

Veritas™ Cluster Server 6.0.4 Release Notes - Linux

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: 6.0.4

Document version: 6.0.4 Rev 0

Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

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

Support agreement resources

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

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

Documentation

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

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

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

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

About Symantec Connect

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

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

Veritas Cluster Server Release Notes

This document includes the following topics:

- [About this document](#)
- [Component product release notes](#)
- [About Veritas Cluster Server](#)
- [About Symantec Operations Readiness Tools](#)
- [Important release information](#)
- [Changes in this release](#)
- [VCS system requirements](#)
- [No longer supported](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation](#)

About this document

This document provides important information about Veritas Cluster Server (VCS) version 6.0.4 for Linux. Review this entire document before you install or upgrade VCS.

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

This is "Document version: 6.0.4 Rev 0" of the *Veritas Cluster Server Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

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

Component product release notes

In addition to reading this Release Notes document, review the component product release notes before installing the product.

Product guides are available at the following location on the software media in PDF formats:

`/docs/product_name`

Symantec recommends copying the files to the `/opt/VRTS/docs` directory on your system.

This release includes the following component product release notes:

- *Veritas Storage Foundation Release Notes* (6.0.4)
- *Veritas Cluster Server Release Notes* (6.0.4)

About Veritas Cluster Server

Veritas™ Cluster Server (VCS) by Symantec provides High Availability (HA) and Disaster Recovery (DR) for mission critical applications running in physical and virtual environments. VCS ensures continuous application availability despite application, infrastructure or site failures.

About VCS agents

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

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

The Veritas High Availability Agent Pack gives you access to agents that provide high availability for various applications, databases, and third-party storage solutions. The Agent Pack is available through Symantec™ Operations Readiness Tools (SORT). For more information about SORT, see <https://sort.symantec.com/home>.

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

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

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

About Symantec Operations Readiness Tools

[Symantec Operations Readiness Tools \(SORT\)](#) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Symantec products.

SORT can help you do the following:

- | | |
|---|--|
| Prepare for your next installation or upgrade | <ul style="list-style-type: none">■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture.■ Analyze systems to determine if they are ready to install or upgrade Symantec products.■ Download the latest patches, documentation, and high availability agents from a central repository.■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems. |
| Manage risks | <ul style="list-style-type: none">■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository.■ Identify and mitigate system and environmental risks.■ Display descriptions and solutions for hundreds of Symantec error codes. |
| Improve efficiency | <ul style="list-style-type: none">■ Find and download patches based on product version and platform.■ List installed Symantec products and license keys.■ Tune and optimize your environment. |

Note: Certain features of SORT are not available for all products. Access to SORT is available at no extra cost.

To access SORT, go to:

<https://sort.symantec.com>

Important release information

- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:
<http://www.symantec.com/docs/TECH164885>
- For the latest patches available for this release, go to:
<https://sort.symantec.com/>
- 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://www.symantec.com/docs/TECH170013>
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.

Changes in this release

This section describes the changes introduced in this release.

Support for Oracle Linux Unbreakable Enterprise Kernel

VCS is enhanced to support Oracle Linux 6 Update 4 on Unbreakable Enterprise Kernel Release 2. Earlier, VCS supported Oracle Linux with RHEL compatible kernel only.

Support for SLES11 SP3

VCS now supports SUSE Linux Enterprise Server 11 Service Pack 3.

See “[Supported Linux operating systems](#)” on page 11.

Changes in the docs directory in the release media

In the VCS 6.0.2 release, the *Veritas High Availability Agent for SAP WebAS Installation and Configuration Guide* and the *Veritas High Availability Agent for WebSphere MQ Installation and Configuration Guide* were packaged in the release media and available in the docs directory. However, in the VCS 6.0.4 release, these

two guides are not packaged in the release media. You can download these guides from the High Availability Agents page here:

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

VCS system requirements

This section describes system requirements for VCS.

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

VCS requires that all nodes in the cluster use the same processor architecture and run the same operating system version. However, the nodes can have different update levels for a specific RHEL or OEL version, or different service pack levels for a specific SLES version.

Note: The system from where you install VCS must run the same Linux distribution as the target systems.

See “[Hardware compatibility list](#)” on page 11.

See “[Supported Linux operating systems](#)” on page 11.

Hardware compatibility list

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

<http://www.symantec.com/docs/TECH170013>

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

Supported Linux operating systems

This section lists the supported operating systems for this release of Veritas products. For current updates, visit the Symantec Operation Readiness Tools Installation and Upgrade page: https://sort.symantec.com/land/install_and_upgrade.

[Table 1-1](#) shows the supported operating systems for this release.

Table 1-1 Supported operating systems

Operating systems	Levels	Kernel version	Chipsets
SUSE Linux Enterprise 11	SP2, SP3	3.0.13-0.27 3.0.76-0.11	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Linux 6	Update 4	2.6.39-400.17.1.el6uek	64-bit x86, EMT*/Opteron 4.1 64-bit only

* Extended Memory Technology

Note: Only 64-bit operating systems are supported.

Symantec supports only Oracle distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel Application Binary Interface (ABI) compatibility.

Required Linux RPMs for VCS

Make sure you install the following operating system-specific RPMs on the systems where you want to install or upgrade VCS. VCS will support any updates made to the following RPMs, provided the RPMs maintain the ABI compatibility.

[Table 1-2](#) lists the RPMs that VCS requires for a given Linux operating system.

Table 1-2 Required RPMs

Operating system	Required RPMs
OL 6	glibc-2.12-1.107.el6.i686.rpm glibc-2.12-1.107.el6.x86_64.rpm ksh-20100621-19.el6.x86_64.rpm libgcc-4.4.7-3.el6.i686.rpm libstdc++-4.4.7-3.el6.i686.rpm mksh-39-7.el6.x86_64.rpm perl-5.10.1-129.el6.x86_64.rpm

Table 1-2 Required RPMs (*continued*)

Operating system	Required RPMs
SLES 11 SP2	glibc-2.11.3-17.31.1.x86_64.rpm glibc-32bit-2.11.3-17.31.1.x86_64.rpm ksh-93u-0.6.1.x86_64.rpm libgcc46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm libstdc++46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm
SLES 11 SP3	glibc-2.11.3-17.54.1.x86_64 .rpm glibc-32bit-2.11.3-17.54.1.x86_64.rpm ksh-93u-0.18.1.x86_64.rpm libgcc_s1-32bit-4.7.2_20130108-0.15.45.x86_64.rpm libstdc++6-32bit-4.7.2_20130108-0.15.45.x86_64.rpm

Supported Linux virtualization technologies in VCS 6.0.4

Veritas Cluster Server supports Kernel-based Virtual Machine (KVM) technology for SUSE Linux Enterprise Server (SLES)

Table 1-3 SUSE system requirements

Supported architecture	<ul style="list-style-type: none"> ■ Intel 64 ■ AMD64
Minimum system requirements	<ul style="list-style-type: none"> ■ 6GB free disk space ■ 2GB of RAM
Recommended system requirements	<ul style="list-style-type: none"> ■ 6GB plus the required disk space recommended by the guest operating system per guest. For most operating systems more than 6GB of disk space is recommended ■ One processor core or hyper-thread for each virtualized CPU and one for the host ■ 2GB of RAM plus additional RAM for virtualized guests
SUSE documentation for more information	http://www.suse.com/documentation/sles11/book_kvm?page=/documentation/sles11/book_kvm/data/book_kvm.html

Table 1-4 VCS system requirements for KVM-supported SUSE Linux Enterprise Server (SLES) configurations

VCS version	6.0.4
-------------	-------

Table 1-4 VCS system requirements for KVM-supported SUSE Linux Enterprise Server (SLES) configurations (*continued*)

Supported OS version in host	SLES 11 SP2 x86_64, SLES 11 SP3
Supported OS in VM guest	SLES 11 SP2, SLES 11 SP3
Hardware requirement	Full virtualization-enabled CPU

Supported software for VCS

VCS supports the following volume managers and file systems:

- ext2, ext3, reiserfs, NFS, and bind on LVM2, raw disks, and VxVM.
- ext4 and xfs on LVM2 and raw disks

VCS supports the following versions of Veritas Storage Foundation:

Table 1-5 Supported Veritas Storage Foundation version

Operating system	Supported Storage Foundation version
SLES	VxVM 6.0.4 with VxFS 6.0.4
OL UEK	No Storage Foundation support

Supported enterprise agents

[Table 1-6](#) lists the agents for enterprise applications and the software that the agents support.

Table 1-6 Supported software for the VCS agents for enterprise applications

Agent	Application	Application version	Linux version
DB2	DB2 Enterprise Server Edition	9.5, 9.7, 10.1	SLES11
Oracle	Oracle	10gR2, 11gR1, 11gR2, 12c	SLES 11
Sybase	Sybase Adaptive Server Enterprise	12.5.x, 15.x	SLES11

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

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

No longer supported

The following features are not supported in this release of VCS products:

No longer supported agents and components

VCS no longer supports the following:

- The `configure_cps.pl` script used to configure CP server is now deprecated and is no longer supported.
- Rolling upgrade is not supported in VCS 6.0.2.

Deprecated attributes

Deprecated DiskGroup agent attribute:

- `DiskGroupType`

Fixed issues

This section includes the issues fixed since the previous major release. The fixed issues are presented in separate tables for each applicable minor release.

LLT, GAB, and I/O fencing fixed issues

This section describes the LLT, GAB, and I/O fencing issues fixed since the previous major release.

LLT, GAB, and I/O fencing fixed issues in 6.0.4

[Table 1-8](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-7 LLT, GAB, and I/O fencing fixed issues

Incident	Description
3106493	If for some reason, kernel components of the Veritas Cluster Server (VCS) software stack are stopped and restarted in quick succession, then during a restart, the cluster communication may fail.

Table 1-7 LLT, GAB, and I/O fencing fixed issues (*continued*)

Incident	Description
3137520	Low Latency Transport (LLT) detects a duplicate node ID incorrectly, even if the nodes are using different ethernet SAP values.

LLT, GAB, and I/O fencing fixed issues in 6.0.1

[Table 1-8](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-8 LLT, GAB, and I/O fencing fixed issues

Incident	Description
2708619	If you set the <code>scsi3_disk_policy</code> attribute to <code>dmp</code> , you cannot enable the Veritas fencing module (VxFEN). The VxFEN source code is updated to pick up the <code>dmp</code> device path that contains the full disk name instead of a partition or slice.
2845244	<code>vxfen</code> startup script gives error <code>grep: can't open /etc/vxfen.d/data/cp_uid_db</code> . The error comes because <code>vxfen</code> startup script tries to read a file that might not be present. This error is typically seen when starting <code>vxfen</code> for the very first time after installation.
2554167	Setting <code>peerinact</code> value to 0 in the <code>/etc/llttab</code> file floods the system log file with large number of log messages.

Installation and upgrades fixed issues

This section describes the installation and upgrade issues fixed since the previous major release.

