Symantec™ Cluster Server 6.2.1 Release Notes - Linux

Platform Release

OL7 UEK R3

May 2015

Symantec

Symantec™ 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: 6.2.1

Document version: 6.2.1 Rev 1

Legal Notice

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_apj@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

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

Symantec Cluster Server Release Notes

This document includes the following topics:

About this document

This document provides important information about Symantec Cluster Server (VCS) version 6.2.1 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: 6.2.1 Rev 1" of the *Symantec 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 website 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:

- *Symantec Cluster Server Release Notes* (6.2.1)

About Symantec Cluster Server

Symantec 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 *Symantec Cluster Server Bundled Agents Reference Guide*.

The Symantec 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 Symantec 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 *Symantec 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 website 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	■ 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 and generate an Installation and Upgrade custom report. ■ List patches by product or platform, and in the order they need to be installed. Display and download the most recent patches or historical patches. ■ Display Array Support Library (ASL) details by vendor, platform, or Storage Foundation and High Availability (SFHA) version. ASLs make it easier to manage arrays that are connected to SFHA-based servers. ■ List VCS and ApplicationHA agents, documentation, and downloads based on the agent type, application, and platform.

Identify risks and get server-specific recommendations	■ Analyze your servers for potential environmental risks. Generate a Risk Assessment custom report with specific recommendations about system availability, storage use, performance, and best practices. ■ Display descriptions and solutions for thousands of Symantec error codes.
Improve efficiency	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), documentation, product releases, Hardware Compatibility Lists (HCLs), and VCS/ApplicationHA agents. ■ Quickly gather installed Symantec product and license key information from across your production environment. Generate a License/Deployment custom report that includes product names, versions, and platforms, server tiers, Symantec Performance Value Units (SPVUs), and End of Service Life dates. ■ List and download Symantec product documentation including product guides, manual pages, compatibility lists, and support articles. ■ Access links to important resources on a single page, including Symantec product support, SymConnect forums, customer care, Symantec training and education, Symantec FileConnect, the licensing portal, and my.symantec.com. The page also includes links to key vendor support sites. ■ Use a subset of SORT features from your iOS device. Download the application at: https://sort.symantec.com/mobile

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

List of RPMs

The section lists the RPMs for Linux.

Table 1-1 lists the RPMs that are updated in this release for Oracle Linux 7 Unbreakable Enterprise Kernel 3.

Table 1-1 RPMs for Oracle Linux 7 Unbreakable Enterprise Kernel 3

Name	Version	Arch	Size in Bytes
VRTSamf	6.2.1.000	x86_64	2366140
VRTScps	6.2.1.000	i686	11634572
VRTSgab	6.2.1.000	x86_64	2467204
VRTSllt	6.2.1.000	x86_64	6113844
VRTSperl	5.16.1.27	x86_64	15836048
VRTSsfcpi62	6.2.1.000	noarch	1445355
VRTSsfmh	6.1.0.400	x86_64	43668551
VRTSspt	6.2.1.000	noarch	25899213
VRTSvbs	6.2.1.000	i686	19925529
VRTSvcs	6.2.1.000	i686	64038472
VRTSvcsag	6.2.1.000	i686	12374192
VRTSvcsea	6.2.1.000	i686	256344
VRTSvcsvmw	6.2.1.000	i686	8491964
VRTSvcswiz	6.2.1.000	i686	7071704
VRTSveki	6.2.1.000	x86_64	36484
VRTSvlic	3.02.62.004	x86_64	407143
VRTSvxfen	6.2.1.000	x86_64	2680044

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/TECH225259
- 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/TECH211575

- The software compatibility list summarizes each Storage Foundation and High Availability (SFHA) Solutions product stack and the product features, operating system versions, and third-party products it supports. For the latest information on supported software, visit the following URL:
  http://www.symantec.com/docs/TECH225258

Note: Before you install or upgrade SFHA Solutions products, review the current compatibility lists to confirm the compatibility of your hardware and software.

Changes introduced in 6.2.1

This section lists the changes in Symantec Cluster Server 6.2.1.

Platform support introduced in this release

The following platform support is introduced in this release.

