRTSvlic` package manually, the product levels that were set using `vxkeyless` may be lost. The output of the `vxkeyless display` command will not display correctly.

To prevent this, perform the following steps while manually upgrading the `VRTSvlic` package.

To manually upgrade the `VRTSvlic` package

1. Note down the list of products configured on the node for keyless licensing.

```
# vxkeyless display
```

2. Set the product level to NONE.

```
# vxkeyless set NONE
```

3. Upgrade the `VRTSvlic` package.


```
# swremove VRTSvlic
```

This step may report a dependency, which can be safely overridden.

```
swinstall -s 'pwd'
```

4. Restore the list of products that you noted in step 1.

```
# vxkeyless set product[|,product]
```

Installer assigns duplicate node ID during `-addnode` procedure

While performing an `-addnode` using a CPI installer to a cluster where a node has failed, VCS appends the new node with a duplicate node ID of its last node. This happens only to the cluster in which any but the last node has failed. In this case, `/etc/llthost` displays two nodes with same node IDs. This is because VCS assigns the node ID by simply counting the number of node entries without checking the assigned node IDs.

Workaround: Instead of running the CPI command, add the new node manually as described in the Veritas Cluster Server Installation Guide.

Issues with keyless licensing reminders after upgrading VRTSvlic

After upgrading from 5.0.1 to 5.1SP1, some keyless licenses may be left in the system. As a result, you may see periodic reminders being logged if the VOM server is not configured.

This happens if you were using keyless keys before upgrading to 5.1SP1. After the upgrade, you install real keys and run `vxkeyless set NONE`. In this case, the keyless licenses may not be completely removed and you see warning messages being logged after two months (if VOM server is not configured). This does not result in any functionality impact.

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.
2. Set the product level to `NONE` with the command:

```
# vxkeyless set NONE
```

3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic`:

- Verify if the key has `VXKEYLESS` feature Enabled using the following command:

```
# vxlicrep -k <license_key> | grep VXKEYLESS
```

- Delete the key if and only if `VXKEYLESS` feature is Enabled.

Note: When performing the search, do not include the `.vxlic` extension as part of the search string.

4. Restore the previous list of products with the command:

```
# vxkeyless set product1[,product]
```

The installer crashes when you add a node using the `-addnode` option

If you manually remove a node from a VCS cluster or SFRAC cluster and then add the node back to the cluster by using the installer, a duplicate node ID is created. The installer terminates abruptly.

Resolution: Add the node manually.

Errors recorded in the `swremove` logs of `VRTSgab` during VCS upgrade from 4.1 to 5.0.1

When VCS is upgraded from 4.1 to 5.0.1 on HP-UX 11i v3 using the Veritas product installer, the installer reports errors for GAB and errors are recorded in the `swremove` logs related to `VRTSgab`. [1719136]

You can safely ignore these error messages.

VCS agents dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the `update-ux` command

On PA-RISC architecture, the VCS agents (Oracle, Netlsnr, Sybase, SybaseBk, MultiNICB, and so on) may dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the `update-ux` command. [1630968]

This is because on HP-UX PA-RISC systems, the default thread stack size is limited to 64k. When the agent requires more than 64k stack memory, it may dump core due to SIGBUS error.

Workaround: Before running the `update-ux` command, edit the `/opt/VRTSvcs/bin/vcsenv` file to append following lines to it:

```
PLATFORM=`uname -s`  
ARCHITECTURE=`uname -m`  
if [ "${PLATFORM}" = "HP-UX" ] && [ "${ARCHITECTURE}" = "9000/800" ]; then
```

```
PTHREAD_DEFAULT_STACK_SIZE=524288
export PTHREAD_DEFAULT_STACK_SIZE
fi
```

Installer cannot split a cluster registered with one or more CP servers [2110148]

Splitting a cluster that uses server-based fencing is currently not supported. You can split a cluster into two and configure VCS on the two clusters using the installer.

For example, you can split a cluster `Clus1` into `clus1A` and `clus1B`. However, if you use the installer to reconfigure the VCS, the installer retains the same cluster UUID of `Clus1` in `clus1A` and `clus1B`. If both `clus1A` and `clus1B` use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Installer enters a loop when cluster is running in secure mode while configuring server-based fencing [2166599]

During server-based fencing configuration with a secure cluster, if `vxfen` fails to start and you retry server-based fencing configuration, the installer keeps asking to enter another system to enable security after you manually start VCS.

Workaround: When `vxfen` fails to start in customized mode for server-based fencing with a secure cluster, do not choose to retry configuring fencing. Select the default option and `vxfen` starts in disabled mode. You can also retry fencing configuration using `-fencing` option.

Issues related to any OS or supported technology

NFS cluster I/O fails when storage is disabled

The I/O from the NFS clusters are saved on a shared disk or a shared storage. When the shared disks or shared storage connected to the NFS clusters are disabled, the I/O from the NFS Client fails and an I/O error occurs.

Workaround: If the application exits (fails/stops), restart the application.

Operational issues for VCS

Volumes outside of VCS control that are mount locked cannot be unmounted without specifying the key

If a VxFS file system has "mntlock=key" in its mount options, then you cannot unmount the file system without specifying the key. Groups having DiskGroup resources configured with UmountVolumes set, may fail to switch or failover if the volumes are mount locked. [1276594]

Saving large configuration results in very large file size for main.cf

If your service groups have a large number resources or resource dependencies, and if the PrintTree attribute is set to 1, saving the configuration may cause the configuration file to become excessively large in size and may impact performance. [616818]

Workaround: Disable printing of resource trees in regenerated configuration files by setting the PrintTree attribute to 0.