Installation related fixed issues in 6.0.2

Table 1-9 Installation related fixed issues

Incident	Description
2622987	If a host is not reporting to any management server but <code>sfmh</code> discovery is running before you upgrade to 6.0, <code>sfmh-discovery</code> may fail to start after the upgrade.

Bundled agents fixed issues

This section describes the bundled agent issues fixed since the previous major release.

Bundled agents fixed issues in 6.0.4

[Table 1-11](#) lists the fixed issues for bundled agents.

Table 1-10 Bundled agents fixed issues

Incident	Description
3065929	VCS agent for VMwareDisks fails to detach or attach the VMware disks if the ESX host and the virtual machine are not in the same network.
3098114	Installation of VRTSvcsag package changes permission of <code>/opt/VRTSvcs/bin</code> directory to 744 from 755.
3125918	AMF driver panics the node when vxconfigd is unresponsive.
3028760	NFSRestart resource does not start NFS processes, such as statd and lockd, during online or offline operations.

Bundled agents fixed issues in 6.0.1

[Table 1-11](#) lists the fixed issues for bundled agents.

Table 1-11 Bundled agents fixed issues

Incident	Description
2850904	Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and PanicSystemOnDGLoss is set to 0.
2794175	Fire drill for Mount agent does not accurately portray a failover scenario.
2728802	If the 'httpd' binary or the 'ab' binary is not present at the location that you specified in the 'httpdDir' attribute, the Apache agent cannot perform detail monitoring or start the HTTP server.
2850905	IMF registration for Mount resource for file systems type other than VxFS and NFS should be blocked.
2850916	Mount resource does not get registered with IMF if the attributes BlockDevice and/or MountPoint have a trailing slash in their values.

Table 1-11 Bundled agents fixed issues (*continued*)

Incident	Description
2822920	DNSAgent goes to UNKNOWN state if the Top Level Domain (TLD) is more than 4 characters in length.
2679251	After stopping the VCS forcefully (hastop -local -force), if disk reservation resource is configured then user may observe a system panic.
2800021	Clean entry point fails due to discrepancy in command path on RHEL 5.
2846389	In releases prior to VCS 6.0.1, the upper bound value of FaultTolerance attribute of the CoordPoint agent was the one less than the number of coordination points. If the majority number of coordination points fault, the entire cluster panicked under network partition scenario. Therefore, the upper bound value of the FaultTolerance attribute of CoordPoint agent had to be set to less than the majority of the coordination points. Subsequent to VCS 6.0.1, the FaultTolerance attribute of CoordPoint agent is less than the majority of coordination points.

VCS engine fixed issues

This section describes the VCS engine issues fixed since the previous major release.

VCS engine fixed issues in 6.0.4

[Table 1-13](#) lists the fixed issues for VCS engine.

Table 1-12 VCS engine fixed issues

Incident	Description
3090710	The High Availability Daemon (HAD) starts and stops before the VxFEN driver configuration completes.
3112608	Resource fails to become online after switch operations fails for a service group.
3079893	The value of LastSuccess attribute of the service group equals the GlobalCounter value of the cluster if the resource faults while you online the service group. Hence the service group fails to come online.
3042450	A parent service group which is frozen and configured with online local hard dependency is brought offline when its child service group faults.

VCS engine fixed issues in 6.0.1

Table 1-13 lists the fixed issues for VCS engine.

Table 1-13 VCS engine fixed issues

Incident	Description
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility <code>hagrp</code> gives incorrect output with the "-clear", "-flush", "-state" options.
2741299	CmdSlave gets stuck in a tight loop when it gets an EBADF on a file descriptor(fd). The CmdSlave process keeps retrying on the FD and eventually dumps core.
2647049	VCS logs do not update the log time lines once the time-zone is changed. Thus, VCS keeps logging messages with old timings even after the time- zone is updated on the system.
2850906	If a group is auto-enabled, the engine clears the Start attribute even if the resource is online.
2692173	Engine does not check whether remote parent is online when <code>-nopro</code> option is selected.
2684818	If the following attributes are specified before SystemList attribute in <code>main.cf</code> , then the value got rejected when HAD started: <ul style="list-style-type: none"> ■ PreOnline ■ ContainerInfo ■ TriggersEnabled ■ SystemZones
2696056	Memory leak occurs in the engine when <code>haclus -status <cluster></code> command is run.
2746802	When failover group is probed, VCS engine clears the MigrateQ and TargetCount.
2746816	The syslog call used in <code>gab_heartbeat_alarm_handler</code> and <code>gabsim_heartbeat_alarm_handler</code> functions is not async signal safe.

Enterprise agents fixed issues

This section describes the issues fixed in enterprise agents since the previous major release.

Enterprise agents fixed issues in 6.0.1

[Table 1-14](#) lists the fixed issues for enterprise agents.

Table 1-14 Enterprise agents fixed issues

Incident	Description
1985093	Ensure that the ohasd process has an entry in the init scripts so that when the process is killed or the machine is rebooted, this automatically restarts the process.
2831044	Sybase agent script entry points must handle large process command line.

Agent framework fixed issues

This section describes the issues fixed in agent framework since the previous major release.

Agent framework fixed issues in 6.0.1

[Table 1-15](#) lists the fixed issues for agent framework.

Table 1-15 Agent framework fixed issues

Incident	Description
2660011	Resource moves to FAULTED state even if value of ManageFaults attribute is set to NONE at service group level. This will cause service group to fault if the resource is Critical.

AMF fixed issues

This section describes the issues fixed in AMF since the previous major release.

Fixed issues related AMF in 6.0.4

[Table 1-16](#) lists the AMF fixed issues.

Table 1-16 AMF fixed issues

Incident	Description
3090229	AMF driver panics the node when vxconfigd is unresponsive.
3097342	AMF driver panics the node when AMF is being stopped.

Table 1-16 AMF fixed issues (*continued*)

Incident	Description
3259682	If vxconfigd hangs, then registration thread of imfd trying to get diskgroup status from vxconfigd also hangs. Therefore, the <code>amfregister</code> command waiting for IMFD gets stuck.

Veritas Cluster Server: issues fixed in 6.0.3

[Table 1-17](#) describes the incidents that are fixed in Veritas Cluster Server in 6.0.3.

Table 1-17 Veritas Cluster Server 6.0.3 fixed issues

Incident	Description
3035052	Add support for special characters . and / in queue manager name, in the WebSphereMQ wizard.
3028989	Add support for the custom log path and data path for WebSphereMQ versions 7.0, 7.0.1 and 7.1 in a unified way.
3013962	Monitor fails to detect DB2 resource online for DB2 version 10.1.
3013940	If DB2 is installed in NON-MPP mode and UseDB2Start attribute is set to 0, we still use db2start command to start the DB2 process instead of using db2gcf command.
2964772	If you take an NFSRestart resource offline, the NFSRestart agent may unexpectedly stop NFS processes in a local container Zones.
2951467	AMF driver panics the machine when a VXFS filesystem is unmounted.
2941155	Group is not marked offline on faulted cluster in a GCO environment after a cluster failure is declared.
2937673	AMF driver panics the machine when amfstat is executed.
2861253	In vxfen driver log statement, jeopardy membership is printed as garbage.
2848009	AMF panicks the machine when an Agent is exiting
2813773	AMF driver panics the box with the message "AMF ASSERT panic: FAILED(dev_name)"
2737653	Incorrect descriptions about the RVGPrimary online script.
2736627	REMOTE CLUSTER STATE remains in INIT state and lcmp heartbeat status is UNKNOWN.

Known issues

This section covers the known issues in this release.

VMware virtual environment specific issues

vCenter Integrated Installer may fail to install VCS on OEL 6.x guests [2905806]

In a VMware environment, if you create a virtual machine by setting **Guest Operating System Version** to **Oracle Linux 4/5/6 (64-bit)**, you cannot use the vSphere Client to install Veritas Cluster Server (VCS) on the virtual machine. The installation wizard is unable to distinguish between OEL 5.x and OEL 6.x guests and therefore does not support installation on OEL 6.x guests.

Workaround: When you create the virtual machine, set the **Guest Operating System Version** to **RedHat Enterprise Linux 6 (64-bit)** and then use vCenter Integrated Installer. Alternatively, install VCS from the command line, using one of the installer-based options. For more information, see the *Veritas Cluster Server Installation Guide*.

Oracle configuration wizard may not allow to add or remove listeners [2868770]

While configuring an Oracle application, the configuration wizard fails to accept new listeners or to remove the existing listeners if you navigate back from the virtual IPs page to the listener page.

Workaround: No workaround.

Symantec High Availability tab does not report LVMVolumeGroup resources as online [2909417]

The Symantec High Availability tab does not automatically report the online status of activated LVMVolumeGroup resources in the following case:

- If you created the VCS cluster as part of the Symantec High Availability Configuration Wizard workflow.

Workaround: Start the LVMVolumeGroup resources from the Symantec High Availability tab. For more information, see the Symantec High Availability Solutions Guide for VMware.

Unable to configure application when a new LSI Logic SCSI controller is added dynamically [2937623]

In case of SLES operating system, it is possible that hot plug drivers do not get loaded and hence the operating system is unable to recognize the newly added SCSI controllers.

Workaround: Load the drivers into the kernel before attempting to configure application or add a new SCSI controller. You can add the drivers by using the command:

```
# modprobe acpihp  
# modprobe pci_hotplug
```

During snapshot reversal, VCS commands fail, and an unrelated error message may appear [2939177]

If a virtual machine is under VCS control, and you revert to the virtual machine configuration snapshot, the VCS configuration file moves to the read- only mode. During this time, if you run a VCS command, the command fails. Also, the following unrelated error message appears in the log entries:

```
Failed to delete the attribute key <value of the key>  
  of attribute DiskPaths in the VCS configuration with error 11335  
(Configuration must be ReadWrite : Use haconf -makerw)
```

Workaround: No workaround. This message is safe to ignore.

vSphere UI may not show applications and their status [2938892]

The vSphere UI may show that it cannot connect to the virtual machine and may prompt for username and password of the virtual machine. Even after entering correct username and password, vSphere UI may not show the list of applications monitored by VCS. This can happen if the VOM MH configuration has authorization problems.

Workaround: Restart the xprtld daemon on the virtual machine with the following commands:

```
/etc/init.d/xprtld stop  
/etc/init.d/xprtld start
```

Application instances are ignored if its dependent storage is configured for other instance [2966123]

Application instances are ignored if its dependent storage is already configure under VCS for other instances. The following message is logged in healthview_A.log when you try to configure a new application instance.

```
Skipping the <Application> instance: <instance_name>,  
because the storage is already configured or an application  
running under VCS control.
```

Workaround: Reconfigure all the instances that use common storage in a single run of the wizard.

VMwareDisks agent may fail to start or storage discovery may fail if SELinux is running in enforcing mode [2854053]

The VMwareDisks agent and discFinder binaries refer to the libvmwarevcs.so shared library. SELinux security checks prevent discFinder from loading the libvmwarevcs.so library, which requires text relocation. If SELinux is running in enforcing mode, these two executables may report the "Permission denied" error and may fail to execute.