Support for Oracle Linux Unbreakable Enterprise Kernel

In this release, Symantec Cluster Server (VCS) is enhanced to support Oracle Linux 7 (OL 7) with Unbreakable Enterprise Kernel Release 3 (UEK R3).

Changes to VCS bundled agents

This section describes changes to the bundled agents for VCS.

See the *Administrator's Guide* and *Bundled Agents Reference Guide* for more information.

Btrfs support in Mount Agent

The Linux Mount agent is enhanced to support the Btrfs (B-tree file system) file system. With this enhancement, you now can use the VCS Mount Agent on the Linux platform to mount the btrfs file system.

Changes to VCS agent framework

The following changes are introduced to the VCS agent framework.

Changes to LLT, GAB, and I/O fencing

This section covers new features or enhancements made to LLT, GAB, and I/O fencing.

Changes to wizard support

You can use the Symantec High Availability Configuration wizard to configure application monitoring for generic applications running on Linux on a physical host.

Changes to database agents

Support for SAP ASE 16 in single instance mode

In this release, Symantec Cluster Server (VCS) is enhanced to support SAP ASE (previously Sybase) 16 in single instance mode.

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.

Note: The system from where you install VCS must run the same Linux distribution as the other cluster nodes.

See "Hardware compatibility list" on page 13.

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/TECH211575

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

Supported Linux operating systems

For current updates, visit the Symantec Operations Readiness Tools Installation and Upgrade page: https://sort.symantec.com/land/install_and_upgrade.

Table 1-2 shows the supported operating systems for this release.

Table 1-2 Supported operating systems

Operating systems	Levels	Kernel version
Oracle Linux 7 UEK R3	3.8.13-35.3.1.el7uek.x86_64	3.8.13-35.3.1.el7uek.x86_64

Note: Only 64-bit operating systems are supported on the AMD Opteron or the Intel Xeon EM64T (x86_64) Processor line.

If your system is running an older version of Oracle Linux, upgrade it before attempting to install the Symantec software. Consult the Oracle documentation for more information on upgrading or reinstalling your operating system.

Symantec supports only Oracle distributed kernel binaries.

Required Linux RPMs for VCS

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

Table 1-3 lists the RPMs that VCS products require for a given Linux operating system.

Table 1-3 Required RPMs

Operating system	Required RPMs
OL 7	glibc-2.17-55.el7.x86_64 glibc-2.17-55.el7.i686 ksh-20120801-19.el7.x86_64 libgcc-4.8.2-16.el7.i686 libstdc++-4.8.2-16.el7.i686 perl-Exporter-5.68-3.el7.noarch perl-5.16.3-283.el7.x86_64

Supported enterprise agents

Refer to the following links for the supported enterprise agent support matrix for each agent:

Oracle	Support matrix for Oracle
DB2	Support matrix for DB2
Sybase	Support matrix for Sybase

See the Symantec Cluster Server agent guides for Oracle, DB2 and Sybase for more details.

For a list of the VCS application agents and the software that the agents support, see the Symantec 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:

- Raw disk I/O fencing policy is no longer supported.

Cluster Manager Java GUI support consideration

The Cluster Manager Java GUI is End of Life but continues to be supported by Symantec. The Java GUI remains available for download and use within clusters with support for all VCS features available in pre-6.0 releases. Customers can manage service groups, generate new configurations, and perform other traditional cluster management operations. Symantec supports the Java GUI only on the Linux and Windows platforms.

Additional feature capabilities and platform support added in VCS 6.0 and later releases are available exclusively through Veritas Operations Manager (VOM). Symantec recommends the use of VOM to manage clusters and for all advanced capabilities.

Deprecated attributes

The following table lists the attributes deprecated in this release.

Table 1-4 Attributes deprecated in this release

Attribute name	Agent type
SecondLevelMonitor	Apache **Note:** The SecondLevelMonitor attribute is deprecated in VCS 6.2. Instead, LevelTwoMonitorFreq attribute at the Apache resource type level may be used
ResLogLevel	Apache **Note:** Use type level attribute LogDbg to enable debug logs. Set LogDbg attribute to DBG_5 to enable debug logs for Apache agent. By default setting the LogDbg attribute to DBG_5, enables the debug logs for all Apache resources in the cluster. If specific Apache resource needs to be enabled for debug logs override LogDbg attribute.
DetailMonitor	Oracle, Sybase **Note:** If you manually upgrade VCS to 6.2 with detail monitoring enabled in the previous version, set the value of LevelTwoMonitorFreq attribute to that of DetailMonitor.
AgentDebug	DB2udb

Fixed issues

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

Installation and upgrades fixed issues

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

Installation and upgrade fixed issues in 6.2.1

Table 1-5 covers the incidents that are fixed related to installation and upgrade in 6.2.1.

