

Veritas™ Cluster Server Release Notes

AIX

5.1



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: 5.1

Document version: 5.1.2

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's 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 automatic software upgrades protection
- 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

Documentation feedback

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

sfha_docs@symantec.com

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 [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

Additional enterprise services

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

Enterprise services that are available include the following:

Managed Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Education Services	Education Services provide a full array of technical training, security education, security certification, and awareness communication programs.

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

www.symantec.com/business/services/

Select your country or language from the site index.

Release Notes

This document includes the following topics:

- [Introduction](#)
- [About Veritas Cluster Server](#)
- [New features](#)
- [VCS system requirements](#)
- [No longer supported agents and components](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation errata](#)
- [VCS documentation](#)

Introduction

Before you start, make sure that you are using the latest version of this guide. It is online at:

http://sfdoccentral.symantec.com/sf/5.1/aix/pdf/vcs_notes.pdf

This is document version 5.1.2.

This document provides important information about Veritas Cluster Server (VCS) version 5.1 for AIX. 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.

For the latest patches available for this release, go to: <http://vos.symantec.com/>.

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/334998>

You can download the latest version of *Veritas Cluster Server Release Notes* from the link that is provided in the TechNote.

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

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

Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

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 bundles agents to 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 third-party storage solutions. Contact your Symantec sales representative for information about agents included in the agent pack, agents under development, and agents that are available through Symantec consulting services.

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 agents for enterprise applications do not meet your needs. You can also request a custom agent through Symantec consulting services.

For more information about the creation of custom agents, refer to the *Veritas Cluster Server Agent Developer's Guide*.

VCS also provides agents to manage key enterprise applications. Before configuring an enterprise agent with VCS, verify that you have a supported version of the agent.

See “[Supported VCS agents](#)” on page 36.

About compiling custom agents

Custom agents developed in C++ must be compiled using the IBM C for AIX Compiler Version 7.0. Use the `-brtl` flag for runtime linking with the framework library.

About Veritas Operations Services

Veritas Operations Services (VOS) is a Web-based application that is designed specifically for Veritas CommandCentral and Veritas Storage Foundation and High Availability products. VOS increases operational efficiency and helps improve application availability.

VOS automates and simplifies administrator tasks, including:

- Determining if systems are ready to install or upgrade Veritas products
- Gathering deployment and usage information on Veritas products
- Receiving notifications about the latest updates for:
 - Patches
 - Hardware Compatibility Lists (HCLs)
 - Array Support Libraries (ASLs)
 - Array Policy Modules (APMs)
- Determining whether your Veritas product configurations conform to best practices
- Managing server and environmental configuration data from a single Website
- Interpreting Unified Message Identifier (UMI) codes and their solutions
- Identifying and downloading patches for Veritas products

To access VOS, go to:

<http://vos.symantec.com/>

New features

This section lists the features introduced in the VCS 5.1 release.

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

Changes related to the VCS installer

This section lists the changes related to the VCS installer.

Changes to VCS fileset location on the product disc

The VCS filesets location on the product disc has changed from the `cluster_server` directory to the following:

```
/cdrom/pkgs
```

The VCS agent fileset for DB2, Oracle, and Sybase (VRTSvcsea) also resides in the same location as the other VCS filesets.

Option to install only the minimal packages

The Veritas Cluster Server product installer provides options for installing only selected packages. You can install the minimal packages, the recommended packages, or all of the packages.

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

Option to create response file templates

You can use the `-makeresponsefile` option of the installer to create response file templates.

The installer also generates a response file after each successful installer task, such as installation, configuration, uninstallation, or upgrade. These response files contain the details that you provided to the installer questions in the form of values for the response file variables. The response file also contains descriptions and explanations of the variables and their values.

See the *Veritas Cluster Server Installation Guide*.

Option to start or stop VCS

After the installation and configuration is complete, the installer starts the processes that the installed products use. You can use the product installer to stop or start the processes and load or unload the drivers, if required.

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

New startup and shutdown environment variables for VCS

VCS now has environment variables to control the startup and shutdown of the following VCS modules during system startup and shutdown:

- LLT: The file `/etc/default/llt` has `LLT_START` and `LLT_STOP` variables.
- GAB: The file `/etc/default/gab` has `GAB_START` and `GAB_STOP` variables.
- I/O fencing: The file `/etc/default/vxfen` has `VXFEN_START` and `VXFEN_STOP` variables.
- VCS engine: The file `/etc/default/vcs` has `VCS_START` and `VCS_STOP` variables.

The installer enables these variables after you configure VCS. You must explicitly enable these variables if you perform a manual configuration.

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

Support for installer resilience

If an installation or upgrade of VCS is interrupted, the next time you re-run it the installer discovers the presence of an installer instance. The installer then gives an option to resume the installation or upgrade.

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

Support for I/O fencing configuration using the installer

You can use the `-fencing` option of the `installvcs` to configure I/O fencing. Based on the fencing mechanism you want to use in the cluster, the installer provides the following options to configure I/O fencing:

- Disk-based I/O fencing - when you want to use disks as coordination points.
- Server-based I/O fencing - when you want to use CP servers as coordination points or a combination of coordinator disks and CP servers as coordination points.

The `installvcs` configures I/O fencing in disabled mode at the end of VCS configuration. However, split-brain prevention capabilities are not activated in disabled mode. Symantec recommends you to configure I/O fencing in enabled mode to use disk-based or server-based I/O fencing.

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

Changes to configuring clusters in secure mode

You can now configure one of the nodes within the cluster as root and authentication broker if you do not want to set up an external root broker system.

You can use the automatic configuration mode that the installer provides to configure a cluster node as the root broker.

The `-security` option of the `installvcs` no longer provides an option to configure root broker. You can only enable or disable AT for clusters using the `-security` option of the `installvcs`. You must use the installer or the `installat` program to install and configure root broker.

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

Installation options for NIM

The installer program can generate sample finish or configuration scripts that integrate with native installation tools. These sample scripts include all of the relevant Symantec packages in the correct order for installation. This feature is supported for installations, not upgrades.

Installer does not proceed with installation in RSH/SSH disabled environments

In the previous releases of VCS, in certain secure enterprise environments where RSH or SSH communication was not enabled, the installer installed and configured VCS only on the local system and the systems with which it could communicate. The installer also generated a response file that you could copy to the other systems in the cluster to identically install and configure VCS on other systems.

With this release of VCS, the installer mandates the availability of either RSH or SSH communication between the systems to perform any installer task.

Support for Web-based installer

This release supports an interactive installation using the Web-based installer. You can use a Web-interface to install and configure VCS.

The Web-installer can perform the following functions:

- Install VCS
- Uninstall VCS
- Configure VCS
- Upgrade VCS
- Start and stop VCS
- Perform an installation precheck

The installer program's default answer is no to configure optional features

The installer's default answer to configure optional features is now no. You must enter `y` if you want to configure certain optional features.

Support to enable rolling upgrades in future releases

VCS 5.1 adds support to enable rolling upgrades in future releases. These changes will help you upgrade to future versions of VCS with a minimal downtime of your infrastructure and applications during the upgrade process.

Gathering requirements using the installer program

You can use the `-requirements` option of the installer to gather the installation requirements. Web-based installer also provides you with a similar option.

The following information is displayed:

- Required operating system level
- Required patches
- Required disk space
- Other requirements

Support to continue installation after a successful precheck

The installer program has more sophisticated precheck, installation, and configuration options, which follow in outline:

- When you perform a successful precheck, you have the option to continue with the installation.
- After a successful installation, you have the option to continue with the configuration, or you can return to the configuration later.

Selecting default systems for installation

From the local system, the installer program checks for the `/etc/llhosts` for node names. When found, the installer program presents these as default nodes for installation. If the `llhosts` file is not present, then no default node names are provided.

Communication modes

By default, the installer program uses SSH for communication. The installer program switches to RSH if password-less SSH is not enabled.

For RSH communication, the `-rsh` option is available.

The installer programs supports mixed RSH and SSH modes on nodes in a cluster. The installation program can install on systems which may have heterogeneous (RSH and/or SSH) communication modes enabled.

IPv6 support for the installer programs

You can now use the installer to install and configure VCS on systems with IPv4, IPv6, or mixed stack configurations.

The following limitations apply:

- RSH is not supported on AIX
- SSH on AIX requires that you enable `sshd`.

Adding a node using the `-addnode` option

The `-addnode` option has been added to the installer to add a node to a running cluster. Based on the existing cluster configuration, the installer also configures the new node to use Symantec Product Authentication service and to use I/O fencing.

The installer also supports adding a node to a single node cluster, but stops the cluster during the addition of the node.

Installer support for alternate boot disk

The installer program supports install, uninstallation, and upgrades on alternate boot disks for AIX.

Refer to the *Veritas Cluster Server Installation Guide*.

Silent and automated installation enhancements for response files

The installer program supports silent installations using response files.

Operations that you can perform using response files follow:

- Fresh installations
- Configurations
- Uninstallations
- Upgrades from previous supported releases

Command options to help troubleshoot installations

You can run the installer with the `-debug` option and the `-trace` option to troubleshoot an installation.

Upgrade changes

The following lists upgrade changes in this release.

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

When you perform a typical upgrade using the installer program from VCS version 5.0 MP3 to VCS version 5.1, a node reboot is not required.

Upgrades that follow any other upgrade paths require a reboot.