AutoStart may violate limits and prerequisites load policy

The load failover policy of Service Group Workload Management may be violated during AutoStart when all of the following conditions are met: [251660]

- More than one autostart group uses the same Prerequisites.
- One group, G2, is already online on a node outside of VCS control, and the other group, G1, is offline when VCS is started on the node.
- The offline group is probed before the online group is probed.

In this scenario, VCS may choose the node where group G2 is online as the AutoStart node for group G1 even though the Prerequisites load policy for group G1 is not satisfied on that node.

Workaround: Persistently freeze all groups that share the same Prerequisites before using the `hastop -force` command to stop the cluster or node where any such group is online. This workaround is not required if the cluster or node is stopped without the force option.

Trigger not invoked in REMOTE_BUILD state

In some situations, VCS does not invoke the injeopardy trigger if the system is a REMOTE_BUILD state. VCS fires the trigger when the system goes to the RUNNING state.

The hagetcf script reports an error

Running the hagetcf script to gather information about the VCS cluster generates the following error:

```
tar: cannot stat ./var/VRTSvcS/log/*.A.log.  Not dumped.
```

Workaround: This message may be safely ignored.

Node cannot join cluster because port v is not ready for configuration

This behavior is observed when a node leaves a cluster and another node tries to join the cluster at the same time. If the GAB thread is stuck in another process, the new node cannot join the cluster and GAB logs the following warning:

```
GAB WARNING V-15-1-20126 Port v not ready  
for reconfiguration, will retry.
```

Using the coordinator attribute

This release contains an attribute for disk groups called coordinator, which configures disks as coordinator disks by the I/O fencing driver. Setting the attribute prevents the coordinator disks from being reassigned to other disk groups. See the Veritas Volume Manager documentation for additional information about the coordinator attribute.

The attribute requires that the disk group contain an odd number of disks. Symantec recommends that you use only three coordinator disks. Using more (five or seven) disks may result in different subclusters.

Some alert messages do not display correctly

The following alert messages do not display correctly [612268]:

51033	Global group %s is unable to failover within cluster %s and AutoFailOver is %s. Administrative action is required.
51032	Parallel global group %s faulted on system %s and is unable to failover within cluster %s. However, group is still online/partial on one or more systems in the cluster
51031	Unable to automatically fail over global group %s remotely because local cluster does not have Authority for the group.
51030	Unable to find a suitable remote failover target for global group %s. Administrative action is required

- 50916 Unable to automatically failover global group %s remotely due to inability to communicate with remote clusters. Please check WAN connection and state of wide area connector.
- 50914 Global group %s is unable to failover within cluster %s and ClusterFailOverPolicy is set to %s. Administrative action is required.
- 50913 Unable to automatically fail over global group %s remotely because clusters are disconnected and ClusterFailOverPolicy is set to %s. Administrative action is required.
- 50836 Remote cluster %s has faulted. Administrative action is required.
- 50761 Unable to automatically fail over global group %s remotely because ClusterList values for the group differ between the clusters. Administrative action is required.

Issues with configuration of resource values

If you configure a resource that has more than 425 values in its **ArgListValues**, the agent managing that resource logs a message such as:

```
VCS WARNING V-16-2-13806 Thread(1437547408) ArgListValues overflow;  
Cannot append values more than upper limit of (425).
```

Normally, the number of values in **ArgListValues** for a resource must not exceed 425. However, in case of a keylist, association or vector type of attribute appears in the ArgList for a resource-type. Since these attributes can take multiple values, there is a chance for the resource values in **ArgListValues** to exceed 425.

Issues with bunker replay

When ClusterFailoverPolicy is set to Auto and the AppGroup is configured only on some nodes of the primary cluster, global cluster immediately detects any system fault at the primary site and quickly fails over the AppGroup to the remote site. VVR might take longer to detect the fault at the primary site and to complete its configuration changes to reflect the fault.

This causes the RVGPrimary online at the failover site to fail and the following message is displayed:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdname  
could not be imported on bunker host hostname. Operation  
failed with error 256 and message VxVM  
VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote server  
unreachable...
```

```
Timestamp VCS ERROR V-16-2-13066 (hostname) Agent is calling
clean for resource (RVGPrimary) because the resource
is not up even after online completed.
```

Resolution: To ensure that global clustering successfully initiates a bunker replay, Symantec recommends that you set the value of the `OnlineRetryLimit` attribute to a non-zero value for `RVGPrimary` resource when the primary site has a bunker configured.

Volume group can cause concurrency violation under VCS control

When you restart the system, it causes automatic activation of the LVM volume groups. The LVM volume can cause concurrency violation issue for VCS if they are under VCS control.

To avoid this issue, you must disable auto-activation of the volume groups. Set the `AUTO_VG_ACTIVATE` variable to 0 in `/etc/lvmrc` file, using the following command:

```
# cat /etc/lvmrc |grep AUTO_VG_ACTIVATE

# AUTO_VG_ACTIVATE and RESYNC which are required by the script in /sbin/lvm
# AUTO_VG_ACTIVATE flag to 0 and customizing the function
# set AUTO_VG_ACTIVATE to 0.
AUTO_VG_ACTIVATE=0
```

Note: This routine is executed only if `AUTO_VG_ACTIVATE` is set to 1.

The `swverify` command displays a note

When you run `swverify` on HP-UX systems with VCS 5.1SP1 installation, the system displays the following note:

```
Note: Volatile file "/var/VRTSat/.VRTSat/Profile/vxatdlog.conf"
```

You can ignore this note as it does not affect the working of VCS. Do not regard it as an error or a warning.

The CmdServer process may not start in IPv6 environments in secure clusters

In an IPv6 environment on secure clusters, the CmdServer process may not start. In addition, security may not function correctly. If it does not start on a particular node, modify that node's `/etc/hosts` file so that the localhost resolves to `::1`.

Workaround: In the `/etc/hosts` file, add the following:

```
::1          localhost
```

Saving large configuration results in very large file size for main.cf

If your service groups have a large number resources or resource dependencies, and if the `PrintTree` attribute is set to 1, saving the configuration may cause the configuration file to become excessively large in size and may impact performance. [616818]

Workaround: Disable printing of resource trees in regenerated configuration files by setting the `PrintTree` attribute to 0.

Issues related to the VCS engine

LinkHbStatus does not reflect the link status correctly

After disabling the LLT links of a node, the `LinkHbStatus` does not reflect the 'DOWN' flag for that node in `'hasys -disp'`. [1831129]

Engine may hang in LEAVING state

When the `hares -online` command is issued for a parent resource when a child resource faults, and the `hares -online` command is followed by the `hastop -local` command on the same node, then the engine transitions to the leaving state and hangs.

Workaround: Issue the `hastop -local -force` command.

Timing issues with AutoStart policy

Consider a case where the service group is offline and engine is not running on node 1. If you restart the engine on node 1 after HAD is killed on node 2 and before the engine is restarted on node 2, then VCS does not initiate the autostart policy of the group.

On a default OEL4U4 install, VCS kernel components cannot start up

By default, OEL4U4 systems boot up in Xen-enabled kernels.

```
# uname -a  
  
Linux host1 2.6.18-164.el5xen #1 SMP Thu March 4 04:41:04 EDT 2010  
x86_64 x86_64 x86_64 GNU/Linux
```

However, VCS kernel modules are built only for the non-Xen kernels:

```
# cat kvers.lst  
  
2.6.18-8.el5v  
  
2.6.18-8.el5
```

Workaround: Set up your system for booting into the non-Xen kernels. For instructions, refer to the OS vendor's documentation.

New nodes get added to SystemList and AutoStartList attributes of ClusterService even if AutoAddSystemToCSG is disabled

The AutoAddSystemToCSG attribute determines whether the newly joined or added systems in a cluster become part of the SystemList of the ClusterService service group if the service group is configured. The value 1 (default) indicates that the new systems are added to SystemList of ClusterService.

AutoAddSystemToCSG has an impact only when you execute the `hasys -add` command or when a new node joins the cluster. [2159139]

However, when you use the installer to add a new node to the cluster, the installer modifies the SystemList and AutoStartList attributes irrespective of whether AutoAddSystemToCSG is enabled or disabled. The installer adds the new system to the SystemList and AutoStartList. To add nodes, the installer uses the following commands that are not affected by the value of AutoAddSystemToCSG:

```
# hagrpl -modify ClusterService SystemList -add newnode n  
# hagrpl -modify ClusterService AutoStartList -add newnode
```

Workaround

The installer will be modified in future to prevent automatic addition of nodes to SystemList and AutoStartList.

As a workaround, use the following commands to remove the nodes from the SystemList and AutoStartList:

```
# hagr -modify ClusterService SystemList -delete newnode  
# hagr -modify ClusterService AutoStartList -delete newnode
```

VCS fails to go to the running state on HP-UX 11.31 with March 2011 release

Due to a regression caused by the patch PHKL_41700 (QXCR1001078659) that went into HP-UX 11.31 March 2011 release, the `select()` call takes long time to return from 'timeout sleep'. Due to this, `_had` misses the heartbeat with GAB resulting in SIGABRT by GAB. [2287383]

Workaround: You must tune 'hires_timeout_enable' kernel parameter to 1 before starting the cluster. Run the following command to set this variable to 1:

```
# kctune hires_timeout_enable=1
```

Note: HP is likely to deliver the resolution for this issue via PHKL_41967 patch post the March 2011 release.

Issues related to the bundled agents

Application agent cannot monitor kernel processes

Application agent cannot monitor processes which have wildcard characters that give a special meaning to `grep` command. [1232043]

RemoteGroup agent's monitor function may time out when remote system is down

If a RemoteGroup agent tries to connect to a system (specified as `IpAddress`) that is down, the monitor function of the RemoteGroup agent times out for the resource. [1397692]

LVMVolumeGroup resources do not depend on DiskReservation resources

An LVMVolumeGroup resource does not depend on a DiskReservation resource. [1179518]

Problem in failing over the IP resource

When a system panics, the IP address remains plumbed to the system for a while. In such a case, VCS may not succeed in failing over the IP resource to another system. This can be observed when a system panics during I/O Fencing.

Workaround: Increase the value of the OnlineRetryLimit attribute for the IP resource type.

LVMLogicalVolume agent may hang

The LVMLogicalVolume agent may hang in some situations, depending on the value of the IOTimeout attribute. Symantec recommends using the LVMLCombo agent instead of the LVMLogicalVolume and LVMLVolumeGroup agents.

LVM agents do not detect disconnected cable

LVM commands continue to function correctly when the cable to disks is pulled. The LVM agent does not detect a fault in this situation.

MultiNICB agent on fails with IPv6 protocol if no network is specified

Description: If you configure MultiNICB agent with IPv6 protocol without specifying a host or by only specifying non-reachable hosts in NetworkHosts, the agent keeps switching the active interface.

Workaround: You must specify at least one reachable host in the NetworkHosts attribute.

Could not write IPMultiNICB Options to file

If you have not specified the Options attribute in IPMultiNICB resource, the following message is logged:

```
Could not write IPMultiNICB Options to file.
```

However, there is no functionality loss. [2234686]

Workaround: Either specify the Options attribute or ignore the log message.

MultiNICB resource goes to faulted state if you do not set the NetworkHosts attribute for IPv6 protocol

While using MultiNICB resource with interfaces configured with IPv6 protocol, the resource goes into faulted state if NetworkHosts attribute is not configured. [2132685]

Workaround: Set the NetworkHosts attribute for IPv6.

Issues related to global service groups

This section covers the issues related to global service groups.

Fault detection takes time in a global cluster running in secure mode

For global clusters running in secure mode, VCS may take a long time to detect a cluster fault on a remote cluster. [1403471]

Switch across clusters may cause concurrency violation

If you try to switch a global group across clusters while the group is in the process of switching across systems within the local cluster, then the group may go online on both the local and remote clusters. This issue affects only global groups. Local groups do not experience this behavior.

Workaround: Ensure that the group is not switching locally before attempting to switch the group remotely.

Global service group does not go online on AutoStart node

At cluster startup, if the last system where the global group is probed is not part of the group's AutoStartList, then the group does not AutoStart in the cluster. This issue affects only global groups. Local groups do not display this behavior.

Workaround: Ensure that the last system to join the cluster is a system in the group's AutoStartList.

Declare cluster dialog may not display highest priority cluster as failover target

When a global cluster fault occurs, the Declare Cluster dialog enables you to fail groups over to the local cluster. However, the local cluster may not be the cluster assigned highest priority in the cluster list.

Workaround: To bring a global group online on a remote cluster, do one of the following:

- From the Java Console, right-click the global group in the Cluster Explorer tree or Service Group View, and use the Remote Online operation to bring the group online on a remote cluster.
- From the Web Console, use the Operations links available on the Service Groups page to bring the global group online on a remote cluster.

Issues related to the VCS database agents

Issues related to the VCS Agent for DB2

This section covers issues related to the VCS agent for DB2.

awk error message

On IA-64, the default awk command may produce this error: Input line /usr/bin:/bin:/usr/s cannot be longer than 3,000 bytes. The source line number is 1.

Workaround: Install GNU awk.

Issues related to the VCS Agent for Oracle

This section covers the issues related to the VCS agent for Oracle.

NOFAILOVER action specified for certain Oracle errors

The Veritas High Availability agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file oraerror.dat, which consists of a list of Oracle errors and the actions to be taken.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for a description of the actions.

Currently, the reference file specifies the NOFAILOVER action when the following Oracle errors are encountered:

```
ORA-00061, ORA-02726, ORA-6108, ORA-06114
```

The NOFAILOVER action means that the agent sets the resource's state to OFFLINE and freezes the service group. You may stop the agent, edit the oraerror.dat file, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

Health check may not work for Oracle 10g R1 and 10g R2

For Oracle 10g R1 and 10g R2, if you set MonitorOption to 1, health check monitoring may not function when the following message is displayed [589934]:

```
Warning message - Output after executing Oracle Health  
Check is: GIM-00105: Shared memory region is corrupted.
```

Workaround: Set MonitorOption to 0 to continue monitoring the resource.

Health check monitoring is not supported for Oracle 11g R1 and 11g R2

The Oracle agent with 11g R1 and 11g R2 does not support Health check monitoring using the MonitorOption attribute. If the database is 11g R1 or 11g R2, the MonitorOption attribute for Oracle resource should be set to 0.

Intentional Offline feature is not supported for Oracle 11g R1 and 11g R2

The Oracle agent with 11g R1 and 11g R2 database does not support the Intentional Offline feature.

Pfile or SPfile is not supported on ASM diskgroups

The ASMInst agent does not support pfile or spfile for ASM Instance on ASM diskgroups in 11g R2. Symantec recommends you to store the file on the local file system.

ASM instance does not unmount VxVM volumes after ASMDG resource is offline

In configurations where ASMInstance resource is part of a separate parallel service group, the ASM instance does not unmount the volumes even after the ASMDG resource is taken offline. Therefore, the Volume resource cannot be taken offline. This issue occurs when you use VxVM volumes as ASM disk groups. [918022]

Workaround: Configure the ASMInstance resource as part of the failover service group where ASMDG resource is configured.

VCS agent for Oracle: Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 (1985055)

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2.

Workaround: Set MonitorOption attribute for Oracle resource to 0.

VCS agent for Oracle: Make sure that the ohasd has an entry in the init scripts

Make sure that the ohasd process has an entry in the init scripts so that when the process is killed or the machine is rebooted, this automatically restarts the process.[1985093]

Workaround: Respawn of ohasd process. Add the ohasd process in the /etc/inittab file to ensure that this process is automatically restarted when killed or the machine is rebooted.

VCS agent for Oracle: Intentional Offline does not work

Due to issues with health check monitoring, Intentional Offline does not work for VCS agent for Oracle.

Issues related to the Cluster Manager (Java Console)

This section covers the issues related to the Cluster Manager (Java Console).

32-bit JRE requirement

This release requires the installation of the 32-bit JRE `ibm-java-ppc-jre-6.0-6.0.ppc`. (1870929)

Cluster Manager (Java Console) may display an error while loading templates

You can access the Template View in the Cluster Manager from the Tools > Templates menu. If you have Storage Foundation configured in a VCS cluster setup, the following error may occur while the Cluster Manager loads the templates. (1433844)

```
VCS ERROR V-16-10-65 Could not load :-  
/etc/VRTSvcs/Templates/DB2udbGroup.tf
```

Workaround: Ignore the error.

Exception when selecting preferences

On Windows systems, selecting the Java (Metal) look and feel of the Java Console may cause a Java exception. [585532]

Workaround: After customizing the look and feel, close restart the Java Console.

Java Console errors in a localized environment

When connected to cluster systems using locales other than English, the Java Console does not allow importing resource types or loading templates from localized directories. [585532]

Workaround: The workaround is to copy the types files or templates to directories with english names and then perform the operation.

Printing to file from the VCS Java Console throws exception

VCS Java Console and Help throw an exception while printing to a file from a system that does not have a printer configured. Also, the content is not written to the file.

Workaround: Before printing, make sure at least one printer is configured on the system where the VCS Java Console is launched.

Common system names in a global cluster setup

If both local and remote systems have a common system name in a global cluster setup, group operations cannot be performed on those systems using the Java console.

Workaround: Use command-line interface to perform group operations.

Some Cluster Manager features fail to work in a firewall setup

In certain environments with firewall configurations between the Cluster Manager and the VCS cluster, the Cluster Manager fails with the following error message [1392406]:

```
V-16-10-13 Could not create CmdClient. Command Server  
may not be running on this system.
```

Workaround: You must open port 14150 on all the cluster nodes.

Issues related to VCS Simulator

This section covers the issues related to VCS Simulator.

Simulator clusters with Windows configurations fail to start on UNIX host platforms

The following clusters are affected: Win_Exch_2K3_primary, Win_Exch_2K3_secondary, Win_Exch_2K7_primary, Win_Exch_2K7_secondary, WIN_NTAP_EXCH_CL1, WIN_NTAP_EXCH_CL2, Win_SQL_Exch_SiteA, Win_SQL_Exch_SiteB, WIN_SQL_VVR_C1, WIN_SQL_VVR_C2. [1363167]

Workaround: For each of these clusters, there is a separate directory named after the cluster under the VCS Simulator installation directory

C:\Program Files\VERITAS\VCS Simulator on Windows

/opt/VRTScssim on Unix

Perform the following steps:

- Navigate to the conf/config directory under this cluster specific directory.
- Open the types.cf file in an editor and change all instances of the string "i18nstr" to "str".
- Open the SFWTypes.cf file in an editor if it exists in this directory and change all instances of the string "i18nstr" to "str".
- Repeat these steps for the following files if they exist: MSSearchTypes.cf, SQLServer2000Types.cf, ExchTypes.cf, SRDFTypes.cf.

VCS Simulator does not start on Windows systems

On Windows systems, starting VCS Simulator displays an error that the required MSVCR70.DLL is not found on the system. [859388]

Workaround: Run the following command:

```
set PATH=%PATH%;%VCS_SIMULATOR_HOME%\bin;
```

Or append %VCS_SIMULATOR_HOME%\bin; to PATH environment variable.

Error in LVMVolumeNFSGroup template for AIX

In the VCS Simulator, the AIX_NFS cluster gives error while loading the LVMVolumeGroupNFS template. [1363967]

This problem can also affect real AIX clusters if they try to load this template.

Workaround: For the Simulator, navigate to the Templates/aix directory under the VCS Simulator installation directory (C:\Program Files\VERITAS\VCS Simulator on Windows, /opt/VRTScsim on Unix). Open the LVMVolumeNFSGroup.tf file and look for all instances of the MajorNumber = "". Remove the empty double-quotes and set the correct integer value for MajorNumber.

For real clusters, make identical changes to /etc/VRTSvcs/Templates/LVMVolumeNFSGroup.tf.

VCS 5.0.1 Rolling Patch 1 known issues

The VCS issues in this release are as follows:

- The ASMInst agent does not support pfile or spfile for the ASM Instance on the ASM diskgroups in 11g Release 2. Symantec recommends that you store the file on the local file system. [1975010]
- The VRTSperl patch takes more than 10 minutes to install on an HP Integrity system node:
On an HP Integrity system node, installing the VRTSperl patch takes more than 10 minutes and requires that VCS is offline during this period. The installation time may vary based on the configuration of the machine on which the VRTSperl patch is being installed.

Issues related to AMF driver

AMF driver fails to unload with the Mount Agent running

If Mount Agent uses IMF to monitor mounts of type VxFS, then you cannot unload AMF driver as long as Mount Agent is running. [2262747]

Workaround: Stop the mount agent before you unload the AMF driver.

Startup or shutdown failure messages reported for LLT, GAB, and VXFEN

If you need to reboot the system when you install VCS, the init scripts for LLT, GAB, and VXFEN report start or stop failure messages. This is because VCS is not yet configured and the required configuration files are not yet generated for these components. These messages may be ignored. [1666327]

Issues related to I/O fencing

This section covers the known issues related to I/O fencing in this release.

All nodes in a sub-cluster panic if the node that races for I/O fencing panics

At the time of a network partition the lowest node in each sub-cluster races for the coordination points on behalf of that sub-cluster. If the lowest node is unable to contact a majority of the coordination points or the lowest node itself unexpectedly panics during the race, then all the nodes in that sub-cluster will panic. [1965954]

Coordination Point agent does not provide detailed log message for inaccessible CP servers

The Coordination Point agent does not log detailed information of the CP servers that are inaccessible. When CP server is not accessible, the agent does not mention the UUID or the virtual IP of the CP server in the engine log. [1907648]

Preferred fencing does not work as expected for large clusters in certain cases

If you have configured system-based or group-based preferred fencing policy, preferred fencing does not work if all the following cases are true:

- The fencing setup uses customized mode with one or more CP servers.

- The application cluster has more than eight nodes.
- The node weight for a single node (say galaxy with node id 0) is more than the sum total of node weights for the rest of the nodes.
- A network fault occurs and the cluster partitions into two with the single node (galaxy) on one part and the rest of the nodes on the other part.

Under such circumstances, for group-based preferred fencing, the single node panics even though more high priority services are online on that node. For system-based preferred fencing, the single node panics even though more weight is assigned to the node. [2161816]

See the *Veritas Cluster Server Administrator's Guide* for more information on preferred fencing.

Server-based I/O fencing fails to start after configuration on nodes with different locale settings

On each (application cluster) node, the vxfen module retrieves and stores the list of the UUIDs of coordination points. When different nodes have different locale settings, the list of UUIDs on one (application) node does not match with that of the other (application) nodes. Hence, I/O fencing does not start after configuration. [2112742]

Workaround: Start I/O fencing after fixing the locale settings to use the same values on all the (application) cluster nodes.

Reconfiguring VCS with I/O fencing fails if you use the same CP servers

When you reconfigure an application cluster that uses server-based I/O fencing (customized fencing mode), the installer does not remove the application cluster information from the CP servers before the reconfiguration. As a result, if you reconfigure the application cluster and choose to configure I/O fencing in customized mode using the same CP servers, then reconfiguration of server-based fencing for the application cluster fails. [2076240]

Workaround: Manually remove the application cluster information from the CP servers after you reconfigure VCS but before you reconfigure server-based I/O fencing for the application cluster.

See the *Veritas Cluster Server Administrator's Guide* for instructions to remove the application cluster information from the CP servers.

CP server cannot bind to multiple IPs (2085941)

Coordination point server (CP server) binds only to a single virtual IP and listens on the same. Application clusters cannot access the CP server if it fails to establish connection to this virtual IP. Therefore, if the connection fails because of the subnet in which the virtual IP of the CP server exists, you cannot access the CP server even if there is another subnet through which the client can connect to the CP server over a different IP.

Resolution: No known resolution for this issue.

Issues related to Symantec Product Authentication Service with VCS

This section covers the known issues related to Symantec Product Authentication Service (AT) in this release.

Issues related to Symantec Product Authentication Service with VCS

This section covers the issues related to Symantec Product Authentication Service with VCS.

The `atldapconf` command fails if the user in the Active Directory does not belong to any group

While using the `atldapconf` command, the user group must be specified. [1596332]

Output of `addldapdomain` returns error

The output of `addldapdomain` returns an error and the help contains incorrect information [1589886]

The `vcsat` and `cpsat` commands may appear to be hung

The following commands may appear to be hung when you invoke them from the command shell:

- `/opt/VRTScps/bin/cpsat`
- `/opt/VRTSvcs/bin/vcsat`

This issue occurs when the command requires some user interaction. [1841185]

Workaround:

- To fix the issue for `vcsat`, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTSvcs
# /opt/VRTSvcs/bin/vssatvcs command_line_argument
# unset EAT_HOME_DIR
```