Table 1-5 Installation and upgrade 6.2.1 fixed issues

Incident	Description
3656701	Oracle RAC supports IPv4 addresses only in Oracle High Availability IP (HAIP) configurations.
3719159	In Oracle RAC version 12.1.0.2, warning messages are reported during the post configuration checks.
3763571	The SFHA product installer reports incorrect minimal version for the required Oracle Linux 7 RPMs.
3739179	The module OpenSSL 1.0.1i in the VRTSperl package has security issues.

Installation and upgrades fixed issues in 6.2

Table 1-6 Installation and upgrades fixed issues

Incident	Description
3536489	If you use install VCS from the VMware vSphere Client menu, the installer does not perform a pre-install check for ksh RPMs.

Symantec Cluster Server fixed issues

This section describes Symantec Cluster Server fixed issues.

Symantec Cluster Server fixed issues in 6.2.1

Table 1-7 covers the fixed issues of Symantec Cluster Server in 6.2.1.

Table 1-7 Symantc Cluster Server 6.2.1 fixed issues

Incident	Description
3520211	The HostMonitor agent reports incorrect memory usage.
3652819	VxFS module fails to unload because the AMF module fails to decrement the reference count.
3662501	The notifier process fails to clearly distinguish whether the CPU, Memory, or Swap has exceeded the warning or critical threshold level.
3662508	The CmdServer log incorrectly reports warning messages relevant to licensing even when keyless licensing is enabled and the host is configured under Veritas Operations Manager (VOM).

Table 1-7 Symantc Cluster Server 6.2.1 fixed issues *(continued)*

Incident	Description
3666049	VCS does not support SAP ASE 16 Sybase agent.
3668853	The halog(1M) command becomes unresponsive when VCS is in the REMOTE_BUILD state.
3699146	The Application agent reports an application resource as offline even when the resource is online.
3754061	Added support for RHEL7.1

Symantec Cluster Server fixed issues in 6.2

This section describes Symantec Cluster Server fixed issues in 6.2.

VCS engine fixed issues

Table 1-8 lists the fixed issues for VCS engine.

Table 1-8 VCS engine fixed issues

Incident	Description
3381042	The checkboot utility core dump or time difference between a system and Network Time Protocol (NTP) time leads to unexpected deletion of the temporary files. The deletion causes the VCS agents to report an incorrect state.
2834247	While initiating a global failover of multiple dependent service groups in a cluster, VCS fails to follow the group dependency order while initiating online of the service groups.
3448510	The `hastatus` command fails and dumps core when it is run from a local zone.
3468891	Inconsistencies in `/etc/VRTSvcs/conf/attributes/cluster_attrs.xml` file and hide resource-level ContainerInfo attribute from `hares -display` command.
3211834	CurrentLimits attribute value is not updated correctly when a service group faults.
3385820	Sometimes the high availability daemon (HAD) crashes if it runs for a long duration.

Table 1-8 VCS engine fixed issues *(continued)*

Incident	Description
3436617	When invoking triggers if some arguments are left unassigned, the `hatrigger` command fails due to compilation errors.
3471819	The service group fails to go online if the CurrentCount attribute value is incorrect.
3464981	The file size of Engine_A.log file could not be increased beyond 32 MB.
3580940	VCS configuration becomes unclean when incorrect filename is provided with the `ha` command for SourceFile attribute.
3603275	HAD and other nodes abort while shutting down two nodes simultaneously.

Bundled agents fixed issues

Table 1-9 lists the fixed issues for bundled agents.