Changes related to the installer for cross-product upgrades

This release includes the following changes related to the cross-product upgrades:

- If you try to perform a cross-product upgrade, the installer now gracefully exits with an error message.
For example, if you choose to upgrade VCS 5.0 MP3 to SFHA 5.1, the installer displays the following error message:

```
VCS 5.0.30.00 is installed.
```

```
Upgrading VCS 5.0.30.00 directly to SFHA 5.1 is not supported.
```

The installer does not support a direct upgrade from a previous VCS version to SFHA, SFCFS, or SF Oracle RAC version 5.1. You must upgrade VCS to version 5.1, and then install the 5.1 version of the stack product.

See the appropriate product Installation Guides for upgrade instructions.

- If a previous version of SFHA is installed, the installer supports partial product upgrade.
You can upgrade only VCS or SF to version 5.1. For example, you can upgrade SFHA 5.0 MP3 to VCS 5.1. If you want to upgrade the complete SFHA stack later, you can run the `installsf` program.

See the *Veritas Cluster Server Installation Guide* for VCS supported upgrade paths.

Other upgrade changes

You can use the `-upgrade` option of the installer to upgrade the product. Web-based installer also provides you with a similar option.

The installer program detects the installed product or products and upgrades them. The installer program detects the products in the following order: Storage Foundation RAC, Storage Foundation CFS, Storage Foundation High Availability, Storage Foundation, Veritas Cluster Server, Veritas Volume Manager/Veritas File System/Veritas Volume Replicator. The installer then upgrades any packages on the node and installs the 5.1 packages.

Changes related to the VCS engine

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

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.

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.

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 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.