Workaround: Enter the following command and relax the security check enforcement on the Symantec libvmwarevcs.so library:

```
# chcon -t textrel_shlib_t '/opt/VRTSvcs/lib/libvmwarevcs.so'
```

SSO configuration fails if the system name contains non-English locale characters

If you install the Symantec High Availability guest components on a system that has non-English locale characters in its name, then the SSO configuration between such a system and the Symantec High Availability Console host fails. (2910613)

Workaround: Ensure that the system name does not contain any non-English locale characters.

Error while unconfiguring the cluster from the Symantec High Availability tab

This issue may occur if you try to unconfigure the cluster from the Symantec High Availability tab in VMware vSphere Client. (3011461)

The cluster unconfiguration task fails with the following error:

```
Failed to stop the HAD service. Node='systemname', Error=00000425."
```


Workaround: Perform the following steps to unconfigure the cluster:

1. Forcefully stop the High Availability Daemon (HAD) process on the system using the following command:

```
taskkill /f had.exe
```

2. If HAD starts again, stop HAD using the following command:

```
hastop -local
```

3. Ensure that the HAD process is in the stopped state.
4. Launch the Cluster Configuration Wizard (VCW) and then delete the cluster using the VCW wizard flow.

Refer to the *Administrator's Guide* for more information about VCW.

Stale NFS file handle on the client across failover of a VCS service group containing LVMLogicalVolume resource (2016627)

A VCS service group for a LVM volume group will be online automatically after a failover. However, the client applications may fail or be interrupted by stale NFS file handle error.

Workaround: To avoid the stale NFS file handle on the client across service group failover, specify "fsid=" in the Options attribute for Share resources.

NFS cluster I/O fails when storage is disabled [2555662]

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

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

Issues related to installing and upgrading VCS

Issue with soft links getting deleted in a manual upgrade

While performing a manual upgrade (from 5.1 to 6.0) of the VRTSvlic RPM, some of the soft links created during your previous installation are deleted. As a result, `vxkeyless` binary is not found in its specified path.

To prevent this, use the `--nopreun` option.

For example: `rpm -Uvh --nopreun VRTSvlic-3.02.61.003-0.x86_64.rpm`

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups [2574731]

The service groups freeze due to upgrading using the product installer if you stopped the installer after the installer already stopped some of the processes and then resumed the upgrade.

Workaround:

You must unfreeze the service groups manually after the upgrade completes.

To unfreeze the service groups manually

- 1 List all the frozen service groups

```
# hagrps -list Frozen=1
```

- 2 Unfreeze all the frozen service groups:

```
# haconf -makerw  
# hagrps -unfreeze service_group -persistent  
# haconf -dump -makero
```

While upgrading the VCS stack from a version prior to VCS 5.1, reconfiguration of MultiNICA IPv4RouteOptions attribute is required (2003864)

The 5.1SP1 MultiNICA agent now uses `ip` command by default. Due to behavioral differences in `ip` and `ifconfig` commands in regards to route configuration, MultiNICA flushes routes and sets them back for the new active device. If the MultiNICA resource configuration is not intended to make use of `ifconfig` command (see table below), you must configure IPv4RouteOptions attribute in MultiNICA resource definition.

Note: RouteOptions values are used by the `route` command where as the IPv4RouteOptions value is used by the `ip route` command. The values to be configured for these two attribute are very specific to their respective commands.

Table 1-18 Whether attributes are configured and required actions that you need to perform during upgrade

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
Configured	May or may not be configured	May or may not be configured	In this case the <code>ifconfig</code> command is used. If <code>RouteOptions</code> is set, attribute value is used to add/delete routes using command <code>route</code> . As the <code>Options</code> attribute is configured, <code>IPv4RouteOptions</code> values are ignored.	No need to configure <code>IPv4RouteOptions</code> .
Not configured	May or may not be configured	Must be configured	In this case the <code>ip</code> command is used. <code>IPv4RouteOptions</code> must be configured and are used to add/delete routes using the <code>ip route</code> command. As <code>Options</code> attribute is not configured, <code>RouteOptions</code> value is ignored.	Configure <code>IPv4RouteOptions</code> and set the IP of default gateway. The value of this attribute typically resembles: <code>IPv4RouteOptions = "default via gateway_ip"</code> For example: <code>IPv4RouteOptions = "default via 192.168.1.1"</code>

Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

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

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

To resolve this issue, perform the following steps:

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

```
# vxkeyless set NONE
```
3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic`:
 - Verify if the key has `VXKEYLESS` feature Enabled using the following command:

```
# vxlicrep -k <license_key> | grep VXKEYLESS
```
 - Delete the key if and only if `VXKEYLESS` feature is Enabled.

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

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

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

The uninstaller does not remove all scripts (2696033)

After removing VCS, some of the RC scripts remain in the `/etc/rc*.d/` folder. This is due to an issue with the `chkconfig` rpm in RHEL6 and updates. You can manually remove the scripts from the `/etc/rc*.d/` folder after removing the VxVM packages.

Workaround:

Install the `chkconfig-1.3.49.3-1` `chkconfig` rpm from the RedHat portal. Refer to the following links:

<http://grokbase.com/t/centos/centos/117pfne4zz/centos-6-0-chkconfig-strange-behavior>
<http://rhn.redhat.com/errata/RHBA-2012-0415.html>

Web installer has no option to remove node from a cluster

Web Installer does not provide the option to remove node from a cluster.

Workaround: Manually remove nodes from a cluster. There is no option to remove nodes available from Web Installer or CPI.

Web installer does not ask for authentication for the same URL after the first session if the browser is still open [2509330]

If you have closed the web installer window after either installing or configuring VCS and you have other windows of the same browser open, the web installer does not ask for authentication in the subsequent sessions. Since there is no option to gracefully log out of the web installer, its session remains open as long as the browser is used by the web installer is open on the system.

However, This issue is URL-specific and is observed only when you use the same URL to perform the subsequent operations. Therefore, if you use different URLs for your purpose, the browser prompts for authentication each time you access the web installer.

Workaround: You can use different URL to access the web installer.

Web installer does not ask for authentication after the first session if the browser is still open (2509330)

If you install or configure VCS and then close the Web installer, if you have other browser windows open, the Web installer does not ask for authentication in the subsequent sessions. Since there is no option to log out of the Web installer, the session remains open as long as the browser is open on the system.

Workaround: Make sure that all browser windows are closed to end the browser session and subsequently log in again.

After finishing a kernel upgrade on a master node the cvm group on a slave node does not come online (2439439)

After successfully finishing a kernel upgrade on one node, the cvm group does not come online on the second node.

Workaround: Check that your cluster is not in a jeopardy state before you perform a rolling upgrade.

Stopping the Web installer causes Device Busy error messages (2633924)

If you start the Web installer, and then perform an operation (such as prechecking, configuring, or uninstalling), you may get an error message saying the device is busy.

Workaround: Do one of the following:

- Kill the start.pl process.
- Start the webinstaller again. On the first Web page you see that the session is still active. Either take over this session and finish it or terminate it directly.

After finishing a kernel upgrade on a master node the cvm group on a slave node does not come online (2439439)

After successfully finishing a kernel upgrade on one node, the cvm group does not come online on the second node.

Workaround: Check that your cluster is not in a jeopardy state before you perform a rolling upgrade.

Operational issues for VCS

LVM SG transition fails in all paths disabled status [2081430]

If you have disabled all the paths to the disks, the `LVM2 vg` commands stop responding and wait until at least one path to the disks is restored. As `LVMVolumeGroup` agent uses `LVM2` commands, this behavior causes online and offline entry points of `LVMVolumeGroup` agent to time out and `clean EP` stops responding for an indefinite time. Because of this, the service group cannot fail over to another node.

Workaround: You need to restore at least one path.

SG goes into Partial state if Native LVMVG is imported and activated outside VCS control

If you import and activate LVM volume group before starting VCS, the `LVMVolumeGroup` remains offline though the `LVMLogicalVolume` resource comes online. This causes the service group to be in a partial state.

Workaround: You must bring the VCS `LVMVolumeGroup` resource offline manually, or deactivate it and export the volume group before starting VCS.

Some VCS components do not work on the systems where a firewall is configured to block TCP traffic

The following issues may occur if you install and configure VCS on systems where a firewall is installed:

- If you set up Disaster Recovery using the Global Cluster Option (GCO), the status of the remote cluster (cluster at the secondary site) shows as "initing".
- If you configure fencing to use CP server, fencing client fails to register with the CP server.
- Setting up trust relationships between servers fails.

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the *Veritas Cluster Server Installation Guide* for the list of ports and services used by VCS.
- Configure the firewall policy such that the TCP ports required by VCS are not blocked. Refer to your respective firewall or OS vendor documents for the required configuration.

Switching service group with DiskGroup resource causes reservation conflict with UseFence set to SCSI3 and powerpath environment set [2749136]

If UseFence is set to SCSI3 and powerpath environment is set, then switching the service group with DiskGroup resource may cause following messages to appear in syslog:

```
reservation conflict
```

This is not a VCS issue. In case UseFence is set to SCSI3, the diskgroups are imported with the reservation. This message gets logged while releasing and reserving the disk.

Workaround: See the tech note available at

<http://www.symantec.com/business/support/index?page=content&id=TECH171786>.

Issues related to the VCS engine

Trigger does not get executed when there is more than one leading or trailing slash in the triggerpath [2368061]

The path specified in TriggerPath attribute must not contain more than one leading or trailing '/' character.

Workaround: Remove the extra leading or trailing '\' characters from the path.

Service group is not auto started on the node having incorrect value of EngineRestarted [2653688]

When HAD is restarted by `hashadow` process, the value of `EngineRestarted` attribute is temporarily set to 1 till all service groups are probed. Once all service groups are probed, the value is reset. If HAD on another node is started at roughly the same time, then it is possible that it does not reset the value of `EngineRestarted` attribute. Therefore, service group is not auto started on the new node due to mismatch in the value of `EngineRestarted` attribute.

Workaround: Restart VCS on the node where `EngineRestarted` is set to 1.

Group is not brought online if top level resource is disabled [2486476]

If the top level resource which does not have any parent dependency is disabled then the other resources do not come online and the following message is displayed:

```
VCS NOTICE V-16-1-50036 There are no enabled
resources in the group cvm to online
```

Workaround: Online the child resources of the topmost resource which is disabled.

NFS resource goes offline unexpectedly and reports errors when restarted [2490331]

VCS does not perform resource operations, such that if an agent process is restarted multiple times by HAD, only one of the agent process is valid and the remaining processes get aborted, without exiting or being stopped externally. Even though the agent process is running, HAD does not recognize it and hence does not perform any resource operations.

Workaround: Terminate the agent process.

Parent group does not come online on a node where child group is online [2489053]

This happens if the `AutostartList` of parent group does not contain the node entry where the child group is online.