Table 1-9 Bundled agents fixed issues

Incident	Description
2490296	Application agent cannot handle a case with user as root, envfile set and shell as `csh`.
2618482	Some agents may fail to come online after full upgrade to VCS 6.0 if they were online before the upgrade.
3354227	The clean entry point operation of the IPMultiNIC agent fails in the absence of a state file that contains current active device information.
3535942	The IP agent shows high CPU utilization and fails to stop even after VCS is stopped when IPv6 address is configured in the resource.
3408712	DiskGroup agent does not fail over the service group if storage connectivity is lost and I/O fencing is configured.
3573876	IP resource fails to go offline when network cable in unplugged.
3505202	VCS does not support non-default value for VCS_LOG environment variable.

Fixed issues related to AMF

Table 1-10 AMF fixed issues

Incident	Description
2848007	The libvxamf library encounters an error condition while doing a process table scan.
3333913	AMF may panic the system if it receives a request to unregister an already unregistered resource.
3407338	If one of the events of a group is in triggered state, then the group fails to unregister because of the triggered event.
3338946	Sometimes when a process offline registration is requested and the system is under heavy load, AMF library fails to verify whether the resource is actually offline. As a result, registration fails.

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 6.2.1

Table 1-11 lists the fixed issues for LLT, GAB, and I/O fencing in 6.2.1.

Table 1-11 LLT, GAB, and I/O fencing 6.2.1 fixed issues

Incident	Description
3567354	The Linux kernel panics when it receives a shared non-linear skb for linearization from the Low Latency Transport (LLT).
3663740	In a rare scenario, divide by zero error is seen when lltshow -p command is run.
3667755	Strict permission check for CPS configuration files.
3691202	Sometimes fencing in the customized mode fails to start when the number of files present in the current working directory is large.
3722178	The rpm --verify command on VXFEN changes the runlevel settings for the VXFEN service.
3728106	On Linux, the value corresponding to 15 minute CPU load average as shown in /proc/loadavg file wrongly increases to about 4.

LLT, GAB, and I/O fencing fixed issues in 6.2

Table 1-12 LLT, GAB, and I/O fencing fixed issues

Incident	Description
3031216	The dash (-) in a disk group name causes vxfentsthdw(1M) and Vxfenswap(1M) utilities to fail.
3335137	Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups.
3410309	LLT driver fails to load and logs the following message in the syslog when a mismatch is observed in the RDMA-specific symbols. llt: disagrees about version of symbol rdma_connect llt: Unknown symbol rdma_connect llt: disagrees about version of symbol rdma_destroy_id llt: Unknown symbol rdma_destroy_id
3471571	Cluster nodes may panic if you stop the HAD process by force on a node and reboot that node.
3473104	When virtual NICs are configured under LLT without specifying the MTU size 1500 in lltttab, cluster does not function properly. For example, VCS engine commands may hang and print below message in the engine logs: `VCS CRITICAL V-16-1-51135 GlobalCounter not updated`
3532859	The Coordpoint agent monitor fails if the cluster has a large number of coordination points.

Known issues

This section covers the known issues in this release.

Issues related to installation

This section describes the known issues during installation.

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

```
# hagrp -list Frozen=1
```

2 Unfreeze all the frozen service groups:

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

Issue with soft links getting deleted in a manual upgrade

While performing a manual upgrade (from 5.1 to 6.2.1) 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-3.02.61.010-0.x86_64.rpm`

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 1-13 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 `ifconfig` command is used. If RouteOptions is set, attribute value is used to add/delete routes using command `route`. As the Options attribute is configured, IPv4RouteOptions values are ignored.	No need to configure IPv4RouteOptions .
Not configured	May or may not be configured	Must be configured	In this case the `ip` command is used. IPv4RouteOptions must be configured and are used to add/delete routes using the `ip route` command. As Options attribute is not configured, RouteOptions value is ignored.	Configure IPv4RouteOptions and set the IP of default gateway. The value of this attribute typically resembles: IPv4RouteOptions = “default via *gateway_ip*” For example: IPv4RouteOptions = “default via 192.168.1.1”

XFS file system is not supported for RDE

The Root Disk Encapsulation (RDE) feature is not supported if the root partition is mounted with XFS file system.

Workaround: There is no workaround available.