- To fix the issue for cpsat, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTScps
# /opt/VRTScps/bin/vssatcps command_line_argument
# unset EAT_HOME_DIR
```

Verification for VRTSat package or patch returns errors

If you run swverify command on VRTSat package or patch, the command returns errors for missing files on VRTSat.CLIENT-PA32. [1244204]

Workaround: This message may be safely ignored.

The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups

The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups.

Workaround:

Have a copy of the pfile/spfile in the default \$GRID_HOME/dbs directory to make sure that this would be picked up during the ASM Instance startup.

Errors observed during partial upgrade of SFHA

While upgrading the VCS packages during an SFHA upgrade from 5.0 MP3 RP2 to 5.1SP1, CPI failed to uninstall the I/O fencing packages (VRTSvxfen, VRTSltt, and VRTSgab). [1779129]

Workaround

Before upgrading SFHA from 5.0 MP3 RP2 to 5.1SP1, you must apply the I/O fencing hotfix 5.0MP3RP2HF2.

Software limitations

This section covers the software limitations of this release.

See “[Documentation](#)” on page 69.

Limitations related to the bundled agents

Limitations of the DiskGroup agent

Volumes in disk group are started automatically if the Veritas Volume Manager default value of `AutoStartVolumes` at system level will be set to `ON` irrespective of the value of the `StartVolumes` attribute defined inside the VCS. Set `AutoStartVolumes` to `OFF` at system level if you do not want to start the volumes as part of import disk group.

Limitations related to global clusters

- Cluster address for global cluster requires resolved virtual IP.
The virtual IP address must have a DNS entry if virtual IP is used for heartbeat agents.
- Total number of clusters in a global cluster configuration can not exceed four.
- Cluster may not be declared as faulted when Symm heartbeat agent is configured even when all hosts are down.
The Symm agent is used to monitor the link between two Symmetrix arrays. When all the hosts are down in a cluster but the Symm agent is able to see the replication link between the local and remote storage, it would report the heartbeat as `ALIVE`. Due to this, DR site does not declare the primary site as faulted.

Systems in a cluster must have same system locale setting

VCS does not support clustering of systems with different system locales. All systems in a cluster must be set to the same locale.

VCS deletes user-defined VCS objects that use the HostMonitor object names

If you had defined the following objects in the `main.cf` file using the reserved words for the HostMonitor daemon, then VCS deletes these objects when the VCS engine starts [1293092]:

- Any group that you defined as `VCShmg` along with all its resources.
- Any resource type that you defined as `HostMonitor` along with all the resources of such resource type.
- Any resource that you defined as `VCSHm`.

GAB panics the systems while VCS gets diagnostic data

On receiving a SIGABRT signal from GAB, VCS engine forks off `vcs_diag` script. When VCS engine fails to heartbeat with GAB, often due to heavy load on the system, the `vcs_diag` script does a `sys req` to dump the stack trace of all processes in the system to collect diagnostic information. The dump of stack trace is intended to give useful information for finding out which processes puts heavy load. However, the dumping puts extra load on the system that causes GAB to panic the system in such heavy loads. See *Veritas Cluster Server User's Guide* for more information. [383970]

Workaround: Disable the `vcs_diag` script. To disable, rename the file `/opt/VRTSvcs/bin/vcs_diag` to `/opt/VRTSvcs/bin/vcs_diag.backup`.

Using agents in NIS

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can hang if the host is disconnected from the network. If such a program is used as an agent entry point, a network disconnect can cause the entry point to hang and possibly time out. For example, if the host is configured to use NIS maps as a client, basic commands such as `ps -ef` can hang if there is network disconnect. Symantec recommends creating users locally. To reflect local users, configure:

```
/etc/nsswitch.conf
```

VxVM site for the diskgroup remains detached after node reboot in campus clusters with fire drill

When you bring the `DiskGroupSnap` resource online, the `DiskGroupSnap` agent detaches the site from the target diskgroup defined. The `DiskGroupSnap` agent invokes VCS action entry points to run VxVM commands to detach the site. These commands must be run on the node where the diskgroup is imported, which is at the primary site.

If you attempt to shut down the node where the fire drill service group or the diskgroup is online, the node goes to a `LEAVING` state. The VCS engine attempts to take all the service groups offline on that node and rejects all action entry point requests. Therefore, the `DiskGroupSnap` agent cannot invoke the action to reattach the fire drill site to the target diskgroup. The agent logs a message that the node is in a leaving state and then removes the lock file. The agent's monitor function declares that the resource is offline. After the node restarts, the diskgroup site still remains detached. [1272012]

Workaround:

You must take the fire drill service group offline using the `hagr -offline` command before you shut down the node or before you stop VCS locally.

If the node has restarted, you must manually reattach the fire drill site to the diskgroup that is imported at the primary site.

Limitations with DiskGroupSnap agent

The DiskGroupSnap agent has the following limitations:

- The DiskGroupSnap agent does not support layered volumes. [1368385]
- If you use the Bronze configuration for the DiskGroupSnap resource, you could end up with inconsistent data at the secondary site in the following cases [1391445]:
 - After the fire drill service group is brought online, a disaster occurs at the primary site during the fire drill.
 - After the fire drill service group is taken offline, a disaster occurs at the primary while the disks at the secondary are resynchronizing.

Symantec recommends that you use the Gold configuration for the DiskGroupSnap resource.

Fire drill does not support volume sets

The fire drill feature for testing fault readiness of a VCS configuration supports only regular Volume Manager volumes. Volume sets are not supported in this release.

Manually removing VRTSat package erases user credentials

Symantec recommends saving user credentials before manually removing the VRTSat package. If you need the credentials again, you can restore them to their original locations.

To save user credentials

- 1 Run the `vssat showbackuplist` command. The command displays the data files and backs them up into the SnapShot directory `/var/VRTSatSnapShot`. Output resembles the following:

```
vssat showbackuplist
B| /var/VRTSat/.VRTSat/profile/VRTSatlocal.conf
B| /var/VRTSat/.VRTSat/profile/certstore
B| /var/VRTSat/RBAuthSource
B| /var/VRTSat/ABAuthSource
B| /etc/vx/vss/VRTSat.conf
Quiescing ...
Snapshot Directory :/var/VRTSatSnapShot
```

- 2 Move the credentials to a safe location. Preserving the directory structure makes restoring the files easier.

To restore user credentials

- 1 Navigate to the SnapShot directory or the safe location where you previously saved credentials:

```
cd /var/VRTSatSnapShot/
```

- 2 Restore the files:

```
cp ABAuthSource /var/VRTSat/
cp RBAuthSource /var/VRTSat/
cp VRTSat.conf /etc/vx/vss/
cd /var/VRTSatSnapShot/
cp -rp profile /var/VRTSat/.VRTSat/
```

Bundled agent limitations

This section covers the software limitations for VCS 5.0 bundled agents.

NFS wizard limitation

The NFS wizard allows only one NFS service group to be created. You need to create additional groups manually.

Volume agent clean may forcibly stop volume resources

When the attribute `FaultOnMonitorTimeouts` calls the Volume agent clean entry point after a monitor time-out, the `vxvol -f stop` command is also issued. This command forcibly stops all volumes, even if they are still mounted.

NFS failover

If the NFS share is exported to the world (*) and the NFS server fails over, NFS client displays the following error, “Permission denied”.

To avoid this error, export NFS shares explicitly using FQDN hostnames.

False concurrency violation when using PidFiles to monitor application resources

The PID files created by an application contain the PIDs for the processes that are monitored by Application agent. These files continue to exist even after a node running the application crashes. On restarting the node, the operating system may assign the PIDs listed in the PID files to other processes running on the node.

Thus, if the Application agent monitors the resource using the `PidFiles` attribute only, the agent may discover the processes running and report a false concurrency violation. This could result in some processes being killed that are not under VCS control.

Networking agents do not support IPv6 protocol

The bundled IP, NIC, IPMultiNIC, MultiNICA, IPMultiNICB, and MultiNICB agents for VCS 5.1 SP1 do not support the IPv6 enhanced IP protocol.

VCS does not provide a bundled agent for volume sets

VCS 5.1 SP1 does not provide a bundled agent to detect Volume Manager volume sets. Problems with volumes and volume sets can only be detected at the `DiskGroup` and `Mount` resource levels.

Workaround: Set `StartVolumes` and `StopVolumes` attributes of the `DiskGroup` resource that contains volume set to 1. If a file system is created on the volume set, use a `Mount` resource to mount the volume set.

Limitations related to I/O fencing

This section covers I/O fencing-related software limitations.

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *Veritas Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

I/O fencing uses SCSI-3 PR keys to implement data protection. Keys are placed on I/O fencing coordinator points and on data disks. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator points and data disks to prevent possible difficulties with subsequent cluster startup.

Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordinator points and data disks. Depending on the order of reboot and subsequent startup events, the cluster may warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown -r command on one node at a time and wait for each node to complete shutdown.

Cluster Management Console limitations

This section covers the software limitations for Cluster Management Console.

Cluster connector not supported on some OS versions

Cluster Management Console does not support cluster connector on AIX 5.1, Solaris 7, and RHEL 3.0. If your cluster runs on any of these platforms, you must use direct connection to manage the cluster from a management server.

Limited peer management server support

Peer management server support is limited to a configuration of two management servers in an enterprise. An enterprise of three or more management servers is not supported in this release.

Management server cannot coexist with GCM 3.5 Master

The Cluster Management Console management server should not be installed on the same system with a GCM 3.5 Master. These two products will conflict with each other and are not supported running on the same system.

Agent info files needed for Agent Inventory report

By design, the Agent Inventory report requires agent info files that supply the information reported on individual agents. These files are shipped with agents in VCS.

Global clusters must be CMC-managed clusters

All clusters forming a global cluster (using the VCS 4.0 Global Cluster Option) must be managed clusters in order for Veritas Cluster Management Console views to display correct and consistent information. Managed clusters are running the cluster connector or have a direct connection with the management server.

HP-UX cluster connector install fails if filesystem mount fails

If an HP-UX system has a mount that is not in `/etc/fstab` or `/etc/checklist`, the HP-UX installer `swinstall` will not work. Be sure the mount has an entry in these files.

Be sure the HP-UX system is configured to not attempt mounting of all the filesystems when performing an install or uninstall. This can be accomplished by adding the following lines to the `/var/adm/sw/defaults` file:

```
swinstall.mount_all_filesystems=false  
swremove.mount_all_filesystems=false
```

Windows Active Directory installation requires NetBIOS

If you install Cluster Management Console management server in a Windows Active Directory domain, NetBIOS must be turned on. A native (non-NetBIOS) Active Directory environment is not supported in this release.

Remote root broker not supported on Windows

If you set up a management server on a Windows system, you must configure a root broker on the management server system. This release does not support specifying a remote root broker during management server install [841739].

The root broker can be changed after install using the `configureRemoteRoot.exe` installed in `C:\Program Files\VERITAS\Cluster Management Console\bin` (default install directory).

Cluster Manager (Java console) limitations

This section covers the software limitations for Cluster Manager (Java Console).

Use the VCS 5.0 Java Console to manage clusters

Cluster Manager (Java Console) from previous VCS versions cannot be used to manage VCS 5.0 clusters. Symantec recommends using the latest version of Cluster Manager.

See the *Veritas Cluster Server Installation Guide* for instructions on upgrading Cluster Manager.

Run Java Console on a non-cluster system

Symantec recommends not running Cluster Manager (Java Console) for an extended period on a system in the cluster. The Solaris version of the Java Virtual Machine has a memory leak that can gradually consume the host system's swap space. This leak does not occur on Windows systems.

Cluster Manager and wizards do not work if the hosts file contains IPv6 entries

VCS Cluster Manager and Wizards fail to connect to the VCS engine if the `/etc/hosts` file contains IPv6 entries.

Workaround: Remove IPv6 entries from the `/etc/hosts` file.

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the `UseFence` attribute is set to the default, "None."

Undocumented commands, command options, and libraries

VCS contains undocumented commands and command options intended for development use only. Undocumented commands are not supported.

Documentation

Product guides are available on the documentation disc in PDF formats. Symantec recommends copying pertinent information, such as installation guides and release notes, from the disc to your system's `/opt/VRTS/docs` directory for reference.

Documentation set

[Table 1-7](#) lists the documents for Veritas Cluster Server.

Table 1-7 Veritas Cluster Server documentation

Title	File name
<i>Veritas Cluster Server Installation Guide</i>	vcs_install_51sp1_hpux.pdf
<i>Veritas Cluster Server Release Notes</i>	vcs_notes_51sp1_hpux.pdf
<i>Veritas Cluster Server Administrator's Guide</i>	vcs_admin_51sp1_hpux.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents_51sp1_hpux.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev_51sp1.pdf
<i>Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide</i>	vcs_vvr_agent_51sp1_hpux.pdf
<i>Veritas Cluster Server Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_agent_51sp1_hpux.pdf
<i>Veritas Cluster Server Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_agent_51sp1_hpux.pdf
<i>Veritas Cluster Server Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_agent_51sp1_hpux.pdf

[Table 1-8](#) lists the documentation for Veritas Volume Replicator.

Table 1-8 Veritas Volume Replicator documentation

Document title	File name
<i>Veritas Volume Replicator Administrator's Guide</i>	vvr_admin_51sp1_hpux.pdf
<i>Veritas Volume Replicator Planning and Tuning Guide</i>	vvr_planning_51sp1_hpux.pdf
<i>Veritas Volume Replicator Advisor User's Guide</i>	vvr_advisor_users_51sp1_hpux.pdf

[Table 1-9](#) lists the documentation for Symantec Product Authentication Service (AT).

Table 1-9 Symantec Product Authentication Service documentation

Title	File name
<i>Symantec Product Authentication Service Release Notes</i>	vxat_notes.pdf

Table 1-9 Symantec Product Authentication Service documentation (*continued*)

Title	File name
<i>Symantec Product Authentication Service Administrator's Guide</i>	vxat_admin.pdf

Manual pages

The manual pages for Veritas Storage Foundation and High Availability Solutions products are installed in the `/opt/VRTS/man` directory.

Set the `MANPATH` environment variable so the `man(1)` command can point to the Veritas Storage Foundation manual pages:

- For the Bourne or Korn shell (`sh` or `ksh`), enter the following commands:

```
MANPATH=$MANPATH:/opt/VRTS/man
export MANPATH
```

- For C shell (`csh` or `tcsh`), enter the following command:

```
setenv MANPATH ${MANPATH}:/opt/VRTS/man
```

See the `man(1)` manual page.