Workaround: Bring the parent group online by specifying the name of the system then use the `hargp -online [parent group] -any` command to bring the parent group online.

Cannot modify temp attribute when VCS is in LEAVING state [2407850]

An `ha` command to modify a temp attribute is rejected if the local node is in a LEAVING state.

Workaround: Execute the command from another node or make the configuration read-write enabled.

If secure and non-secure WAC are connected the engine_A.log receives logs every 5 seconds [2653695]

Two WACs in GCO must always be started either in secure or non-secure mode. The secure and non-secure WAC connections cause log messages to be sent to `engine_A.log` file.

Workaround: Make sure that WAC is running in either secure mode or non-secure mode on both the clusters in GCO.

Oracle group fails to come online if Fire Drill group is online on secondary cluster [2653695]

If a parallel global service group faults on the local cluster and does not find a failover target in the local cluster, it tries to failover the service group to the remote cluster. However, if the firedrill for the service group is online on a remote cluster, offline local dependency is violated and the global service group is not able to failover to the remote cluster.

Workaround: Offline the Firedrill service group and online the service group on a remote cluster.

Service group may fail to come online after a flush and a force flush operation [2616779]

A service group may fail to come online after flush and force flush operations are executed on a service group where offline operation was not successful.

Workaround: If the offline operation is not successful then use the force flush commands instead of the normal flush operation. If a normal flush operation is already executed then to start the service group use `-any` option.

Elevated TargetCount prevents the online of a service group with `hagrps -online -sys` command [2871892]

When you initiate an offline of a service group and before the offline is complete, if you initiate a forced flush, the offline of the service group which was initiated

earlier is treated as a fault. As start bits of the resources are already cleared, service group goes to OFFLINE|FAULTED state but TargetCount remains elevated.

Workaround: No workaround.

Auto failover does not happen in case of two successive primary and secondary cluster failures [2858187]

In case of three clusters (clus1, clus2, clus3) in a GCO with steward not configured, if clus1 loses connection with clus2, it sends the inquiry to clus3 to check the state of clus2 one of the following condition persists:

1. If it is able to confirm that clus2 is down, it will mark clus2 as FAULTED.
2. If it is not able to send the inquiry to clus3, it will assume that a network disconnect might have happened and mark clus2 as UNKNOWN

In second case, automatic failover does not take place even if the ClusterFailoverPolicy is set to Auto. You need to manually failover the global service groups.

Workaround: Configure steward at a geographically distinct location from the clusters to which the above stated condition is applicable.

GCO clusters remain in INIT state [2848006]

GCO clusters remain in INIT state after configuring GCO due to :

- Trust between two clusters is not properly set if clusters are secure.
- Firewall is not correctly configured to allow WAC port (14155).

Workaround: Make sure that above two conditions are rectified. Refer to *Veritas Cluster Server Administrator's Guide* for information on setting up Trust relationships between two clusters.

The `ha` commands may fail for non-root user if cluster is secure [2847998]

The `ha` commands fail to work if you first use a non-root user without a home directory and then create a home directory for the same user.

Workaround

- 1 Delete `/var/VRTSat/profile/<user_name>`.
- 2 Delete `/home/user_name/.VRTSat`.
- 3 Delete `/var/VRTSat_lhc/<cred_file>` file which same non-root user owns.
- 4 Run `ha` command with same non-root user (this will pass).

Older ClusterAddress remains plumbed on the node while modifying ClusterAddress [2858188]

If you execute `gcoconfig` to modify ClusterAddress when ClusterService group is online, the older ClusterAddress remains plumbed on the node.

Workaround: Un-plumb the older ClusterAddress from the node manually or offline ClusterService group by executing the following command before running `gcoconfig`:

```
hagrp -offline -force ClusterService -any
```

or

```
hagrp -offline -force ClusterService -sys <sys_name>
```

Auto failover does not happen in case of two successive primary and secondary cluster failures [2858187]

In case of three clusters (clus1, clus2, clus3) in a GCO with steward not configured, if clus1 loses connection with clus2, it sends the inquiry to clus3 to check the state of clus2 one of the following condition persists:

1. If it is able to confirm that clus2 is down, it will mark clus2 as FAULTED.
2. If it is not able to send the inquiry to clus3, it will assume that a network disconnect might have happened and mark clus2 as UNKNOWN

In second case, automatic failover does not take place even if the ClusterFailoverPolicy is set to Auto. You need to manually failover the global service groups.

Workaround: Configure steward at a geographically distinct location from the clusters to which the above stated condition is applicable.

GCO clusters remain in INIT state [2848006]

GCO clusters remain in INIT state after configuring GCO due to :

- Trust between two clusters is not properly set if clusters are secure.
- Firewall is not correctly configured to allow WAC port (14155).

Workaround: Make sure that above two conditions are rectified. Refer to *Veritas Cluster Server Administrator's Guide* for information on setting up Trust relationships between two clusters.

The `ha` commands may fail for non-root user if cluster is secure [2847998]

The `ha` commands fail to work if you first use a non-root user without a home directory and then create a home directory for the same user.

Workaround

- 1 Delete `/var/VRTSat/profile/<user_name>`,
- 2 Delete `/home/user_name/.VRTSat`.
- 3 Delete `/var/VRTSat_lhc/<cred_file>` file which same non-root user owns.
- 4 Run `ha` command with same non-root user (this will pass).

Older ClusterAddress remains plumbed on the node while modifying ClusterAddress [2847997]

If you execute `gcoconfig` to modify ClusterAddress when ClusterService group is online, the older ClusterAddress remains plumbed on the node.

Workaround: Un-plumb the older ClusterAddress from the node manually or offline ClusterService group by executing the following command before running `gcoconfig`:

```
hagrp -offline -force ClusterService
```

Extremely high CPU utilization may cause HAD to fail to heartbeat to GAB [1744854]

When CPU utilization is very close to 100%, HAD may fail to heartbeat to GAB.

The `hacf -cmdtocf` command generates a broken `main.cf` file [1919951]

The `hacf -cmdtocf` command used with the `-dest` option removes the include statements from the types files.

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

Issues related to the bundled agents

System panics after starting KVM virtualized guest or initiating KVMGuest resource online [2337626]

System panics when the KVM guest is started or when the KVMGuest resource online is initiated. This issue is rarely observed.

The issue is observed due to the file descriptor leak in the libvirtd process. The maximum file open limit of file descriptor for libvirtd process is 1024. You may sometimes observe that more than 1024 file descriptors are opened when the KVM guest is started. Therefore, if the maximum file open limit is crossed, any attempt to start the KVM guest or to open a new file causes the system to panic. VCS cannot control this behavior as it suspects a file descriptor leak in the libvirtd process.

Workaround: There is no definite resolution for this issue; however, you can check the number of files opened by the libvirtd process in `/proc/<pid of libvirtd>/fd/`. If the file count exceeds 1000, restart libvirtd with the following command:

```
/etc/init.d/libvirtd restart
```

KVMGuest monitor entry point reports resource ONLINE even for corrupted guest or with no OS installed inside guest [2394235]

The VCS KVMGuest monitor entry point reports resource state as ONLINE in spite of the operating system inside the guest being corrupted or even if no operating system is installed inside the guest. The VCS KVMGuest agent uses `virsh` utility to determine the state of the guest. When the guest is started, the `virsh` utility reports the state of the running guest as running. Based on this running state, VCS KVMGuest agent monitor entry point reports the resource state as ONLINE.

In case the operating system is not installed inside the guest or the installed operating system is corrupted, `virsh` utility still reports the guest state as running. Thus, VCS also reports the resource state as ONLINE. Since RedHat KVM does not provide the state of the operating system inside guest, VCS cannot detect the guest state based on the state of the operating system.

Workaround: No workaround for this known issue.

LVM logical volume may get stuck with reiserfs file system on SLES11 [2120133]

LVM logical volume may get stuck with reiserfs file system on SLES11 if the service group containing the logical volume is switched continuously between the cluster node.

This issue may be observed:

- During the continuous switching of the service group having the LVM logical volume with reiserfs file system.
- On SLES11 and with reiserfs file system only.
- Due to the behavior of device-mapper on SLES11.

However, the issue is not consistent. Sometimes, the device-mapper gets stuck while handling the logical volumes and causes the logical volume to hang. In such a case, LVM2 commands also fail to clear the logical volume. VCS cannot handle this situation as the LVM2 commands are unable to deactivate the hung logical volume.

Resolution: You must restart the system on which the logical volumes are stuck in this situation.

KVMGuest resource comes online on failover target node when started manually [2394048]

The VCS KVMGuest resource comes online on failover target node when VM guest started manually, even though the resource is online on the primary node.

Kernel-based virtual machine (KVM) allows you to start the guest using same guest image on multiple nodes. The guest image is residing on the cluster file system. If the guest image is stored on the cluster file system, then it becomes available on all the cluster nodes simultaneously.

If the KVMGuest resource of VCS has made the guest online on one node by starting it using the guest image on cluster file system and if you manually start the same guest on the other node, Red Hat KVM does not prevent you from doing so. However, as this particular guest is under VCS control, VCS does not allow the resource to be ONLINE on multiple nodes simultaneously (unless it is in parallel service group configuration). VCS detects this concurrency violation and brings down the guest on the second node.

Note: This issue is also observed with CVM raw volume.

Workaround: No workaround required in VCS. VCS concurrent violation mechanism handles this scenario appropriately.

Application agent cannot handle a case with user as root, envfile set and shell as csh [2490296]

Application agent does not handle a case when the user is root, envfile is set, and shell is csh. The application agent uses the `system` command to execute the `Start/Stop/Monitor/Clean` Programs for the root user. This executes `Start/Stop/Monitor/Clean` Programs in `sh` shell, due to which there is an error when root user has `csh` shell and `EnvFile` is written accordingly.

Workaround: Do not set `csh` as shell for root user. Use `sh` as shell for root instead.

DiskReservation agent may call clean if you configure large number of its resources in a single service group [2336391]

DiskReservation agent may call clean if you configure large number of DiskReservation resource (more than 400 resources) in a single service group and try to offline the service group.

In a single service group configuration with more than 400 DiskReservation resources and equal number of Mount resources, the service group offline may cause the DiskReservation agent to call clean entry point. This issue is not observed if you configure about 150 resources.

Workaround: No workaround.

IMF registration fails for Mount resource if the configured MountPoint path contains spaces [2442598]

If the configured MountPoint of a Mount resource contains spaces in its path, then the Mount agent can online the resource correctly, but the IMF registration for ONLINE monitoring fails. This is due to the fact that the AMF driver does not support spaces in the path. Leading and trailing spaces are handled by the Agent and IMF monitoring can be done for such resources.

Workaround: Symantec recommends to turn off the IMF monitoring for a resource having spaces in its path. For information on disabling the IMF monitoring for a resource, refer to Veritas Cluster Server Administrator's Guide.

DiskGroup agent is unable to offline the resource if volume is unmounted outside VCS

DiskGroup agent is unable to offline the resource if volume is unmounted using the `umount -l` command outside VCS.