Erroneous resstatechange trigger warning [2277819]

You may encounter the following warning when you restart resources:

```
CPI WARNING V-9-40-4317 The installer has detected that resstatechange
trigger is configured by setting TriggerResStateChange attributes.
```

Workaround: In future releases, the resstatechange trigger will not be invoked when a resource is restarted. Instead, the resrestart trigger will be invoked if you set the TriggerResRestart attribute. The resrestart trigger is available in the current release. Refer to the VCS documentation for details.

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 RPMs.

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/117pfhe4zz/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 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.

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.

Rolling upgrade of VCS from pre-6.0 versions fails with CP server in secure mode [3262900]

If the CP server is configured in secure mode, rolling upgrade of VCS from versions lower than 6.0 to 6.1 is not supported. Since the `vxcpserv` process is not compatible with shared authentication, CP server service group fails to come online after performing phase 1 of the rolling upgrade.

Workaround: Use full upgrade or phased upgrade instead of rolling upgrade.

The tuneable values cannot persist after upgrade from the releases prior to 6.0 [3736830]

If you upgrade releases that are prior to 6.0 to 6.0 or later release, the tunable values which were set by vxtune will be replaced by default values.

Workaround: There is no workaround.

The Install<product> script displays 6.2 EULA verison instead of 6.2.1 [3768997]

Oracle Linux 7 UEK3, the `install<product>` script displays the following EULA message:

```
Do you agree with the terms of the End User License Agreement
as specified in the cluster_server/EULA/en/EULA_SFHA_Ux_6.2.pdf file
present on media?
```

The EULA file version 6.2 is not correct. It should be:
cluster_server/EULA/en/EULA_SFHA_Ux_6.2.1.pdf

Workaround:

This is no workaround for this issue.

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 *Symantec 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.

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.

VVR configuration may go in a primary-primary configuration when the primary node crashes and restarts [3314749]

The AutoResync attribute of the RVGPrimary and RVGSharedPri agent control whether the agent must attempt to automatically perform a fast-failback resynchronization of the original primary after a takeover and after the original

primary returns. The default value of this attribute is 0, which instructs the agent not to perform a fast-failback resynchronization of the original primary after a takeover and after the original primary returns. The takeover is performed automatically since the default value of the AutoTakeover attribute of the RVGPrimary and RVGShared agents is 1. Thus, the default settings of AutoTakeover and AutoResync set to 1 and 0 respectively cause the first failover to succeed when the original primary goes down, and on return of the original primary, the Replicated Data Set (RDS) ends up with a primary-primary configuration error.

Workaround: Set the default value of the AutoResync attribute of the RVGPrimary agent to 1 (one) when you want the agent to attempt to automatically perform a fast-failback resynchronization of the original primary after a takeover and after the original primary returns. This prevents the primary-primary configuration error. Do not set AutoResync to 1 (one) if you intend to use the Primary-Elect feature.

Moreover, if you want to prevent VCS from performing an automatic takeover and fast-failback resynchronization, set AutoTakeover and AutoResync attributes to 0 for all the RVGPrimary and RVGSharedPri resources in your VCS configuration. For more information, refer to the RVGPrimary and RVGSharedPri agent sections of the *Replication Agents chapter*in the *Symantec Cluster Server Bundled Agents Reference Guide*.

CP server does not allow adding and removing HTTPS virtual IP or ports when it is running [3322154]

CP server does not support adding and removing HTTPS virtual IPs or ports while the CP server is running. However, You can add or remove the IPM virtual IPs or ports.

Workaround: No workaround. If you want to add a new virtual IP for HTTPS, you must follow the entire manual procedure for generating HTTPS certificate for the CP server (server.crt), as documented in the *Symantec Cluster Server Installation Guide*.

CP server does not support IPv6 communication with HTTPS protocol [3209475]

CP server does not support IPv6 communication when using the HTTPS protocol. This implies that in VCS 6.1, CP servers listening on HTTPS can only use IPv4. As a result, VCS 6.1 fencing clients can also use only IPv4.

Workaround: No workaround.