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

The command displays the resource of a global group on a remote cluster *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.

```
hagr -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 HAWPARVERIFY(1M) and HAWPARSETUP(1M) man pages

New man pages for WPARs have been added. The `HAWPARVERIFY(1M)` man page describes how to verify the WPAR configuration. The `HAWPARSETUP(1M)` man page describes which configure VCS service group to include WPAR resource.

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.

VCS support for IPv6

VCS components that support IPv6 are as follows:

- VCS engine information follows:
 - Supports IPv6 and IPv4 in a dual stack configuration and in a pure stack configuration (either IPv4 or IPv6).
 - Simulator on Windows supports IPv4 only.
 - You can use an IPv6 address as the value for the ClusterAddress attribute in the "Cluster" object.
- Wide-Area Connector (WAC) information follows:
 - 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.

New virtualization features

This release introduces the following virtualization features.

VCS supports Live Partition Mobility

You can use Live Partition Mobility to perform a stateful migration of an LPAR in a VCS environment.

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

WPAR-awareness

In this release a new model to enable VCS to work with WPARs has been introduced. A ContainerInfo service group level attribute and a ContainerOpts type level static attribute work together to provide WPAR support.

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

Configuring failovers among physical and virtual servers

You can configure VCS to fail over from a physical system to a virtual system and vice versa. A physical to virtual failover gives an N + N architecture in an N + 1 environment. For example, several physical servers with applications can fail over to containers on another physical server.

See [“Configuring for failovers—a typical setup”](#) on page 73.

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.
- **MemCPUAllocator**—Allocates CPU and memory to an IBM AIX dedicated partition.
- **WPAR**—Brings online, takes offline, and monitors workload partitions.

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*.

See [“No longer supported agents and components”](#) on page 37.

Intelligent resources

The monitor entry points (agent functions) of intelligent resources provides the following two extra pieces of information:

- Confidence level—This is the confidence the agent has when the monitor entry point returns a state. For example, an agent can return an online state with a 50% confidence level.
- Confidence message—This is the message that explains the reason for a less than 100% confidence level.

You can use the `hares -display` command to determine the confidence level and confidence message for a given resource.

This is a reporting feature and VCS triggers no action based on these reports.

The following agents on AIX support this feature:

- NIC
- MultiNICA
- MultiNICB

The Mount agent supports loopback mounts

The Mount agent now supports loopback mounts.

IPv6 support for networking agents

Networking agents—You can configure all of the networking agents to bring online, take offline, and fail over IPv6 addresses.

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 Cluster Server Management Server.

To download the most current version of VCS Management Console, go to <http://www.symantec.com/business/cluster-server> and click **Utilities**.

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 Cluster Server Management Console to manage, monitor and report on multi-cluster environments. You can download this utility at no charge from <http://go.symantec.com/vcsmc>.
If you wish to manage a single cluster using Cluster Manager (Java Console), a version is available for download from the same website.
You can download VCS Java Console from <http://go.symantec.com/vcsmc>
- The Java-based configuration wizards (hawizards) for Oracle, NFS, application agents, and logical domains (LDoms) are not supported for this release. Use VCS Management Console, 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://www.symantec.com/business/cluster-server> and click **Utilities**.

New attributes

The following sections introduce attributes for VCS 5.1 and VCS 5.0 MP3.

Attributes introduced in VCS 5.1

VCS 5.1 introduces the following new attributes. See the *Veritas Cluster Server Administrator's Guide* for more information.

Resource type attributes:

- **ContainerOpts**: Specifies the behavior of the agent in a container environment.
- **CleanRetryLimit**: Number of times to retry the clean function before moving a resource to ADMIN_WAIT state.
- **EPClass**: Enables you to control the scheduling class for the agent functions (entry points) except the online entry point.
- **EPPriority**: Enables you to control the scheduling priority for the agent functions (entry points) except the online entry point.
- **FaultPropogation**: Specifies if VCS should propagate the fault up to parent resources and take the entire service group offline when a resource faults.
- **OnlineClass**: Enables you to control the scheduling class for the online agent function (entry point).
- **OnlinePriority**: Enables you to control the scheduling priority for the online agent function (entry point).

Service group attribute:

- **ContainerInfo**: Specifies information about the container that the service group manages.

Cluster attributes:

- **CID**: The CID provides universally unique identification for a cluster.
- **DeleteOnlineResource**: Defines whether you can delete online resources.
- **HostMonLogLvl**: Controls the behavior of the HostMonitor feature.

Attributes introduced in VCS 5.0 MP3

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 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 attributes:

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

Service group 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.

Removed attributes

The following attributes are obsolete for VCS 5.1:

- **OfflineProcScanInterval**
- **ProcScanInterval**

The following attributes are obsolete for VCS 5.0 MP3:

- **DiskHbStatus**
- **MajorVersion**
- **MinorVersion**

Packaging updates

[Table 1-1](#) lists the updates related to packages for this release.

Table 1-1 List of packages

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
N/A	VRTScps	New package.	Veritas Cluster Server Coordination Point Server
N/A	VRTSvcsea	New package.	Veritas Cluster Server Enterprise Agents
N/A	VRTSsfmh	New package.	Veritas Storage Foundation Managed Host
N/A	VRTSaslapm	New package.	Volume Manager ASL/APM
SYMClma	N/A	Obsolete in 5.0MP3. Functionality dropped.	Symantec License Inventory Agent
VRTSaa	VRTSsfmh	Consolidated into VRTSsfmh.	Veritas Enterprise Administrator action agent
VRTSacclib	N/A	Obsolete in 5.1. Not available for fresh installation. Only available to upgrade customers.	Veritas Cluster Server ACC Library 5.0 by Symantec
VRTSalloc	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation Intelligent Storage Provisioning
VRTSat	VRTSat	No change.	Symantec Product Authentication Service
VRTScavf	VRTScavf	No change.	Veritas Cluster Server Agents for Storage Foundation Cluster File System
VRTSccg	VRTSsfmh	Consolidated into VRTSsfmh.	Veritas Enterprise Administrator Central Control Grid
VRTScfsdc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	Veritas Cluster File System Documentation

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTScmccc	N/A	Obsolete in 5.1. Delivered with 5.x CMC release.	Veritas Cluster Management Console Cluster Connector
VRTScmcdc	N/A	Obsolete in 5.0MP3. Delivered with 5.x CMC release.	User Documentation for Veritas Cluster Management Console
VRTScmcm	N/A	Obsolete in 5.0MP3. Delivered with 5.x CMC release.	Veritas Cluster Management Console for multicluster environments
VRTScmcs	N/A	Obsolete in 5.1. Delivered with 5.x CMC release.	Veritas Cluster Management Console for single cluster environments
VRTScs	N/A	Obsolete in 5.0MP3. Delivered with SFM release.	Veritas Centralized Management for Storage Foundation Management Server
VRTScscm	N/A	Obsolete in 5.1. Available for download from http://go.symantec.com/vcsmc	Veritas Cluster Server Cluster Manager
VRTScscw	N/A	Obsolete in 5.1.	Veritas Cluster Server Configuration Wizards
VRTScsdoc	N/A	Obsolete in 5.0MP3. Delivered with SFM release.	Veritas Enterprise Administrator Central Server Documentation
VRTScsocw	N/A	Obsolete in 5.1.	Veritas Cluster Server Oracle and RAC Configuration Wizards
VRTScssim	N/A	Obsolete in 5.1. Available for download from http://go.symantec.com/vcsmc .	Veritas Cluster Server Simulator
VRTScutil	VRTScutil	Expanded to include few VCS packages.	Veritas Cluster Utility
VRTScweb	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Enterprise Administrator Central Server Documentation

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSd2gui	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation Graphical User Interface for DB2
VRTSdb2ed	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation for DB2
VRTSdbac	VRTSdbac	No change.	Veritas Oracle Real Application Cluster Support Package
VRTSdbcom	VRTSdbed	Consolidated into VRTSdbed.	Veritas Storage Foundation Common Utilities for Databases
VRTSdbdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	Veritas Storage Foundation Documentation for Databases
VRTSdbed	VRTSdbed	Expanded to include DBED packages.	Veritas Storage Foundation for Oracle
VRTSdbms3	N/A	Obsolete in 5.1. Sybase ASA repository no longer used in 5.1.	Symantec Shared DBMS
VRTSdcli	N/A	Obsolete in 5.1.	Veritas Distributed Command Line Interface
VRTSdcp	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Disk Correlator Provider
VRTSddlpr	N/A	Obsolete in 5.1. Functionality merged into VRTSob.	Veritas Device Discovery Layer Services Provider
VRTSdsa	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Datacenter Storage Agent
VRTSdsm	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Datacenter Storage Manager

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSfas	N/A	Obsolete in 5.0MP3. Functionality dropped.	Veritas FlashSnap Agent for Symmetrix
VRTSfasag	N/A	Obsolete in 5.0MP3. Functionality dropped.	Veritas Cluster Server Agents for Veritas FlashSnap Agent for Symmetrix
VRTSfasdc	N/A	Obsolete in 5.0MP3. Functionality dropped.	Veritas FlashSnap Agent for Symmetrix Documentation
VRTSfsdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	Veritas File System Documentation
VRTSfsman	VRTSvxfs	Consolidated into VRTSvxfs.	Veritas File System - Manual Pages
VRTSfsmnd	VRTSfssdk	Consolidated into VRTSfssdk.	Veritas File System SDK - Manual Pages
VRTSfspro	VRTSob	Consolidated into VRTSob.	Veritas File System Management Services Provider
VRTSfssdk	VRTSfssdk	No change.	Veritas File System SDK
VRTSfsweb	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas File System Provider Web Client Extension
VRTSgab	VRTSgab	No change.	Veritas Group Membership and Atomic Broadcast
VRTSgapms	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Generic Array Plug-in for Mapping Services
VRTSgcscha	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas GCS high availability agents
VRTSgcspr	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas SAN Global Configuration Server Object Bus Provider

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSglm	VRTSglm	No change.	Veritas Global Lock Manager
VRTSgms	VRTSgms	No change.	Veritas Group Messaging Services
VRTSicsco	N/A	Obsolete in 5.1.	Symantec Infrastructure Core Services Common
VRTSjre	N/A	Obsolete in 5.0MP3.	Veritas Java Runtime Environment Redistribution
VRTSjre15	N/A	Obsolete in 5.1.	Symantec Java Runtime Environment Redistribution
VRTSllt	VRTSllt	No change.	Veritas Low Latency Transport
VRTSmapro	N/A	Consolidated into VRTSob.	Veritas Storage Mapping Provider
VRTSmh	VRTSsfmh	Consolidated into VRTSsfmh.	Veritas Storage Foundation Management host
VRTSob	VRTSob	No change.	Veritas Enterprise Administrator Service
VRTSobc33	N/A	Obsolete in 5.1. Functionality Delivered with SFM release	Veritas Enterprise Administrator Core
VRTSobgui	N/A	Obsolete in 5.1. Functionality Delivered with SFM release.	Veritas Enterprise Administrator
VRTSobweb	N/A	Obsolete in 5.1. Functionality Delivered with SFM release.	Veritas Enterprise Administrator Web Console
VRTSodm	VRTSodm	No change	Veritas Oracle Disk Manager
VRTSorgui	N/A	Obsolete in 5.1. No longer supported.	Veritas Storage Foundation Graphical User Interface for Oracle
VRTSspb	N/A	Obsolete in 5.1.	Symantec Private Branch Exchange

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSperl	VRTSperl	No change.	Veritas Perl 5.8.8 redistribution
VRTSsmf	N/A	Obsolete in 5.0MP3.	Symantec Service Management Framework
VRTSspt	VRTSspt	No change.	Veritas Software Support Tools
VRTSsybed	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation for Sybase
VRTSvail	N/A	Obsolete in 5.1. Functionality Delivered with SFM release.	Veritas Array Providers
VRTSvcsc	VRTSvcsc	Expanded to include few VCS packages.	Veritas Cluster Server
VRTSvcscag	VRTSvcscag	Expanded to include agents previously included in VRTSvcscvr.	Veritas Cluster Server Bundled Agents
VRTSvcscdb	VRTSvcscsea	Consolidated into VRTSvcscsea.	Veritas High Availability Agent for DB2
VRTSvcscdc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	User Documentation for Veritas Cluster Server
VRTSvcscmsg	VRTSvcsc	Consolidated into VRTSvcsc.	Veritas Cluster Server English Message Catalogs