A service group contains DiskGroup, Volume and Mount resources and this service group is online. Volume is mounted by Mount resource with `VxFSMountLock` enabled. An attempt to manually unmount the volume using `umount -l system`

command causes the mount point to go away; however, the file system lock remains as it is. The volume cannot be stopped as it is mount locked and hence the disk group cannot be imported. This causes the disk group resource to go into UNABLE to OFFLINE state. Also, any attempt to again mount the file system fails, because it is already mount locked. This issue is due to file system behavior on Linux.

Workaround: Do not use `umount -l` command to unmount the VxFS file system when the mount lock is enabled. Instead, first unlock the mount point using the `/opt/VRTS/bin/fsadm` command and then unmount the file system.

RemoteGroup agent does not failover in case of network cable pull [2588807]

A RemoteGroup resource with ControlMode set to OnOff may not fail over to another node in the cluster in case of network cable pull. The state of the RemoteGroup resource becomes UNKNOWN if it is unable to connect to a remote cluster.

Workaround:

- Connect to the remote cluster and try taking offline the RemoteGroup resource.
- If connection to the remote cluster is not possible and you want to bring down the local service group, change the ControlMode option of the RemoteGroup resource to MonitorOnly. Then try taking offline the RemoteGroup resource. Once the resource is offline, change the ControlMode option of the resource to OnOff.

VVR setup with FireDrill in CVM environment may fail with CFSSMount Errors [2564411]

When you try to bring the FireDrill service group online through Java Console or `hagrps -online` command, the CFSSMount resource goes into faulted state.

Workaround: Run the `fsck` command.. You can find these commands in the engine logs.

CoordPoint agent remains in faulted state [2852872]

The CoordPoint agent remains in faulted state because it detects `rfsm` to be in replaying state.

Workaround: After HAD has stopped, reconfigure fencing.

RVGsnapshot agent does not work with volume sets created using vxvset [2553505]

RVGsnapshot agent does not work with volume sets created using `vxvset`. This happens during FireDrill in a VVR environment.

Workaround: No workaround.

No log messages in engine_A.log if VCS does not find the Monitor program [2563080]

No message is logged in the `engine_A.log`, when VCS cannot find the Monitor program with KVM guest with service group online.

Workaround: In case resource state is unknown, also refer to agent log files for messages.

No IPv6 support for NFS [2022174]

IPv6 is not supported for NFS.

Workaround: No workaround.

Some agents may fail to come online after full upgrade to VCS 6.0 if they were online before the upgrade [2618482]

Resources of type NFSRestart, DNS and LogicalVolumeGroup do not come online automatically after a full upgrade to VCS 6.0 if they were previously online.

Workaround: Online the resources manually after the upgrade, if they were online previously.

SambaShare agent clean entry point fails when access to configuration file on shared storage is lost [2858183]

When the Samba server configuration file is on shared storage and access to the shared storage is lost, SambaShare agent clean entry point fails.

Workaround: No workaround.

SambaShare agent fails to offline resource in case of cable pull or unplumbing of IP [2848020]

When IP is unplumbed or in case of cable pull scenario, agent fails to offline SambaShare resource.

Workaround: No workaround.

Concurrency violation observed during migration of monitored virtual machine [2755936]

If a VCS service group has more than one KVMGuest resource monitoring virtual machine and one of the virtual machines is migrated to another host, a service group level concurrency violation occurs as the service group state goes into PARTIAL state on multiple nodes.

Workaround: Configure only one KVMGuest resource in a Service group.

KVMGuest agent fails to recognize paused state of the VM causing KVMGuest resource to fault [2796538]

In a SUSE KVM environment, when a virtual machine is saved, its state is changed to paused and then shut-off. The paused state remains for a very short period of time, due to timing in case that the KVMGuest agent misses this state. Then the resource state will be returned as OFFLINE instead of INTENTIONAL OFFLINE, which causes the KVMGuest resource to fault and failover.

This is due to the limitation of SUSE KVM as it does not provide a separate state for such events.

Workaround: No workaround.

Concurrency violation in the service group [2870982]

Concurrency violation and data corruption of a volume resource may occur if storage connectivity is lost or all paths under VxDMP are disabled and PanicSystemOnDGLoss is set to 0 This happens when:

- In a cluster configuration, if cluster-wide UseFence attribute is set to SCSI3 and service group contains Volume resource and DiskGroup resource with the PanicSystemOnDGLoss attribute set to 0 (zero).
- If storage connectivity is lost or all paths under VxDMP are disabled, VCS fails over the service group. If storage connectivity is restored on the node on which the service group was faulted and disk group is not deported manually, then volume may get started if disk group is not deported during the service group failover. Thus, volume resource shows its state as online on both the nodes and thus causes concurrency violation. This may lead to data corruption.

Workaround: Ensure that the disk group is deported soon after storage connectivity is restored. Always configure volume resource whenever disk group resource is configured and set the attribute PanicSystemOnDGLoss to 1 or 2 as required.

Resource does not appear ONLINE immediately after VM appears online after a restart [2735917]

During a VM restart the resource does not come ONLINE immediately after the VM starts running. As the VM state is 'Reboot in Progress' it reports INTENTIONAL OFFLINE and after VM is UP the resource cannot immediately detect it as the next monitor is scheduled after 300 seconds.

Workaround: Reduce the OfflineMonitorInterval and set it to suitable value.

Logical volume resources fail to detect connectivity loss with storage when all paths are disabled in KVM guest [2871891]

In a KVM environment if all storage paths are disabled, then LVMLogicalVolume and LVMVolumeGroup resources fails to detect the loss of connectivity with storage. This is because even if the storage paths are disabled, native LVM commands return success, which causes VCS to report the resource state ONLINE and hence no SG failover is initiated. If any application monitored by VCS is doing a read/write I/O to this volumes, then it can detect the fault and initiate service group failover.

Workaround: No workaround.

LVM Logical Volume will be auto activated during I/O path failure [2140342]

LVM Logical Volume gets auto activated during the I/O path failure. This causes the VCS agent to report "Concurrency Violation" errors, and make the resource groups offline/online temporarily. This is due to the behavior of Native LVM.

Workaround: Enable the LVM Tagging option to avoid this issue.

VMwareDisks resource faults during data store migration of VM or during creation, deletion, or reversal of the VM snapshot (3318335)

VMwareDisks agent updates its DiskPaths attribute when the underlying disk path changes in VMware because of data store migration or snapshot creation of the virtual machine. In some environments where virtual machine snapshots are frequently created and deleted, stale entries of previous disk paths may persist in the DiskPaths attribute. This leads to multiple entries in the attribute. Such multiple entries with the same SCSI target, port and UUID cause the VMwareDisks resource to fault.

Workaround: No workaround.

System having DiskReservation resource online panics with the loss of storage connectivity (3321322)

If the storage connectivity is lost on the system where DiskReservation resource is online, that system panics.

Workaround: No workaround.

Issues related to the VCS database agents

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 [1985055]

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

Workaround: Set MonitorOption attribute for Oracle resource to 0.

Health check monitoring does not work with VCS agent for Oracle [2101432]

The health check monitoring in Oracle agent for VCS does not work due to incompatibility of the health check APIs provided by Oracle.

Workaround: Disable health check monitoring by setting the MonitorOption attribute to 0 (zero).

Intentional Offline does not work for VCS agent for Oracle [1805719]

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

No health check monitoring for Oracle agent on SLES11 platform [1919203]

Oracle agent does not support health check monitoring on SLES11 platform.

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

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

Workaround:

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

VCS agent for ASM: Health check monitoring is not supported for ASMinst agent

The ASMinst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

NOFAILOVER action specified for certain Oracle errors

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

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

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

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

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

Issues related to the agent framework

Agent may fail to heartbeat under heavy load [2073018]

An agent may fail to heartbeat with the VCS engine under heavy load.

This may happen when agent does not get enough CPU to perform its tasks and when the agent heartbeat exceeds the time set in the AgentReplyTimeout attribute. The VCS engine therefore stops the agent and restarts it. The VCS engine generates a log when it stops and restarts the agent.

Workaround: If you are aware that the system load is likely to be high, then:

- The value of AgentReplyTimeout attribute can be set to a high value
- The scheduling class and scheduling priority of agent can be increased to avoid CPU starvation for the agent, using the AgentClass and AgentPriority attributes.

Agent framework cannot handle leading and trailing spaces for the dependent attribute (2027896)

Agent framework does not allow spaces in the target resource attribute name of the dependent resource.

Workaround: Do not provide leading and trailing spaces in the target resource attribute name of the dependent resource.

IMF related error messages while bringing a resource online and offline [2553917]

For a resource registered with AMF, if you run `hagrps -offline` or `hagrps -online` explicitly or through a collective process to offline or online the resource respectively, the IMF displays error messages in either case.

The errors displayed is an expected behavior and it does not affect the IMF functionality in any manner.

Workaround: No workaround.

Issues with configuration of resource values (1718043)

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

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

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

The agent framework does not detect if service threads hang inside an entry point [1442255]

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.

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.

Issues related to global clusters

Application group attempts to come online on primary site before fire drill service group goes offline on the secondary site (2107386)

The application service group comes online on the primary site while the fire drill service group attempts to go offline at the same time, causing the application group to fault.

Workaround: Ensure that the fire drill service group is completely offline on the secondary site before the application service group comes online on the primary site.

The engine log file receives too many log messages on the secure site in global cluster environments [1919933]

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

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 log file with the above messages.

LLT known issues

This section covers the known issues related to LLT in this release.

LLT connections are not formed when a vlan is configured on a NIC (2484856)

LLT connections are not formed when a vlan is configured on a NIC that is already used to configure an LLT link.

Workaround: Do not specify the MAC address of a NIC in the `llttab` file while configuring LLT if you want to configure a vlan later. If you have already specified the MAC address of a NIC, then delete the MAC address from the `llttab` file, and update the file before you restart LLT.

LLT may fail to detect when bonded NICs come up (2604437)

When LLT is configured over a bonded NIC and that bonded NIC is DOWN with the `ifconfig` command, LLT marks the corresponding link down. When the bonded

NIC is UP again using the `ifconfig` command, LLT fails to detect this change and marks the link up.

Workaround: Close all the ports and restart LLT, then open the ports again.

Cannot use CPI response files to add nodes to a cluster that is using LLT over UDP (2869763)

When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/llttab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files.

Workaround: None

LLT port stats sometimes shows recvcnt larger than recvbytes (1907228)

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

LLT may incorrectly declare port-level connection for nodes in large cluster configurations [1810217]

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.

GAB known issues

This section covers the known issues related to GAB in this release.