VRTSvcscmn	VRTSvcsc	Consolidated into VRTSvcsc.	Manual Pages for Veritas Cluster Server
VRTSvcscor	VRTSvcscsea	Consolidated into VRTSvcscsea.	Veritas High Availability Agent for Oracle
VRTSvcscsy	VRTSvcscsea	Consolidated into VRTSvcscsea.	Veritas High Availability Agent for Sybase
VRTSvcscvr	VRTSvcscag	Consolidated into VRTSvcscag.	Veritas Cluster Server Agents for Veritas Volume Replicator

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSvdid	N/A	Obsolete in 5.1.	Veritas Device Identification API
VRTSveki	VRTSveki	No change.	Veritas Kernel Interface
VRTSvlic	VRTSvlic	No change.	Symantec License Utilities
VRTSvmdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	User Documentation for Veritas Volume Manager
VRTSvmman	VRTSvxvm	Consolidated into VRTSvxvm.	Manual Pages for Veritas Volume Manager
VRTSvmpro	N/A	Consolidated into VRTSob.	Veritas Volume Manager Management Services Provider
VRTSvmweb	N/A	Obsolete in 5.0MP3. Delivered with SFM release.	Veritas Volume Manager Management Services Web Client Extensions
VRTSvrdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	User Documentation for Veritas Volume Replicator
VRTSvrpro	N/A	Consolidated into VRTSob.	Veritas Volume Replicator Management Services Provider
VRTSvrw	N/A	Obsolete in 5.1. Delivered with SFM release.	Veritas Volume Replicator Web Console
VRTSsvsc	N/A	Obsolete in 5.0MP3.	Veritas Volume Server and Client Provider
VRTSvxfen	VRTSvxfen	No change.	Veritas I/O Fencing

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSvxfs	VRTSvxfs	Expanded to include VRTSfsman (man pages). On Linux: VRTSvxfs-common and VRTSvxfs-platform packages are consolidated into single VRTSvxfs package.	Veritas File System
VRTSvxmsa	N/A	Obsolete in 5.1. Functionality dropped.	Veritas VxMS Mapping Service, Application Libraries
VRTSvxvm	VRTSvxvm	Expanded to include VRTSvxman (man pages). On Linux: VRTSvxvm-common and VRTSvxvm-platform packages are consolidated into single VRTSvxvm package.	Veritas Volume Manager binaries
VRTSweb	N/A	Obsolete in 5.1.	Symantec Web Server

Changes to I/O fencing feature

This section lists the new features and changes related to the I/O fencing feature.

Support for Coordination Point server

This release adds 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.

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

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 does not support sharing of coordinator disks across multiple clusters.

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

Support for I/O fencing in VIO server environment

This release supports I/O fencing with N_Port ID Virtualization (NPIV) in Virtual IO (VIO) server environment.

See the hardware compatibility list for more information about supported hardware. Visit the following URL:

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

New command options for vxfcntlpre

The vxfcntlpre 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 vxfcntlconfig command

The vxfcntlconfig command now has a -W option. You can use this option to display the supported and the current I/O fencing protocol versions.

New vxfcntl_vxfnd_tmt tunable parameter

I/O fencing introduces a new tunable parameter vxfcntl_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.

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

Changes to LLT

Changes to LLT are as follows:

- 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 OS 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.

Changes to GAB

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

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 `vcscenv` 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 5.1 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 Guide* for more details.

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.

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.

All nodes in the cluster must run the same VCS version. Each node in the cluster may run a different version of the operating system, as long as the operating system is supported by the VCS version in the cluster.

Before you install your Symantec products on AIX 6.1, you must read the Late Breaking News TechNote:

<http://entsupport.symantec.com/docs/282024>

Supported hardware

The compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit 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 AIX operating systems

This section lists the supported operating systems for this release of Veritas products.

Product installation scripts verify the required update levels. The installation process terminates if the target systems do not meet the maintenance level requirements.

For any Veritas cluster product, all nodes in the cluster must have the same operating system version and update level.

The minimum system requirements for this release are as follows:

AIX 5.3 at one of the following levels:

- TL7 with SP6
- TL8 with SP4
- or any higher TLs.

AIX 6.1 at one of the following levels:

- TL0 with SP6
- TL1 with SP2
- or any higher TLs.

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/334998>

Storage Foundation and High Availability Solutions 5.1 patches

Symantec strongly recommends that you install Storage Foundation and High Availability Solutions (SFHA) 5.1 Patch 1 immediately after you install SFHA 5.1.

The patch for AIX is available at the following URL:

<https://vos.symantec.com/patch/detail/2956>

Supported software

VCS supports the following volume managers and files systems:

- Logical Volume Manager (LVM)
- Journaled File System (JFS) and Enhanced Journaled File System (JFS2) on LVM

- Veritas Storage Foundation (SF): Veritas Volume Manager (VxVM) with Veritas File System (VxFS)

VCS 5.1 supports the following versions of SF:

- SF 5.0 MP3
 - VxVM 5.0 with VxFS 5.0
 - VxVM 5.0 MP1 with VxFS 5.0 MP1
 - VxVM 5.0 MP1 Update 1 with VxFS 5.0 MP1 Update1 (AIX 6.1 only)
 - VxVM 5.0 MP3 with VxFS 5.0 MP3
- SF 5.1
 - VxVM 5.1 with VxFS 5.1

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

Supported VCS agents

The Veritas agents for enterprise applications released with version 5.1 (VRTSvcsea) support VCS 5.1.

Veritas agents support a specified application version on AIX if the application vendor supports that version on AIX.

[Table 1-2](#) lists the agents for enterprise applications and the software that the agents support.

Table 1-2 Supported software for the VCS agents for enterprise applications

Agent	Application	Application version	AIX version
DB2	DB2 Enterprise Server Edition	8.1, 8.2 9.1, 9.5, 9.7	AIX 5.3, AIX 6.1
Oracle	Oracle	9i 10g R1 10g R2 11g R1	AIX 5.3, AIX 6.1
Sybase	Sybase Adaptive Server Enterprise	12.5.x, 15.x	AIX 5.3, AIX 6.1

See the Installation and Configuration 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.

No longer supported agents and components

VCS no longer supports the following:

- For a list of removed attributes:
See “[Removed attributes](#)” on page 23.
- Configuration wizards
- ServiceGroupHB agent.
This release does not support disk heartbeats. Symantec recommends using I/O fencing.
- Disk heartbeats (GABDisk).
This release does not support disk heartbeats. Symantec recommends using I/O fencing.
- The updated Oracle agent does not support Oracle 8.0.x and Oracle 8.1.x.
- The updated DB2 Agent does not support DB2 7.2
- VCS documentation package (VRTSvcscd)
The VCS documentation package (VRTSvcscd) 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.

Fixed issues

The following presents incidents that have been fixed.

Fixed issues for VCS 5.1

The following issues are fixed for VCS 5.1.

[Table 1-3](#) lists the fixed issues for VCS 5.1.

Table 1-3 VCS 5.1 fixed issues

Incident	Description
1836922	The application agent dumped userinfo to the engine log during every monitor cycle when monitorProgram is configured with non- root user.
1830978	Had may crash while sending notifications to notifier if the send fails. This happens due to incorrect data-access.
1829180	The VCSAG_SU() function from vcsag_i18n_inc.sh file, has incorrect options to execute the su command.
1808754	When bringing the Sybase resource online, the SqlTest.pl generates error in the agent log.
1789808	Cluster does not accept HA commands without reboot of whole cluster.
1780722	For a SAMBA GROUP, "netbios" resource, with CIDR address for interface, fails to come ONLINE.
1710470	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.
1634031	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.
1599129	Fixed an issue due to which the service group could not fail over after node panic.
1588784	VCS engine does not support system names starting with numbers.
1587173	The LC_ALL value was set to empty string by the hastart script even though it was not present in the environment.
1539089	The agent framework leaked memory if there is continuous logging into agent's log file.
1537433 1403471	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.
1531720	Seeding of a port a does not seed other ports.

Table 1-3 VCS 5.1 fixed issues (*continued*)

Incident	Description
1531512	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.
1479349	On disabling the array side switch port, the system panics intermittently.
1458338	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. VCS NOTICE V-16-1-53021 Diagnostics directory moved to /var/VRTSvcs/diag/had.1227075736, please check its contents and contact VERITAS Support. This was due to incorrect handling of directories.
1404384 1456802	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.
1142970	VCS logs an error "SSL Handshake failed" if the client creates a channel and then disconnects.
1403471 1397692	The VCS clients may hang in connect() call if the target system or IP is down.
1363967	In the VCS Simulator, the AIX_NFS cluster gives error while loading the LVMVolumeGroupNFS template.

Fixed issues for VCS 5.0 MP3 RP2

The following issues were fixed in VCS MP3 5.0 RP2.

Table 1-4 VCS 5.0 MP3 RP2 fixed issues

Incident	Description
1709129	[GAB] Fixed an issue in which symbolic links that were created during the installation of the GAB package were not removed during the uninstallation of the GAB package.

Table 1-4 VCS 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1703756	[VCS] Fixed an issue in which a warning message is displayed even when a parallel global group was brought online successfully. This happens because after a suitable target is determined, an internal variable is not incremented. This results in a re-visiting of the target selection algorithm, which causes error because the action is already initiated on the suitable target.
1678110	[GAB] Fixed an issue in which high priority processes running in the cluster may result in delayed response to the GAB timer function. the GAB timer function may not run as quickly as required if there are higher priority processes in the system. The priority for the gab timer is made into a tunable "gab_timer_pri". This tunable can have a value within a range.
1677412	[Agents] Fixed an issue so that when the SystemList of the service group is modified, you do not start all agents but only the required agents. The agent that was stopped by a user on a system gets restarted even if the group has no resource of that agent type, when the SystemList is modified to add that system. On SystemList modification to add new systems in SystemList, the engine starts all the agents without ensuring if the group has a resource of that type. Code changes so that only agents for which the group has resources are started whenever the SystemList is modified to add a new system.
1675815	[HAD] Fixed an issue so that the HostMonitor objects like VCShmg (Group), VCSHM (Resource), and HostMonitor (Type) are not counted in each object's number.

Table 1-4 VCS 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1672405	<p>[VCS] Fixed an issue in which a switch operation on a child service group with an OLH (Online Local Hard) and OLF (Online Local Firm) parent results in a switch of the OLH parent and the child group even though the OLF parent was online. In a situation, where two service groups depend on one child and one parent has an online local hard dependency (OLH) while the other parent has an online local firm dependency (OLF):</p> <p>The command: <code>hagrp -switch Hard_ParentSG -any</code> switches both the parents. The command: <code>hagrp -switch Hard_ParentSG -to sysB</code> switches only the hard parent group along with the child group. When the <code>hargp -switch</code> command is executed with any of the following options:</p> <p><code>hagrp -switch SG_parent -any</code></p> <p><code>hagrp -switch SG_parent -to <sys></code></p> <p>The parent group switches (while the child group is online) only in the case of a hard dependency. The switch does not happen in the case of soft or firm dependency. The switch operation succeeds for an OLH parent, if only the parent group is online. The child group has no other parents online. The OLH parent and child group can have other parents. However, the OLH child group is always a leaf node.</p>
1672335	<p>[Fencing] The <code>vxfenclearpre</code> command now outputs the progress of its execution.</p>
1670337	<p>[LLT] Added a mechanism to track the operating system timeouts that are registered by LLT with the operating system.</p>
1668609	<p>[Agents] Fixed an issue in which the Proxy agent is updated to allow the target resource to be probed before scheduling the first probe of the Proxy resource.</p>
1665036	<p>[Fencing] The <code>vxfen</code> startup process on retry no longer displays the "RFSM GAB err 16" error when the cluster is fencing a node out of the cluster.</p>
1640292	<p>[HAD] Fixed an issue to increase the resilience of HAD to extreme load conditions.</p>

Table 1-4 VCS 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1638725	[LLT] Fixed an issue in which the LLT timer function may not run as quickly as required if there are higher priority processes in the system. LLT uses the heartbeat mechanism between nodes to ensure and identify that the other node is alive. Any node in VCS/SFRAC sends heartbeat packets to all the other nodes in the cluster after every 50 millisecond. This heartbeat is sent with the LLT timer thread. Under a heavy load condition, LLT timer thread may not be scheduled to send heartbeat. If the LLT thread is on the busy node, it is not able to send a heartbeat for 16 seconds. The other node considers the busy node failed and this results in panic whenever the load of the busy node goes down and it starts communicating with the other node of cluster. The LLT heartbeat code has been moved from an llt thread context to a timer interrupt context. This ensures that the heartbeat is sent as soon as timer returns after 50 milliseconds. Interrupt handler will run real time and this removes scheduling delays.
1638240	[Agents] Fixed an issue in which the Sybase agent is unable to bring the Sybase resource online if the RUN_<servername> file is moved to some other (non default) location. The non default location for the Sybase dataserver RUN_<servername> file is not supported by the Sybase agent. Hence, if you move the RUN_<servername> file to some other location, the agent is unable to bring the Sybase resource online. A new attribute named Run_ServerFile of type string was introduced for the the Sybase and SybaseBk agents. The value of this attribute can be set to the absolute path of the RUN_<servername> file.
1635044	[Fencing] The vxfen scripts now determine the location of ssh and scp from the environment variables if sshor scp are not available in the standard path.
1634924	[VCS] Fixed an issue in which the engine logs indicated CPU usage even after the HostMonitor resource is deleted.
1633973	[VCS] Fixed an issue in which the node does not test the Authority attribute before bringing the faulted service group online, leading to concurrency violations and the service group being taken offline on the disaster recovery site.
1631012	[LLT] Fixed an issue in a configuration with LLT over UDP4, in which when an ip address is not plumbed on a link and an IP address is specified in llttab file, lltconfigwould configure that link successfully when it should not.
1603120	[VCS] Fixed an issue where NFSRestart triggers were called despite no configured NFSRestart resources, which was detrimental to performance.

Table 1-4 VCS 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1600786	<p>[Fencing] Fixed an issue in which I/O errors occur in case of a network partition at any point when the keys on the coordinator disks are being refreshed using the vxfenswap command. If the keys on coordinator disks are accidentally cleared, they can be refreshed using the vxfenswap command. However if there is a network partition at a particular point in the operation, it could result in I/O errors. If the keys that are registered on the coordinator disks are lost, the cluster may panic when a split-brain occurs. Using the vxfenswap script to replace the coordinator disks with the same disks will register the missing keys again without any risk of data corruption. However there is a possibility of seeing I/O errors because the algorithm registers the keys in the modify phase and if there is a network partition then the register(s) could override preempt(s) without synchronization. If the vxfenswap utility is run on existing coordinator disks, then the registrations are done in the commit phase instead of the modify phase.</p>
1600484	<p>[VCS] Fixed an issue so that user names are checked and validated while verifying the configuration and modifying the UserNames attribute. A user with a special character in the userid is accepted if it is the second or later user in the UserNames attribute within the main.cf file. Only the first user name is checked for valid names. If the attribute UserNames has more than one user defined in the main.cf file or the command haclus -modify UserNames u1 p1 u2 p2 is run, then even invalid user names were accepted.</p>
1598940	<p>[Agents] Fixed an issue in which the NIC and MultiNICA resources did not go online because the broadcast ping failed due to changes in the AIX operating system. Note that these agents will break if used with SP3+. Modify agent to correctly parse netstat output for both pre and post AIX 6.1 TL2 SP3 levels.</p>
1590726	<p>[VCS] Fixed an issue in which VCS generated notifications about high CPU/SWAP usage when notifications were configured. The HostMonitor feature is enhanced to give control to the user for enabling or (fully / partially) disabling the feature through the cluster object attribute - HostMonLogLvl. VCS has the HostMonitor feature enabled by default through the VCShmg group with a HostMonitor type resource VCShm. If notification is configured in VCS, you see the notifications whenever the CPU/SWAP usage is beyond critical levels. A new attribute HostMonLogLvl is added. The values can be 'ALL', 'HMAgentLog' or 'DisableHMAgent', with 'ALL' as default.</p>
1589851	<p>[GAB] Fixed the cause of a system panic that was due to depleted memory reserves.</p>

Table 1-4 VCS 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1545229	[Agents] Fixed an issue to allow control of entry point scheduling priorities and scheduling class using the new attributes EPPriority, EPClass, OnlinePriority, and OnlineClass.
1545222	[Agents] Fixed an issue to provide the ability to pass the entry point timeout value as a parameter to agent entry points in their argument list.
1544263	[Agents] Fixed an issue in which the Oracle agent performs an action corresponding to the last error even when it encounters multiple errors, thereby ignoring the previous error numbers. This happens because when the list of errors was parsed by the agent, it moved to the last error and got its state to perform the action corresponding to that error. The priority of actions are: FAILOVER, UNKNOWN, and IGNORE. If any error has FAILOVER/NOFAILOVER, the resource is FAULTED. If any error has UNKNOWN action, the resource is moved to UNKNOWN state. Else, we safely ignore the error and return the state as ONLINE.
1542334	[VCS] Fixed an issue in which the nfs_restart trigger was issuing too many hares -list commands, which impacted the response time of other HA commands invoked from the command line. The HA commands in nfs_postoffline trigger were replaced with more efficient HA commands. Also, the nfs_postoffline and nfs_preonline triggers were moved to the sample_triggers directory so that they are not invoked by default. Users are required to copy both the triggers from /opt/VRTSvcs/bin/sample_triggers to /opt/VRTSvcs/bin/triggers, if the configuration has the NFSRestart agent.
1540807	[GAB] Fixed an issue in which the error number returned by the gab_receive() function in the GAB library is wrong. The gab_receive() function returns -1, but the error number was set to 0.
1539087	[Agents] Fixed an issue in which the agent framework seems to be leaking memory during message logging.
1522568	[Agents] Fixed an issue in which the agent framework crashed while setting the resource name for the dependent attribute.
1513967/ 1638722	[LLT/GAB] Fixed an issue in which one or more nodes would halt if LLT or GAB attempted to submit timer requests during a heavy I/O workload. This leads to an abend exception and a system panic.

Table 1-4 VCS 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1509742	<p>[GAB] Fixed an issue in which GAB membership to VCS (Port h) may not occur, if VCS is brought online before the Port a membership occurs. Clients of the GAB service may not get cluster membership. Symantec recommends that GAB must be configured to provide membership only after a minimum quorum number of nodes join the cluster. If a client of GAB comes up before GAB Port a forms membership on that node, then this client may not get cluster membership until it starts up on at least the configured quorum number of nodes, not even if Port a or any other GAB Ports receive cluster membership. Previously, seeding of Port a would trigger seeding on all other ports by seeding a CONNECTS message on all those ports. However, this led to a race which was fixed via e1424927. The fix disabled CONNECTS which used to propagate the SEED bit to other ports. SEED bit is now propagated to other ports after Port 'a' reconfigures. The master for each port just runs the reconfiguration calculation after Port a reconfigures there.</p>
1504693	<p>[GAB/LLT] Fixed an issue in which LLT cannot provide backenable to GAB. This resulted in an error being produced from the GAB module gabwrite() function.</p>
1478488	<p>[Agents] Fixed an issue in which file systems failed to mount on volume sets and required one or two failovers between nodes to work correctly. When there are no Volume resources and DiskGroup resource has StartVolumes=1, all the volumes of the disk group are started by the DiskGroup agent. If there are lots of volumes/volume sets in a disk group, the volume recovery for all the volumes can take time. The file system on them cannot be mounted before the volumes are started. The fix involved modifying the DiskGroup agent to start all the volumes before initiating recovery.</p>
1377324	<p>[Agents] Fixed a parsing error which caused an error message to appear in the /var/VRTSvcs/log/tmp/Oracle-0 file.</p>
1368385	<p>[Agents] Fixed an issue in which DiskGroupSnap does not work if layered volumes are used. VxVM creates layered volumes by default, in case of larger volumes spanning multiple disks. The agent expects each volume to have a plex at each site but VxVM does not assign a site tag to plex and there is only one top level plex. Thus, the agent reports that the configuration is invalid. This was a limitation in the original agent when no layered volumes were supported.</p>

Table 1-4 VCS 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1070177	[Agents] Fixed an issue to include a new attribute to use the db2start command. There was no option to use the db2start command. Added optional attribute UseDB2start to allow users to start DB2 using the db2start command.

Fixed issues for VCS 5.0 MP3 RP1

The following issues were fixed in VCS 5.0 MP3 RP1.

Table 1-5 VCS 5.0 MP3 RP1 fixed issues

Incident	Description
1510002	Fixed an issue where the HostMonitor tool misreported CPU as busy when it is idle in an AIX 5.0 MP3 Micropartition environment.
1469381	<p>Fixed an issue where the Share agent was 10x slower on 5.0 MP1 with 300+ Share resources in a service group.</p> <p>Note: This fix changes basic VCS functionality, it is critically important for you to implement these changes for all service groups that contain NFSRestart resources.</p> <p>You must set the value of the PreOnline attribute to 1 for all service groups that contain NFSRestart resources. Failure to set the service group's PreOnline attribute to a value of 1 results in broken NFSRestart resource configurations.</p> <p>The ha commands to change this attribute are:</p> <pre># haconf -makerw # hagr -modify servicegroup_name PreOnline 1 # haconf -dump -makero"</pre>
1457429	Removed the VCS NOTICE V-16-1-53021 message after the hastart command is run.
1424927	Optimized GAB connect messages.
1414709	The hagr -offline command and hares -offline command now behave similarly when you bring the last resource in a service group offline.
1404384	Global groups can switch over to a node where WAC is not running, when PreSwitch is set to 1 and HAD runs properly.

Table 1-5 VCS 5.0 MP3 RP1 fixed issues (*continued*)

Incident	Description
1403471	Reduced time for global cluster fault detection.
1398750	Added the MemCPUAllocator agent.
1397692	Removed a condition where VCS engine clients hung in connect when the target system was down.
1396639	Return code for SCSI commands are now logged before re-using a variable.
1379299	LLT: fixed llt_recordmac() messages.
1395905	Changes implemented to close device file for device vxdmconfig.
1394624	LLT: fixed an issue where the lltctl thread spun indefinitely.

Known issues

The following issues are open for this release of VCS.

Issues related to installing and upgrading VCS

This section covers the issues related to installing and upgrading VCS.

Errors observed during partial upgrade of SFHA

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

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

Operational issues for VCS

This section covers the operational issues for VCS.

SMTP notification emails should contain the entity's name in the subject line

The SMTP notification emails should contain the subject's name in the subject line. [1836562]

VCS engine may continuously restart in secure clusters with IPv6 configurations

In IPv6 or mixed IP stack configurations, VCS engine may continuously restart in the clusters that run in secure mode. VCS looks up the `/etc/hosts` file to resolve the host names and IPv6 addresses before it looks up the DNS. If the `/etc/hosts` file does not have an IPv6 entry or if DNS does not return in 30 seconds, GAB times out and kills VCS engine. In such cases, VCS engine may continuously restart. [1826870]

Workaround: Ensure that DNS has IPv6 configured. Add the following lines to the `/etc/hosts` file.