While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373)

After you unregister the `gtx` port with `-nodeinit` option, the `gabconfig -C` command shows `refcount` as 1. But when forceful `deinit` option (`gabdebug -R`

`GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinited on user request
```

The `refcount` value is incremented by 1 internally. However, the `refcount` value is shown as 2 which conflicts with the `gabconfig -C` command output.

Workaround: There is no workaround for this issue.

Cluster panics during reconfiguration (2590413)

While a cluster is reconfiguring, GAB broadcast protocol encounters a race condition in the sequence request path. This condition occurs in an extremely narrow window which eventually causes the GAB master to panic.

Workaround: There is no workaround for this issue.

I/O fencing known issues

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

Installer is unable to split a cluster that is registered with one or more CP servers (2110148)

Splitting a cluster that uses server-based fencing is currently not supported.

You can split a cluster into two and reconfigure VCS on the two clusters using the installer. For example, you can split a cluster `clus1` into `clus1A` and `clus1B`.

However, if you use the installer to reconfigure the VCS, the installer retains the same cluster UUID of `clus1` in both `clus1A` and `clus1B`. If both `clus1A` and `clus1B` use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: There is no workaround for this issue.

CoordPoint agent does not report the addition of new disks to a Coordinator disk group [2727672]

The LevelTwo monitoring of the CoordPoint agent does not report a fault even if the constituent of a coordinator disk group changes due to addition of new disks in the coordinator disk group

Workaround: There is no workaround for this issue.

Fencing does not come up on one of the nodes after a reboot (2573599)

If VxFEN unconfiguration has not finished its processing in the kernel and in the meantime if you attempt to start VxFEN, you may see the following error in the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1007 Vxfen already configured
```

However, the output of the `gabconfig -a` command does not list port b. The `vxfenadm -d` command displays the following error:

```
VXFEN vxfenadm ERROR V-11-2-1115 Local node is not a member of cluster!
```

Workaround: Start VxFEN again after some time.

The vxfenswap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)

The `vxfenswap` utility runs the `vxfenconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxfenswap` command using RSH (with the `-n` option), then RSH does not detect the failure of validation of coordination points on a node. From this point, `vxfenswap` proceeds as if the validation was successful on all the nodes. But, it fails at a later stage when it tries to commit the new coordination points to the VxFEN driver. After the failure, it rolls back the entire operation, and exits cleanly with a non-zero error code. If you run `vxfenswap` using SSH (without the `-n` option), then SSH detects the failure of validation of coordination of points correctly and rolls back the entire operation immediately.

Workaround: Use the `vxfenswap` utility with SSH (without the `-n` option).

In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)

When you start server-based I/O fencing, the node may not join the cluster and prints error messages in logs similar to the following:

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1043  
Detected a preexisting split brain. Unable to join cluster.
```

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
operation failed.  
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@sys1,  
domaintype vx; not allowing action
```

The `vxfsend` daemon on the application cluster queries the coordination point server (CP server) to check if the cluster members as seen in the GAB membership are registered with the CP server. If the application cluster fails to contact the CP server due to some reason, then fencing cannot determine the registrations on the CP server and conservatively assumes a pre-existing split-brain.

Workaround: Before you attempt to start VxFEN on the application cluster, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

The `cpsadm` command fails if LLT is not configured on the application cluster (2583685)

The `cpsadm` command fails to communicate with the coordination point server (CP server) if LLT is not configured on the application cluster node where you run the `cpsadm` command. You may see errors similar to the following:

```
# cpsadm -s 10.209.125.200 -a ping_cps  
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using  
environment variable  
CPS_NODEID  
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

However, if you run the `cpsadm` command on the CP server, this issue does not arise even if LLT is not configured on the node that hosts CP server. The `cpsadm` command on the CP server node always assumes the LLT node ID as 0 if LLT is not configured.

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily, or if the node is a single-node VCS configuration where LLT is not configured, then the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

Workaround: Set the value of the `CPS_NODEID` environment variable to 255. The `cpsadm` command reads the `CPS_NODEID` variable and proceeds if the command is unable to get LLT node ID from LLT.

Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)

Even if the cluster nodes have no registration on the CP server and if you provide coordination point server (CP server) information in the `vxfenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

Workaround: Manually add the cluster information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster information to the CP server during configuration.

CP server repetitively logs unavailable IP addresses (2530864)

If coordination point server (CP server) fails to listen on any of the IP addresses that are mentioned in the `vxcps.conf` file or that are dynamically added using the command line, then CP server logs an error at regular intervals to indicate the failure. The logging continues until the IP address is bound to successfully.

```
CPS ERROR V-97-51-103 Could not create socket for host
10.209.79.60 on port 14250
CPS ERROR V-97-1400-791 Coordination point server could not
open listening port = [10.209.79.60]:14250
Check if port is already in use.
```

Workaround: Remove the offending IP address from the listening IP addresses list using the `rm_port` action of the `cpsadm` command.

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

Hostname and username are case sensitive in CP server (2846392)

The hostname and username on the CP server are case sensitive. The hostname and username used by fencing to communicate with CP server must be in same case as present in CP server database, else fencing fails to start.

Workaround: Make sure that the same case is used in the hostname and username on the CP server.

Cannot run the vxfentsthdw utility directly from the install media if VRTSvxfen package is not installed on the system (2858190)

If VRTSvxfen package is not installed on the system, then certain script files that are needed for the vxfentsthdw utility to function are not available. So, without the

VRTSvxfen package installed on the system you cannot run the utility from the install media.

Workaround: Install VRTSvxfen package, then run the utility from either the install media or from the `/opt/VRTSvcs/vxfen/bin/` location.

Server-based fencing may fail to start after reinstalling the stack (2802682)

Server-based fencing may fail to start if you use the existing configuration files after reinstalling the stack.

Workaround:

After reinstalling the stack, add the client cluster information on the coordination point server because the client cluster information is removed when the stack is uninstalled. For more details, see the Setting up server-based I/O Fencing manually section in the Veritas Cluster Server Installation Guide. Alternatively, you can manually modify the `/etc/vxfenmode` file and the `main.cf` file to start fencing in disable mode and then configure fencing.

Fencing may show the RFSM state as replaying for some nodes in the cluster (2555191)

Fencing based on coordination point clients in Campus cluster environment may show the RFSM state as replaying for some nodes in the cluster.

Workaround:

Restart fencing on the node that shows RFSM state as replaying.

After you run the vxfenswap utility the CoordPoint agent may fault (2846389)

After you run the `vxfenswap` utility, if the value of the `FaultTolerance` attribute of the CoordPoint agent is more than the majority (more than 50%) of the coordination points then the Coordination Point agent faults.

Workaround: Manually set the value of the `FaultTolerance` attribute of CoordPoint agent to be less than the majority (more than 50%) of the coordination points.

The cpsadm command fails after upgrading CP server to 6.0 or above in secure mode (2846727)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTS`at` RPM is not removed from the system, the `cpsadm` command loads the old security libraries present on the system.

As the installer runs the `cpsadm` command on the CP server to add or upgrade the VCS cluster (application cluster), the installer also fails.

Workaround: Perform the following procedure on all of the nodes of the CP server.

To resolve this issue

1 Rename `cpsadm` to `cpsadmbin`:

```
# mv /opt/VRTScps/bin/cpsadm /opt/VRTScps/bin/cpsadmbin
```

2 Create a file `/opt/VRTScps/bin/cpsadm` with the following content:

```
#!/bin/sh
EAT_USE_LIBPATH="/opt/VRTScps/lib"
export EAT_USE_LIBPATH
/opt/VRTScps/bin/cpsadmbin "$@"
```

3 Change the permissions of the new file to 775:

```
# chmod 755 /opt/VRTScps/bin/cpsadm
```

Server-based fencing comes up incorrectly if default port is not mentioned (2403453)

When you configure fencing in customized mode and do not provide default port, fencing comes up. However, the `vxfenconfig -l` command output does not list the port numbers.

Workaround: Retain the "port=<port_value>" setting in the `/etc/vxfenmode` file, when using customized fencing with at least one CP server. The default port value is 14250.

Secure CP server does not connect from localhost using 127.0.0.1 as the IP address (2554981)

The `cpsadm` command does not connect to the secure CP server on the localhost using 127.0.0.1 as the IP address

Workaround: Connect the secure CP server using any of the virtual IPs that is configured with the CP server and is plumbed on the local node.

Unable to customize the 30-second duration (2551621)

When the vxcpsserv process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: There is no workaround for this issue.

Coordination point server-based fencing may fail if it is configured on 5.1SP1RP1 using 6.0.1 coordination point servers (2824472)

The 5.1SP1 installer (CPI) cannot set up trust between a 5.1SP1 client and a 6.0 or later server, because there are no separate directories for truststores in the 5.1SP1. When trust cannot be setup, the 5.1SP1 installer cannot configure 5.1SP1 clients to work with 6.0 or later CPS in secure mode.

Workaround:

Set up trust manually between the CPS and clients using the cpsat or the vcsat command. After that, CPS and client will be able to communicate properly in the secure mode.

Veritas Cluster Server agents for Veritas Volume Replicator known issues in 6.0.4

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in 6.0.4 release.

fdsetup cannot correctly parse disk names containing characters such as "-" (1949294)

The fdsetup cannot correctly parse disk names containing characters such as "-".

Stale entries observed in the sample main.cf file for RVGLogowner and RVGPrimary agent [2872047]

Stale entries are found in sample main.cf file for RVGLogowner agent and RVGPrimary agent.

The stale entries are present in the main.cf.seattle and main.cf.london files on the RVGLogowner agent which includes CFSQlogckd resource. However, CFSQlogckd is not supported since VCS 5.0.

On RVGPrimary agent, the stale entries are present in file `main.cf.seattle` and `main.cf.london` and the stale entry includes the DetailMonitor attribute.

Workaround

1 For main.cf.seattle for RVGLogowner agent in the cvm group:

- Remove the following lines.