```
127.0.0.1 localhost
::1 localhost
```

Installer fails with a message that VRTSvlic is not installed

If you run the installer to configure, upgrade, or uninstall a cluster with three or more nodes, the installer may fail with a message that the VRTSvlic fileset is not installed. This failure occurs because of an installer timing issue. [1865116]

Workaround: You can either try to start the installer script again which resolves the timing issue or you can use the `-serial` option of the installer.

Question marks in Application agent logs

Questions marks may appear in the Application agent logs. You can safely ignore them. [1220322]

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.

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.
```

HAD cannot join clusters if ulimit is low

The default ulimit or the ulimit that the root administrator sets may restrict the amount of memory that can be allocated by a process. Run the following command to verify this limit:

```
ulimit -a
      data seg size      (kbytes -d) 131072
```

In case of very large main.cf configurations, the VCS engine HAD may fail to join clusters. This could be a result of HAD not being able to allocate enough memory to facilitate transfer of main.cf across cluster nodes.

Workaround: Increase the data seg size to an appropriate value by running the `ulimit -d value` command.

Issues related to the VCS engine

This section covers the issues related to the VCS engine.

Systems with multiple CPUs and copious memory shut-down time may exceed the ShutdownTimeout attribute

The time taken by the system to go down may exceed the default value of the ShutdownTimeout attribute for systems that have a large numbers of CPUs and memory. [1472734]

Workaround: Increase the value of the ShutdownTimeout attribute based on your configuration.

The hacf -cmdtofc command generates a broken main.cf file

The `hacf -cmdtofc` command used with the `-dest` option removes the include statements from the types files. [1728738]

Workaround: Add include statements in the main.cf files that are generated using the `hacf -cmdtofc` command.

VCS engine may get stuck in LEAVING state

VCS engine may transition to the LEAVING state and may get stuck if you perform the following operations in succession:

- Run the `hares -online` command for a resource.
While the resource comes online, its child resource faults.
- Run the `hastop -local` command on the same node from where you tried to bring the resource online.

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

Parent group faulting in zone 1 results in the child group being automatically failed over to zone 2

Parent group faulting in zone 1 results in the child group being automatically failed over to zone 2. [1859387]

Issues related to the VCS bundled agents

This section covers issues related to the VCS bundled agents.

DiskGroupSnap agent assumes all nodes are part of a campus cluster configuration

The agent currently assumes that all nodes in the cluster are part of a campus cluster configuration. [1852521]

RemoteGroup agent faults when set up for monitor only the local service group is taken offline

The agent returns offline for a resource when the service group goes offline, but has not yet called an offline entry point. This faults the resource. [1851078]

Poor agent performance and inability to heartbeat to the engine

If the system has more than 200 configured resources, the agent may not get enough CPU cycles to function properly. This can prevent the agent from producing a heartbeat synchronously with the engine. If you notice poor agent performance and an agent's inability to heartbeat to the engine, check for the following symptoms.

Navigate to `/var/VRTSvcs/diag/agents/` and look for files that resemble:

```
FFDC_AGFWMMain_729_agent_type.log  FFDC_AGFWTimer_729_agent_type.log  core
FFDC_AGFWSvc_729_agent_type.log   agent_typeAgent_stack_729.txt
```

Where *agent_type* is the type of agent, for example Application or FileOnOff. If you find these files, perform the next step.

Navigate to `/var/VRTSvcs/log/` and check the `engine_*.log` file for messages that resemble:

```
2009/10/06 15:31:58 VCS WARNING V-16-1-10023 Agent agent_type
not sending alive messages since Tue Oct 06 15:29:27 2009
2009/10/06 15:31:58 VCS NOTICE V-16-1-53026 Agent agent_type
ipm connection still valid
2009/10/06 15:31:58 VCS NOTICE V-16-1-53030 Termination request sent to
agent_type agent process with pid 729
```

Workaround: If you see that both of the above criteria are true, increase the value of the AgentReplyTimeout attribute value. (Up to 300 seconds or as necessary.) [1853285]

The agent framework does not detect if service threads hang inside an entry point

In rare cases, the agent framework does not detect if all service threads hang inside a C entry point. In this case it may not cancel them successfully. [1511211]

Workaround: If the service threads of the agent are hung, send a kill signal to restart the agent. Use the following command: `$ kill -9 hung_agent's_pid`. The `haagent -stop` command does not work in this situation.

The ArgListValues attribute values for dependent resources may not populate correctly when a target resource is deleted and re-added

For resource attributes, deleting a resource prevents a dependent attribute's value from refreshing in the dependent resource's value.

For example, you have resource (*rD*), which depends on a resource's attribute value (*rT:Attr_rt*). When you delete the target resource (*rT*), and readd it (*rT*), the dependent resource (*rD*) does not get the correct value for the attribute (*Attr_rt*). [1539927]

Workaround: Set the value of the reference attribute (*target_res_name*) to an empty string.

```
# hares -modify rD target_res_name ""
```

Where *rD* is the name of the dependent resource, and *target_res_name* is the name of the reference attribute that contains the name of the target resource.

Set the value of the reference attribute (*target_res_name*) to the name of the target resource (*rT*).

```
# hares -modify rD target_res_name rT
```

Monitoring multiple IPv6 MultiNICA resources requires PingOptimize to equal 1 for one resource

If you use two IPv6 MNICA resources, set the value of the PingOptimize attribute to exactly 1 for one of the resources. When the value of PingOptimize is equal to 0 for both resources, the agent sends pings for both interfaces. In IPv6 environments, these pings can sometimes interfere with each other to create false failures. To avoid this issue, set PingOptimize to 1 for at least one resource. [1766048]

Application agent cannot monitor kernel processes

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

Issues related to the I/O fencing for VCS

This section covers the issues related to I/O fencing feature for VCS.

The vxfenclearpre script displays error messages

The vxfenclearpre script displays error messages if the /etc/vxfentab file is commented. (1512956)

The /etc/vxfentab file has the following comments:

```
#  
# /etc/vxfentab  
# DO NOT MODIFY this file it is generated by the  
# VXFEN rc script from the file /etc/vxfendg  
#
```

When you run the vxfenclearpre script, the following errors are displayed:

```
VXFEN vxfenadm ERROR V-11-2-1116 Cannot open:  
# VXFEN vxfenadm ERROR V-11-2-1132 Open of file failed, errno =  
-15344  
VXFEN vxfenadm ERROR V-11-2-1205 READ_KEYS failed for:  
# VXFEN vxfenadm ERROR V-11-2-1133 Error returned  
VXFEN vxfenadm ERROR V-11-2-1116 Cannot open:  
# VXFEN vxfenadm ERROR V-11-2-1132 Open of file failed, errno =  
-15856
```

However, the vxfenclearpre script operation is successful.

Workaround: To avoid these error messages, delete the comments from the /etc/vxfentab file before you run the vxfenclearpre script.

The vxfenswap and the vxfentsthdw utilities fail when rsh or ssh communication is not set to the same node

The vxfenswap and the vxfentsthdw utilities fail if you do not set up rsh or ssh communication to the same node. In addition to the passwordless rsh or ssh communication requirement between the systems, these utilities also require passwordless ssh or rsh configuration to the same node. [1846387]

Workaround: Make sure you have ssh or rsh configured for passwordless logins to the node where you run these utilities.

Preexisting split brain after rebooting nodes

If I/O fencing is configured in dmp mode, the fencing driver in VCS 5.0 uses Veritas DMP to handle SCSI commands to the disk driver. This allows fencing to use Veritas DMP for access to the coordinator disks. With certain disk arrays, when paths are failed over due to a path failure, the SCSI-3 persistent reservation keys for the previously active paths are not removed. If the nodes in a cluster are all rebooted at the same time, then the cluster will not start due to a preexisting split-brain condition with a `Preexisting split brain` message. [609407]

Workaround: Use the `vxfcntlpre` script to remove the keys from the coordinator disks as well as from the data disks.

Some vxfenadm options do not work with DMP paths

Some options of the `vxfenadm` utility do not work well with DMP paths such as `/dev/vx/rdmp/sdt3`.

Workaround: Use the `-a` option to register keys instead of `-m` option for DMP paths.

The vxfenswap utility has an incorrect usage message for -n option

When you invoke the `vxfenswap` utility without the `-g` option, the utility displays a usage message of which the following line is incorrect [1847517]:

```
-n use /usr/bin/ssh for communication  
instead of the default /usr/bin/ssh
```

The correct message is:

```
-n use /usr/bin/rsh for communication  
instead of the default /usr/bin/ssh
```

Issues related to LLT

This section covers the issues related to LLT.