```
CFSQlogckd qlogckd (
    Critical = 0
)
```

```
cvm_clus requires cvm_vxconfigd
qlogckd requires cvm_clus
vxfscd requires qlogckd
```

```
// resource dependency tree
//
//   group cvm
//   {
//     CFSfsckd vxfscd
//     {
//       CFSQlogckd qlogckd
//       {
//         CVMCluster cvm_clus
//         {
//           CVMVxconfigd cvm_vxconfigd
//         }
//       }
//     }
//   }
// }
```

- Replace the above lines with the following:

```
cvm_clus requires cvm_vxconfigd
vxfscd requires cvm_clus
```

```
// resource dependency tree
//
//   group cvm
//   {
//     CFSfsckd vxfscd
//     {
//       CVMCluster cvm_clus
//       {
```



```
//          CVMVxconfigd cvm_vxconfigd
//          }
//      }
//  }
```

2 For main.cf.london for RVGLogowner in the cvm group:

- Remove the following lines

```
CFSQlogckd qlogckd (
    Critical = 0
)
```

```
cvm_clus requires cvm_vxconfigd
qlogckd requires cvm_clus
vxfsckd requires qlogckd
```

```
// resource dependency tree
//
//      group cvm
//      {
//          CFSfsckd vxfsckd
//          {
//              CFSQlogckd qlogckd
//              {
//                  CVMcluster cvm_clus
//                  {
//                      CVMVxconfigd cvm_vxconfigd
//                  }
//              }
//          }
//      }
```

- Replace the above lines with the following:

```
cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus
```

```
// resource dependency tree
//
//      group cvm
//      {
```

```
//      CFSfsckd vxfscckd
//      {
//          CVMCluster cvm_clus
//          {
//              CVMVxconfigd cvm_vxconfigd
//          }
//      }
```

- 3 For main.cf.seattle for RVGPrimary agent in the cvm group:
 - In the group ORAGrp and for the Oracle resource database, remove the line: `DetailMonitor = 1`
- 4 For main.cf.london for RVGPrimary agent in the cvm group:
 - In the group ORAGrp and for the Oracle resource database, remove the line: `DetailMonitor = 1`

Issues related to Intelligent Monitoring Framework (IMF)

Registration error while creating a Firedrill setup [2564350]

While creating the Firedrill setup using the `Firedrill setup` utility, VCS encounters the following error:

```
AMF amfregister ERROR V-292-2-167
Cannot register mount offline event
```

During Firedrill operations, VCS may log error messages related to IMF registration failure in the engine log. This happens because in the firedrill service group, there is a second CFSMount resource monitoring the same MountPoint through IMF. Both the resources try to register for online/offline events on the same MountPoint and as a result, registration of one fails.

Workaround: No workaround.

Perl errors seen while using haimfconfig command

Perl errors seen while using `haimfconfig` command:

```
Perl errors seen while using haimfconfig command
```

This error is due to the absolute path specified in main.cf for type-specific configuration files. Currently, `haimfconfig` does not support absolute path for type-specific configuration file in main.cf.

Workaround: Replace the actual path with the actual file name and copy the file from its absolute location to `/etc/VRTSvcs/conf/config` directory.

For example, if `OracleTypes.cf` is included in `main.cf` as:

```
include "/etc/VRTSagents/ha/conf/Oracle/OracleTypes.cf"
```

It should be replaced as follows in `main.cf`:

```
include "OracleTypes.cf"
```

IMF does not provide notification for a registered disk group if it is imported using a different name [2730774]

If a disk group resource is registered with the AMF and the disk group is then imported using a different name, AMF does not recognize the renamed disk group and hence does not provide notification to DiskGroup agent. Therefore, the DiskGroup agent keeps reporting the disk group resource as offline.

Workaround: Make sure that while importing a disk group, the disk group name matches the the one registered with the AMF.

Direct execution of `linkamf` displays syntax error [2858163]

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

```
# /opt/VRTSperl/bin/perl /opt/VRTSamf/imf/linkamf <destination-directory>
```

Error messages displayed during reboot cycles [2847950]

During some reboot cycles, the following message might get logged in the engine log:

```
AMF libvxamf ERROR V-292-2-149 Cannot unregister event: no rid -1 found  
AMF libvxamf ERROR V-292-2-306 Unable to unregister all events (errno:405)
```

This does not have any effect on the functionality of IMF.

Workaround: No workaround.

Error message displayed when ProPCV prevents a process from coming ONLINE to prevent concurrency violation does not have I18N support [2848011]

The following message is seen when ProPCV prevents a process from coming ONLINE to prevent concurrency violation. The message is displayed in English and does not have I18N support.

```
Concurrency Violation detected by VCS AMF.  
Process <process-details> will be prevented from startup.
```

Workaround: No Workaround.

Error message seen during system shutdown [2954309]

During some system shutdowns, you might see the following message in the syslog.

```
Stopping AMF...  
AMF amfconfig ERROR V-292-2-405 AMF_UNCONFIG failed, return value = -1
```

The system continues to proceed with the shutdown.

Workaround: No workaround.

System panics when `getnotification` requests access of groups cleaned by AMF [2848009]

AMF while handling an agent that has faulted due to external or internal activity, cleans up the groups monitored by the agent. Simultaneously, if the agent notification is in progress and the `getnotification` thread requests to access an already deleted group, the system panics.

Workaround: No workaround.

The `libvxamf` library encounters an error condition while doing a process table scan [2848007]

Sometimes, while doing a process table scan, the `libvxamf` encounters an error condition. As a result, the process offline registration with AMF fails. In most cases, this registration succeeds when tried again by the agent during the next monitor cycle for this resource. This is not a catastrophic failure as the traditional monitoring continues for this resource.

Workaround: No workaround.

AMF displays StartProgram name multiple times on the console without a VCS error code or logs [2872064]

When VCS AMF prevents a process from starting, it displays a message on the console and in syslog. The message contains the signature of the process that was prevented from starting. In some cases, this signature might not match the signature visible in the PS output. For example, the name of the shell script that was prevented from executing will be printed twice.

Workaround: No workaround.

Terminating the `imfd` daemon orphans the `vxnotify` process [2728787]

If you terminate `imfd` daemon using the `kill -9` command, the `vxnotify` process created by `imfd` does not exit automatically but gets orphaned. However, if you stop `imfd` daemon with the `amfconfig -D` command, the corresponding `vxnotify` process is terminated.

Workaround: The correct way to stop any daemon is to gracefully stop it with the appropriate command (which is `amfconfig -D` command in this case), or to terminate the daemon using Session-ID. Session-ID is the -PID (negative PID) of the daemon.

For example:

```
# kill -9 27824
```

Stopping the daemon gracefully stops all the child processes spawned by the daemon. However, using `kill -9 pid` to terminate a daemon is not a recommended option to stop a daemon, and subsequently you must kill other child processes of the daemon manually.

Core dump observed when `amfconfig` is run with `set` and `reset` commands simultaneously [2871890]

When you run `amfconfig -S -R` on a node, a command core dump is observed, instead of displaying the correct usage of the command. However, this core dump has no effect on the AMF functionality on that node. You need to use the correct command syntax instead.

Workaround: Use the correct commands:

```
# amfconfig -S <options>
# amfconfig -R <options>
```

Issues related to the Cluster Manager (Java Console)

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

Cluster Manager (Java Console) may display an error while loading templates (1433844)

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

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

Workaround: Ignore the error.

Some Cluster Manager features fail to work in a firewall setup [1392406]

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

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

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

Issues related to virtualization

Switching KVMGuest resource fails due to inadequate swap space on the other host [2753936]

Virtual machine fails to start on a host if the host does not have the sufficient swap space available.

Workaround: Please make sure that each host has sufficient swap space available for starting the virtual machine.

Policies introduced in SLES 11SP2 may block graceful shutdown if a VM in SUSE KVM environment [2792889]

In a SUSE KVM environment, virtual machine running SLES11 SP2 inside may block the virtual machine graceful shutdown request due to some policies introduced in SLES 11SP2. SUSE recommends turning off the policy with `polkit-gnome-authorization` for a virtual machine.

Workaround: Make sure that all the policies blocking any such request are turned off.

Load on `libvirtd` may terminate it in SUSE KVM environment [2824952]

In a SUSE KVM environment, occasionally `libvirtd` process may get terminated and `/etc/init.d/libvirtd status` command displays:

```
#/etc/init.d/libvirtd status
Checking status of libvirtd                                dead
```

This may be due to heavy load on `libvirtd` process.

Workaround: Restart the `libvirtd` process and run:

```
# service libvirtd stop
# service libvirtd start
```

VCS may detect the migration event during the regular monitor cycle due to the timing issue [2827227]

In a virtualization environment, VCS detects the virtual machine migration initiated outside VCS and changes the state accordingly. However, occasionally, VCS may miss the migration event due to timing issue and detect the migration during the regular monitor cycle. For example, if you set `OfflineMonitorInterval` as `300sec`, it takes up to 5 minutes for VCS to report `ONLINE` on the node where the virtual machine got migrated.

Workaround: No workaround available.

Resource faults when it fails to `ONLINE` VM because of insufficient swap percentage [2827214]

In virtualization environment, if VCS fails to start the virtual machine due to unavailability of required virtualization resources such as CPU, memory, or disks, the resource goes into `FAULTED` state.

Workaround: Make sure that the required virtualization resources are always available in a virtualization environment.

Offline or switch of `KVMGuest` resource fails if the VM it is monitoring is undefined [2796817]

In a SUSE KVM environment, if a running virtual machine is undefined using `virsh undefine` command, an attempt to offline or switch the KVM guest resource

monitoring that VM fails because the agent is not able to get the information from the KVM hypervisor.

Workaround: To undefine the VM on a particular node, first switch the service group containing the KVMGuest resource to another node and then undefine the VM on the first node.

Migration of guest VM on native LVM volume may cause libvirtd process to terminate abruptly (2582716)

When the guest VM image is on native LVM volume, then the migration of that guest initiated by the administrator may cause `libvirtd` process to terminate abruptly.

Workaround: Start the `libvirtd` process manually.

Wizard, Dashboard, and Symantec High Availability tab specific issues

The vCenter Server tasks shown in the vSphere Client may fail to display the correct status if the Symantec High Availability Guest Components installation fails

This issue occurs in case of VMware vCenter Server version 5.1.

If you choose to install the Symantec High Availability Guest Components using the vSphere Client menu and if the installation fails, then the vCenter Server task shown in the vSphere Client gets held-up at 10%. It however, does not display that the installation has failed. (2824900)

Workaround:

Refer to the installer logs at the following locations on the Symantec High Availability Console server. Rectify the cause and re-run the wizard.

```
% ALLUSERSPROFILE %\Symantec\ApplicationHA\Log\ApplicationHA.log
```

```
% ALLUSERSPROFILE %\Symantec\ApplicationHA\Log\VII\
```

You can also refer to logs on the system at the following location:

```
(1) /opt/INSTALLER/vii_logs
```

Alternatively, choose to install the guest components using the product installer or the CLI. For details refer to the product-specific installation and upgrade guide.

Software limitations

This section covers the software limitations of this release.

See the corresponding Release Notes for a complete list of software limitations related to that component or product.

See [“Documentation”](#) on page 74.

Limitations related to VMware virtual environment

Symantec High Availability Configuration Wizard does not support nested mount points [2874775]

Symantec High Availability Configuration Wizard does not discover nested mount points. As a result, you cannot use the wizard to monitor applications that access a database via a nested mount point.

Workaround: Use VCS commands to configure the mount resources for monitoring nested mount points. Ensure that you set appropriate resource dependencies to configure the nested mounts. For more information, see the *Veritas Cluster Server Administrator's Guide*.

Symantec High Availability tab does not provide an option to remove a node from a VCS cluster [2944997]

The **Symantec High Availability** tab does not support removal of a system from a VCS cluster.

Workaround: You can use VCS commands to remove the system from the VCS cluster. You can also unconfigure the entire VCS cluster from the **Symantec High Availability** tab.

VMware fault tolerance does not support adding or removing non-shared disks between virtual machines

VMware fault tolerance does not support adding or removing of non-shared disks between virtual machines. During a failover, disks that contain application data cannot be moved to alternate failover systems. Applications that are being monitored, thus cannot be brought online on the failover systems.

Symantec High Availability wizard does not support configuration of parallel service groups

Symantec High Availability wizard does not support configuration of parallel service groups. You can, however, configure parallel groups from the command line. Such parallel service groups are visible in the Symantec High Availability tab. You cannot, however, modify the configuration from the GUI.

Symantec High Availability wizard limitations

Wizard fails to configure VCS resources if storage resources have the same name [3024460]

Naming storage resources like disk group and volumes with the same name is not supported as the Symantec High Availability wizard fails to configure the VCS resources correctly.

Workaround: No workaround.

Limitations related to installing and upgrading VCS

Remote and target systems must have the same OS and architecture while using installer from a remote system [589334]

If you use the installer from a remote system, then the remote system must have the same operating system and architecture as that of the target systems where you want to install VCS.

Limitations related to bundled agents

Volumes in a disk group start automatically irrespective of the value of the StartVolumes attribute in VCS

Volumes in a disk group are started automatically when the disk group is imported, irrespective of the value of the StartVolumes attribute in VCS. This behavior is observed if the value of the system-level attribute `autostartvolumes` in Veritas Volume Manager is set to On.

Workaround: If you do not want the volumes in a disk group to start automatically after the import of a disk group, set the `autostartvolumes` attribute to Off at the system level.

Programs using networked services may stop responding if the host is disconnected

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can stop responding 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 stop responding and possibly time out.

For example, if the host is configured to use NIS maps as a client, basic commands such as `ps -ef` can hang if there is network disconnect.

Symantec recommends creating users locally. To reflect local users, configure:

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

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.

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 may 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 stopped that are not under VCS control.

Mount agent limitations

The Mount agent has the following limitations:

- The Mount agent mounts a block device at only one mount point on a system. After a block device is mounted, the agent cannot mount another device at the same mount point.
- Mount agent does not support:
 - ext4 filesystem configured on VxVM
 - xfs filesystem configured on VxVM

Share agent limitations

To ensure proper monitoring by the Share agent, verify that the `/var/lib/nfs/etab` file is clear upon system reboot. Clients in the Share agent must be specified as fully qualified host names to ensure seamless failover.

Driver requirements for DiskReservation agent

The `VRTSvcsdr` package ships the `scsiutil` utility. DiskReservation agent supports only those drivers supported by the `scsiutil` utility.

Application agent limitations

- ProPCV fails to prevent execution of script-based processes configured under `MonitorProcesses`.

Limitations related to IMF

- IMF registration on Linux for “bind” file system type is not supported.
- In case of SLES11 and RHEL6:
 - IMF should not be enabled for the resources where the `BlockDevice` can get mounted on multiple `MountPoints`.
 - If `FSType` attribute value is `nfs`, then IMF registration for “nfs” file system type is not supported.

Agent directory base name must be type name for an agent using out-of-the-box `imf_init` IMF entry point to get IMF support [2858160]

To get IMF support for an agent which uses the out-of-the-box `imf_init` IMF entry point, the base name of agent directory must be the type name. When `AgentFile` is set to one of the out-of-the-box agents like `Script51Agent`, that agent will not get IMF support.

Workaround:

- 1 Create the following symlink in agent directory (for example in `/opt/VRTSagents/ha/bin/WebSphereMQ6` directory).

```
# cd /opt/VRTSagents/ha/bin/<ResourceType>
# ln -s /opt/VRTSvcs/bin/Script51Agent <ResourceType>Agent
```
- 2 Run the following command to update the `AgentFile` attribute based on value of `VCS_HOME`.

- If VCS_HOME is /opt/VRTSvcs:

```
# hatype -modify <ResourceType> AgentFile  
/opt/VRTSvcs/bin/<ResourceType>/<ResourceType>Agent
```
- If VCS_HOME is /opt/VRTSagents/ha:

```
# hatype -modify <ResourceType> AgentFile  
/opt/VRTSagents/ha/bin/<ResourceType>/<ResourceType>Agent
```

Limitations related to the VCS database agents

DB2 RestartLimit value [1234959]

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.

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]

Sybase agent does not perform qrmutil based checks if Quorum_dev is not set (2724848)

If you do not set the Quorum_dev attribute for Sybase Cluster Edition, the Sybase agent does not perform the qrmutil-based checks. This error in configuration may lead to undesirable results. For example, if qrmutil returns failure pending, the agent does not panic the system. Thus, the Sybase agent does not perform qrmutil-based checks because the Quorum_dev attribute is not set.

Therefore, setting `Quorum_Dev` attribute is mandatory for Sybase cluster edition.

Security-Enhanced Linux is not supported on SLES distributions

VCS does not support Security-Enhanced Linux (SELinux) on SLES10 and SLES11. [1056433]

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.

VxVM site for the disk group remains detached after node reboot in campus clusters with fire drill [1919317]

When you bring the `DiskGroupSnap` resource online, the `DiskGroupSnap` agent detaches the site from the target disk group 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 disk group is imported, which is at the primary site.

If you attempt to shut down the node where the fire drill service group or the disk group 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 disk group. 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 disk group site still remains detached.

Workaround:

You must take the fire drill service group offline using the `hagrp -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 disk group that is imported at the primary site.

If the secondary node has crashed or restarted, you must manually reattach the fire drill site to the target disk group that is imported at the primary site using the following command: `/opt/VRTSvcs/bin/hares -action $targetres joindg -actionargs $fdsitename $is_fenced -sys $targetsys`.

Limitations with DiskGroupSnap agent [1919329]

The `DiskGroupSnap` agent has the following limitations:

- The DiskGroupSnap agent does not support layered volumes.
- 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:
 - 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.

System reboot after panic

If the VCS kernel module issues a system panic, a system reboot is required [293447]. The supported Linux kernels do not automatically halt (CPU) processing. Set the Linux “panic” kernel parameter to a value other than zero to forcibly reboot the system. Append the following two lines at the end of the `/etc/sysctl.conf` file:

```
# force a reboot after 60 seconds
kernel.panic = 60
```

Cluster Manager (Java console) limitations

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

Cluster Manager (Java Console) version 5.1 and lower cannot manage VCS 6.0 secure clusters

Cluster Manager (Java Console) from versions lower than VCS 5.1 cannot be used to manage VCS 6.0 secure clusters. Symantec recommends using the latest version of Cluster Manager.

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

Cluster Manager does not work if the `hosts` file contains IPv6 entries

VCS Cluster Manager fails 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”.

Limited support from Cluster Manager (Java console)

Features introduced in VCS 6.0 may not work as expected with Java console. However, CLI option of the simulator supports all the VCS 6.0 features. You are recommended to use Veritas Operations Manager (VOM) since all new features are already supported in VOM. However, Java console may continue to work as expected with features of releases prior to VCS 6.0.

Port change required to connect to secure cluster [2615068]

In order to connect to secure cluster, the default port must be changed from 2821 to 14149. You must choose **Advanced settings** in the **Login** dialog box and change **IP: 2821** to **IP: 14149** for secure cluster login.

Using the KDE desktop

Some menus and dialog boxes on Cluster Manager (Java Console) may appear misaligned or incorrectly sized on a KDE desktop. To ensure the proper appearance and functionality of the console on a KDE desktop, use the Sawfish window manager. You must explicitly select the Sawfish window manager even if it is supposed to appear as the default window manager on a KDE desktop.

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

The preferred fencing feature gives preference to more weighted or larger subclusters by delaying the smaller subcluster. This smaller subcluster delay is effective only if the initial RACER node in the larger subcluster is able to complete the race. If due to some reason the initial RACER node is not able to complete the race and the VxFEN driver activates the racer re-election algorithm, then the smaller subcluster delay is offset by the time taken for the racer re-election and the less weighted or smaller subcluster could win the race. This limitation though not desirable can be tolerated.

Uninstalling VRTSvxvm causes issues when VxFEN is configured in SCSI3 mode with dmp disk policy (2522069)

When VxFEN is configured in SCSI3 mode with dmp disk policy, the DMP nodes for the coordinator disks can be accessed during system shutdown or fencing arbitration. After uninstalling VRTSvxvm RPM, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm RPM is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

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.

In a cluster using SCSI-3 based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on both the data disks and coordinator disks. In a cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. 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 coordination 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 coordination 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.

Limitations related to global clusters

- Cluster address for global cluster requires resolved virtual IP.
The virtual IP address must have a DNS entry if virtual IP is used for heartbeat agents.
- Total number of clusters in a global cluster configuration can not exceed four.
- Cluster may not be declared as faulted when Symm heartbeat agent is configured even when all hosts are down.
The Symm agent is used to monitor the link between two Symmetrix arrays. When all the hosts are down in a cluster but the Symm agent is able to see the

replication link between the local and remote storage, it would report the heartbeat as ALIVE. Due to this, DR site does not declare the primary site as faulted.

Documentation

Product guides are available in the PDF format on the software media in the `/docs/product_name` directory. Additional documentation is available online.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The latest product documentation is available on the Symantec website.

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

Documentation set

Table 1-19 lists the documents for Veritas Cluster Server.

Table 1-19 Veritas Cluster Server documentation

Title	File name
<i>Veritas Cluster Server Installation Guide</i>	vcs_install_604_lin.pdf
<i>Veritas Cluster Server Release Notes</i>	vcs_notes_604_lin.pdf
<i>Veritas Cluster Server Administrator's Guide</i>	vcs_admin_604_lin.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents_604_lin.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i> (This document is available online, only.)	vcs_agent_dev_601_unix.pdf
<i>Veritas Cluster Server Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_agent_604_lin.pdf
<i>Veritas Cluster Server Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_agent_604_lin.pdf
<i>Veritas Cluster Server Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_agent_604_lin.pdf

Table 1-20 lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-20 Veritas Storage Foundation and High Availability Solutions products documentation

Document title	File name
<i>Veritas Storage Foundation and High Availability Solutions Solutions Guide</i>	sphas_solutions_604_lin.pdf
<i>Veritas Storage Foundation and High Availability Solutions Virtualization Guide</i>	sphas_virtualization_604_lin.pdf

If you use Veritas Operations Manager (VOM) to manage Veritas Storage Foundation and High Availability products, refer to the VOM product documentation at:

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

Manual pages

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

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

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

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

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

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

See the `man(1)` manual page.

Manual pages are divided into sections 1, 1M, 3N, 4, and 4M. Edit the `man(1)` configuration file `/etc/man.config` to view these pages.

To edit the man(1) configuration file

- 1 If you use the man command to access manual pages, set `LC_ALL` to “C” in your shell to ensure that the pages are displayed correctly.

```
export LC_ALL=C
```

See incident 82099 on the Red Hat Linux support website for more information.

- 2 Add the following line to `/etc/man.config`:

```
MANPATH /opt/VRTS/man
```

where other man paths are specified in the configuration file.

- 3 Add new section numbers. Change the line:

```
MANSECT          1:8:2:3:4:5:6:7:9:tcl:n:l:p:o
```

to

```
MANSECT          1:8:2:3:4:5:6:7:9:tcl:n:l:p:o:3n:1m
```

The latest manual pages are available online in HTML format on the Symantec website at:

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