Unloading of LLT is broken

All operations that call for unloading LLT require the nodes to be rebooted. [1111685]

Workaround: You cannot reload LLT and dependent drivers.

lltconfig -T query command displays a partially incorrect output

When the `lltconfig -T query` command is executed, the following output is displayed. [1859023]

```
# lltconfig -T query

Current LLT timer values (.01 sec units):

    heartbeat      = 50

    heartbeatlo    = 100

    peertrouble    = 200

    peertroublelo  = 400

    peerinact      = 1600

    oos            = 10

    retrans        = 10

    service        = 100

    arp            = 30000

    arpreq         = 3000

Current LLT flow control values (in packets):

    lowwater       = 40

    highwater      = 200
```

The expected output is as follows:

```
# lltconfig -T query

Current LLT timer values (.01 sec units):

    heartbeat      = 50

    heartbeatlo    = 100
```

```
peertrouble = 200

peertroublelo= 400

peerinact = 1600

oos = 10

retrans = 10

service = 100

arp = 30000

arpreq = 3000

timetoreqhb = 1400

reqhbtime = 40

timetosendhb = 200
```

This discrepancy does not affect the functioning of LLT.

Workaround: Map the last three incorrect lines of the actual output to the expected output as follows:

Actual output	Expected output	Value
Current LLT flow control values (in packets):	timetoreqhb	200
lowwater	reqhbtime	40
highwater	timetosendhb	200

LLT port stats sometimes shows recvcnt larger than rcvbytes

With each received packet, LLT increments the following variables:

- recvcnt (increment by one for every packet)
- rcvbytes (increment by size of packet for every packet)

Both these variables are integers. With constant traffic, rcvbytes hits and rolls over MAX_INT quickly. This can cause the value of rcvbytes to be less than the value of recvcnt. [1788315]

This does not impact the LLT functionality.

LLT may incorrectly declare port-level connection for nodes in large cluster configurations

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This occurs in some corner cases even though a port is not even registered on the peer node. [1809827]

Issues related to Symantec Product Authentication Service with VCS

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

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/VRTSvcsvcs/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/VRTSvcsvcs
# /opt/VRTSvcsvcs/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
```

Issues related to global clusters

This section covers the issues related to global service groups.

Global group fails to come online on the DR site with a message that it is in the middle of a group operation

When the node that runs the global group faults, VCS internally sets the MigrateQ attribute for the group and attempts to fail over the global group within another node in the local cluster. The MigrateQ attribute stores the node name on which the group was online. If the failover within the cluster does not succeed, then VCS clears the MigrateQ attribute for the groups. However, if the groups have dependencies which are more than one-level deep, then VCS does not clear the MigrateQ attribute for all the groups. [1795151]

This defect causes VCS to misinterpret that the group is in the middle of a failover operation within the local cluster and prevents the group to come online on the DR site with the following message:

```
VCS WARNING V-16-1-51042 Cannot online group global_group.  
Group is in the middle of a group operation in cluster local_cluster.
```

Workaround: Perform the following steps on a node in the local cluster which is in the RUNNING state.

To bring the global group online on the DR site

- 1 Check whether the MigrateQ attribute is set for the global group you want to bring online on the remote cluster.

```
# hagrps -display -all | grep -i migrateq
```

This command displays the name of the faulted node on which the group was online.

- 2 Flush the global group that you want to bring online on the remote cluster.

```
# hagrps -flush global_group -sys faulted_node -clus local_cluster
```

where:

- *global_group* is the group that you want to bring online on the remote cluster.
- *faulted_node* is the node in the local cluster that hosted the global group and has faulted.
- *local_cluster* is the cluster at the local site.

The flush operation clears the node name from the MigrateQ attribute.

- 3 Bring the service group online on the remote cluster.

```
# hagrps -online global_group -any -clus remote_cluster
```

Clusters are stuck in INIT and LOST_CONN states after enabling AT after cluster migration to IPv6

Clusters are stuck in INIT and LOST_CONN states after enabling Symantec Product Authentication Service (AT) for the first time without secure WAC. [1836428/1847556]

Workaround: Reboot the clusters or restart VCS on both the clusters. The command `hastop -all` may not stop WAC. In this case, you have to manually kill WAC.

The engine.logs file receives too many log messages on the secure site in global cluster environments

When the WAC process runs in secure mode on one site, and the other site does not use secure mode, the engine.logs file on the secure site gets logs every five seconds. [1539646]

Workaround: The two WAC processes in global clusters must always be started in either secure or non-secure mode. The secure and non-secure WAC connections will flood the engine_A.logs with the above messages.

The haclus -state command displays inconsistent output in four-node cluster configurations

In four-node cluster configurations, the `haclus -state` command displays an inconsistent error message after a fault. [1179782]

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 Agent for DB2

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

All partitions fault even if there are errors on only one partition with the IndepthMonitor database

This issue occurs in an MPP environment when multiple partitions use the same database. If the Databasename attribute is changed to an incorrect value, all partitions using the database fault. [568887]

Db2udb resource faults when IndepthMonitor is configured with a Japanese database

For locales other than English, you need to add the following lines to the \$INSTHOME/sql/lib/userprofile file. [590010]

The following example adds Japanese support for AIX:

```
export LANG=Ja_JP
```

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 Cluster Server 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

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.

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.

Software limitations

The following limitations apply to this release.

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.

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]

Connecting to the database outside VCS control using sqlplus takes too long to respond

Connecting to start the database outside VCS control, using sqlplus takes more than 10 minutes to respond after pulling the public network cable. [704069]

LLT over IPv6 UDP cannot detect other nodes while VCS tries to form a cluster

LLT over IPv6 requires link-local scope multicast to discover other nodes when VCS tries to form a cluster. If multicast networking is undesirable, or unavailable in your environment, use the address of the peer nodes to eliminate the need for the multicast traffic. [1533308]

Workaround: Add the set-addr entry for each local link into the /etc/llttab file. You add the entry to specify the address of the peer nodes that are available on the corresponding peer links. For example, you add the following lines into the llttab file to specify the set-addr entry for a node. In this example, the node's IPv6 address is fe80::21a:64ff:fe92:1d70.

```
set-addr 1 link1 fe80::21a:64ff:fe92:1d70
set-arp 0
```

The operating system does not distinguish between IPv4 and IPv6 packet counts

In a dual-stack configuration, when you use packet counts and the IPv6 network is disabled, the NIC agent might not detect a faulted NIC. It might not detect a fault because while the IPv6 network is down its packet count still increases. The packet count increases because the operating system does not distinguish between the packet counts for IPv4 and IPv6 networks. The agent then concludes that the NIC is up. If you are using the same NIC device for IPv4 as well as IPv6 resources, set PingOptimize to 0 and specify a value for the NetworkHosts attribute for either the IPv6 or the IPv4 NIC resource. [1061253]

DB2 RestartLimit value

When multiple DB2 resources all start at the same time with no dependencies, they tend to interfere or race with each other. This is a known DB2 issue.

The default value for the DB2 agent RestartLimit is 3. This higher value spreads out the re-start of the DB2 resources (after a resource online failure), which lowers the chances of DB2 resources all starting simultaneously. [1231311]

HAD reports incorrect CPU utilization

API issues cause HAD to report incorrect CPU use. [1214140]

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.

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.

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/netsvc.conf
```

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

When you bring the DiksGroupSnap 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 `hagrps -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/
```

A service group that runs inside of a WPAR may not fail over when its network connection is lost

For a WPAR configuration when the WPAR root is on NFS, the WPAR service group may not fail over if the NFS connection is lost. This issue is due to an AIX operating system limitation. [1637430]

I/O fencing limitations

Limitation with RDAC driver and FAS*T* array for coordinator disks that use raw disks

For multipathing to connected storage, AIX uses the RDAC driver for FAS*T* arrays. Since it is an active/passive array, only the current active path is exposed to clients. The I/O fencing driver, vxfen, can use only a single active path and has no foreknowledge of the passive paths to the coordinator disks on an array. If the single active path fails, all nodes in the cluster lose access to the coordinator disks.

The loss of the path to the coordinator disks can potentially go unnoticed until a reboot, split brain, or any other reason that leads to a cluster membership change occurs. In any of these conditions, the cluster cannot form, and all nodes panic to prevent data corruption. No data loss occurs.

Workaround: Use DMP and specify paths to coordinator disks as DMP paths rather than raw disks to avoid this limitation.

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.

Cannot modify VxFEN tunable parameters

You cannot change the VxFEN tunable parameters due to a software limitation. [1863916]

Virtualizing shared storage using VIO servers and client partitions

AIX 5.3, with proper patches to the operating system and client partitions, is capable of running multiple virtualized partitions within a single frame. You can split the CPU, memory, and certain adapters (networking and storage), into smaller virtual units that the partitions can then use.

In an Advanced POWER™ Virtualization (APV) environment, AIX uses the VIO Server to monitor and manage the I/O paths for the virtualized client partitions. At a very high level, the VIO server provides a partition's access to storage that is external to the physical computer. The VIO server encapsulates the physical hardware into virtual adapters called virtual SCSI adapters (server adapter). On the client side, you can create virtual adapters (client adapters) that map to the server adapter and enable a partition to connect to external storage.

The VIO server provides similar mechanisms to share limited networking resources across partitions. Refer to the manual that came with your system to help set up partitions, and to configure and use the various components such as VIO server and HMC, which are integral parts of IBM's APV environment.

The minimum patch level for using VIO servers with VCS is: Fix Pack 7.1.2.0.0.

Supported storage

Refer to the IBM data sheet:

<http://techsupport.services.ibm.com/server/vios/documentation/datasheet.html>

Disk Restrictions

When using VCS in combination with VIO servers and their client partitions, you need to ensure that no reservations are placed on the shared storage. This enables client partitions on different systems to access and use the same shared storage.

- If the shared storage is under MPIO control, set the `reserve_policy` attribute of the disk to `no_reserve`.
- If the shared storage is not under MPIO control, look up the array documentation to locate a similar attribute to set on the disk.

Internal testing on EMC disks shows that this field maps as the `reserve_lock` attribute for EMC disks. In this case, setting it to `no` achieves the same result.

Accessing the same LUNs from Client Partitions on different Central Electronics Complex (CEC) modules

This section briefly outlines how to set shared storage so that it is visible from client partitions on different CEC modules.

With the VIO server and client partitions set up and ready, make sure that you have installed the right level of operating system on the client partitions, and that you have mapped the physical adapters to the client partitions to provide access to the external shared storage.

To create a shareable diskgroup, you need to ensure that the different partitions use the same set of disks. A good way to make sure that the disks (that are seen from multiple partitions) are the same is to use the disks serial numbers, which are unique.

Run the following commands on the VIO server (in non-root mode), unless otherwise noted.

Get the serial number of the disk of interest:

```
$ lsdev -dev hdisk20 -vpd
hdisk20
U787A.001.DNZ06TT-P1-C6-T1-W500507630308037C-
L401 0401A00000000 IBM FC 2107
```

```
Manufacturer.....IBM
Machine Type and Model.....2107900
Serial Number.....7548111101A
EC Level.....131
Device Specific. (Z0).....10
Device Specific. (Z1).....0100
...
```

Make sure the other VIO server returns the same serial number. This ensures that you are viewing the same actual physical disk.

List the virtual SCSI adapters.

```
$ lsdev -virtual | grep vhost
vhost0 Available Virtual SCSI Server Adapter
vhost1 Available Virtual SCSI Server Adapter
```

Note: Usually vhost0 is the adapter for the internal disks. vhost1 in the example above maps the SCSI adapter to the external shared storage.

Prior to mapping hdisk20 (in the example) to a SCSI adapter, change the reservation policy on the disk.

```
$ chdev -dev hdisk20 -attr reserve_policy=no_reserve
hdisk20 changed
```

For hdisk20 (in the example) to be available to client partitions, map it to a suitable virtual SCSI adapter.

If you now print the reserve policy on hdisk20 the output resembles:

```
$ lsdev -dev hdisk20 attr reserve_policy
value
no_reserve
```

Next create a virtual device to map hdisk20 to vhost1.

```
$ mkvdev -vdev hdisk20 -vadapter vhost1 -dev mp1_hdisk5
mp1_hdisk5 Available
```

Finally on the client partition run the `cfgmgr` command to make this disk visible via the client SCSI adapter.

You can use this disk (hdisk20 physical, and known as mp1_hdisk5 on the client partitions) to create a diskgroup, a shared volume, and a eventually a shared file system.

Perform regular VCS operations on the clients vis-a-vis service groups, resources, resource attributes, etc.

Bundled agent limitations

This section covers the software limitations for VCS bundled agents.

Mount resources can cause core dumps

Due to a known Solaris issue, certain system calls create memory leaks that can lead to a core dump. This happens in situations where the Mount resource's FSType attribute has a value of nfs, and is exacerbated when the resource is for a non-global zone and the value of the SecondLevelMonitor attribute is 1. [1827036]

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.

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.

VCS does not provide a bundled agent for volume sets

VCS 5.1 does not provide a bundled agent to manage and monitor Volume Manager volume sets. Problems with 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.

Limitation with intentional offline functionality of VCS agent for Oracle

The Oracle resource never faults after an intentional offline.

Intentional offline functionality of VCS agent for Oracle requires you to enable health check monitoring. The agent uses Oracle's Health Check API to find the state of the database. If the API returns a graceful shutdown for the database, then the agent marks the resource state as INTENTIONAL OFFLINE. Later if the Oracle agent's online function does not succeed, the agent does not mark the resource as FAULTED. The state remains as INTENTIONAL OFFLINE because the agent receives the database state from the API as graceful shutdown during each monitor cycle. [1805719]

Cluster Manager (Java console) limitations

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

Use the VCS 5.1 Java Console to manage clusters

Cluster Manager (Java Console) from previous VCS versions cannot be used to manage VCS 5.1 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 node 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 does not work if the hosts file contains IPv6 entries

VCS Cluster Manager 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 errata

Veritas Cluster Server Installation Guide

This section covers the additions or corrections to the *Veritas Cluster Server Installation Guide* for document version 5.1.0.

These additions or corrections may be included in later versions of the *Veritas Cluster Server Installation Guide* that can be downloaded from the Symantec Support website and sfdoccentral.

Corrections for I/O fencing procedure

Topic: Testing the disks using vxfcntl utility

Issue: The procedure in the *Veritas Cluster Server Installation Guide* has a missing step after step 3.

Use the following info for the missing step:

Enter the names of the disks that you want to check. Each node may know the same disk by a different name:

```
Enter the disk name to be checked for SCSI-3 PGR on node
IP_adrs_of_galaxy in the format:
for dmp: /dev/vx/rdmp/DiskXX
for raw: /dev/rhdiskXX
Make sure it's the same disk as seen by nodes IP_adrs_of_galaxy
and IP_adrs_of_nebula
/dev/rhdisk75
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
IP_adrs_of_nebula in the format:
for dmp: /dev/vx/rdmp/DiskXX
for raw: /dev/rhdiskXX
```


Make sure it's the same disk as seen by nodes `IP_adrs_of_galaxy`
and `IP_adrs_of_nebula`
`/dev/rhdi.s75`

If the serial numbers of the disks are not identical, then the test terminates.

Veritas Cluster Server Administrator's Guide

This section covers the additions or corrections to the *Veritas Cluster Server Administrator's Guide* for 5.1.

VCS environment variables

In the VCS environment variables section, in the VCS environment variables table, in the row that starts with `VCS_GAB_RMACTION`, replace `PANIC` (in uppercase) with `panic` (in lowercase.) The following excerpt correctly describes its use.

Controls the GAB behavior when `VCS_GAB_RMTIMEOUT` exceeds.

You can set the value as follows:

- `panic`—GAB panics the system
- `SYSLOG`—GAB logs an appropriate message

Registration monitoring

In Registration monitoring section, replace `PANIC` (in uppercase) with `panic` (in lowercase.) The following excerpt correctly describes its use.

You can control GAB behavior in this situation by setting the `VCS_GAB_RMACTION` parameter in the `vcseenv` file.

- To configure GAB to panic the system in this situation, set:

```
VCS_GAB_RMACTION=panic
```

In this configuration, killing the `HAD` and `hashadow` processes results in a panic unless you start `HAD` within the registration monitoring timeout interval.

Configuring for failovers—a typical setup

In the Configuring for failovers—a typical setup section, replace the word `zone` with `WPAR` in all instances.

ContainerOpts is a resource type attribute not a resource attribute

In VCS attributes appendix, the ContainerOpts resource is incorrectly placed in the table for resource attributes. This is a resource type attribute.

Veritas Cluster Server Bundled Agents Reference Guide

This section covers the additions or corrections to the *Veritas Cluster Server Bundled Agents Reference Guide* for 5.1.

MultiNICB agent, LinkTestRatio attribute

In the description of the optional LinkTestRatio attribute, remove or ignore the line that reads, "Symantec does not recommend setting the value to zero."

DiskGroup agent

Under the section, "Setting the noautoimport flag for a disk group" the sub-section that reads:

For VxVM version 5.0 on AIX

Should instead read:

For VxVM version 5.0 or later on AIX

Veritas Cluster Server Agent for DB2 Installation and Configuration Guide

This section covers the additions or corrections to the *Veritas Cluster Server Agent for DB2 Installation and Configuration Guide* for 5.1.

Db2udb resource type attributes

In the Resource type attributes for DB2 appendix, the Db2udb resource type attributes optional attributes table, update the description of the UseDB2start attribute to include the following information.

If the setup is MPP, you must set the value of the UseDB2start attribute to 0. For MPP, VCS uses the actual hostname and needs the ability to start and stop each partition. For this purpose, VCS uses the `db2gcf` command and requires a value of 0 for the UseDB2start attribute.

VCS documentation

The software disc contains the documentation for VCS in Portable Document Format (PDF) in the *cluster_server/docs* directory.

VCS documentation set

[Table 1-6](#) lists the documents that VCS includes.

Table 1-6 VCS documentation

Title	File name
<i>Veritas Cluster Server Installation Guide</i>	vcs_install.pdf
<i>Veritas Cluster Server Release Notes</i>	vcs_notes.pdf
<i>Veritas Cluster Server Administrator's Guide</i>	vcs_admin.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf
<i>Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide</i>	vcs_vvr_agent.pdf
<i>Veritas Cluster Server Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_agent.pdf
<i>Veritas Cluster Server Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_agent.pdf
<i>Veritas Cluster Server Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_agent.pdf

[Table 1-7](#) lists the documentation for the VCS component - Symantec Product Authentication Service.

Table 1-7 Documentation for VCS components

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

VCS manual pages

The manual pages for the VRTSllt, VRTSgab, and VRTSvcS are installed in /opt/VRTS/man. Set the MANPATH environment variable so the man(1) command can point to the VCS manual pages.

For Bourne or Korn shell (sh or ksh), type:

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

For C shell (csh or tcsh), type:

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

For more information, refer to the man(1) manual page.

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to sfha_docs@symantec.com.

Include the document's title, its document version (located on the top of page two), the chapter title, and the section title of the text.