VCS fails to stop volume due to a transaction ID mismatch error [3292840]

If VCS imports a disk group *A* on node *sys1*, which implies that the DiskGroup resource is online on *sys1*. If you run `vxdg -C import <dg_name>` outside VCS on node *sys2*, then the disk group gets imported on node *sys2* and `-C` clears the import locks and host tag. However on node *sys1*, disk group *A* continues to appear as imported and enabled, and hence, VCS continues to report the resource state as ONLINE on node *sys1*. Subsequently, when VCS detects the imported disk group on *sys2*, it deports the disk group from *sys2*, and imports it on *sys1* to resolve concurrency violation. At this point, the disk group deported from node *sys2* is shown as imported and enabled on node *sys1*. If you stop any volume from within or outside VCS, it fails with the `Transaction ID mismatch` error, but the read and write operations continue to function so the data continues to be accessible. This situation may lead to data corruption if the disk group appears enabled on multiple nodes. This issue is due to the Volume Manager behavior.

Workaround: Do not import a disk group using `-C` option if that diskgroup is under VCS control.

Some VCS components do not work on the systems where a firewall is configured to block TCP traffic [3545338]

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 *Symantec 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.

Issues related to the VCS engine

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.

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 dependancy 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.

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 `hagrp -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 *Symantec 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 for one of the following reasons:

- If you first use a non-root user without a home directory and then create a home directory for the same user.
- If you configure security on a cluster and then un-configure and reconfigure it.

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).

Running `-delete -keys` for any scalar attribute causes core dump [3065357]

Running `-delete -keys` for any scalar attribute is not a valid operation and must not be used. However, any accidental or deliberate use of this command may cause engine to core dump.

Workaround: No workaround.

VCS enters into admin_wait state when Cluster Statistics is enabled with load and capacity defined [3199210]

VCS enters into admin_wait state when started locally if:

1. Statistics attribute value is set to Enabled, which is its default value.
2. Group Load and System Capacity values are defined in units in `main.cf`.

Workaround:

1. Stop VCS on all nodes in the cluster.
2. Perform any one of the following steps:
   - Edit the `main.cf` on one of the nodes in the cluster and set the Statistics attribute to Disabled or MeterHostOnly.
   - Remove the Group Load and System Capacity values from the `main.cf`.

3. Run `hacf -verify` on the node to verify that the configuration is valid.
4. Start VCS on the node and then on the rest of the nodes in the cluster.

Agent reports incorrect state if VCS is not set to start automatically and `utmp` file is empty before VCS is started [3326504]

If you have not configured VCS to start automatically after a reboot and have tmptied the `utmp` file before starting VCS manually with the hastart command, some agents might report an incorrect state.

The `utmp` file (file name may differ on different operating systems) is used to maintain a record of the restarts done for a particular machine. The checkboot utility used by `hastart` command uses the functions provided by the OS which in turn use the `utmp` file to find if a system has been restarted so that the temporary files for various agents can be deleted before agent startup. If OS functions do not return correct value, High Availability Daemon (HAD) starts without deleting the stale agent files. This might result in some agents reporting incorrect state.

Workaround: If a user wishes to delete the `utmp` file this should be done only when VCS is already running or the customer should delete the temporary files in`/var/VRTSvcs/lock/volatile/` manually before starting VCS.

Site preference fencing policy value fails to set on restart of a site-aware cluster [3380586]

If you restart VCS on a site-aware cluster, the PreferredFencingPolicy fails to reset to the value ‘Site’ assigned to it before the restart.

Workaround: Reassign the fencing policy value manually to the cluster.

Log messages are seen on every systemctl transaction on RHEL7 [3609196]

On RHEL7 systems, a log message stating `VCS dependency is not met` is logged in system logs with every `systemctl` transaction. Currently, all the `init` scrips for VCS modules (such as LLT, GAB, I/O fencing, and AMF) bypass systemctl. However, `systemctl` attempts to validate the dependency check before bypassing the service start operation. This generates the log messages in the system log.

Workaround: You can ignore the log messages as they do not affect the `init` script operation.

VCS crashes if feature tracking file is corrupt [3603291]

VCS keeps a track of some specific features used in the VCS cluster. For example, if a Global service group is brought online then the feature is logged in a specific feature tracking file. If the file however is corrupt, then VCS may dump core when attempting to write data to the file.

Workaround: Delete the corrupt feature tracking file (`/var/vx/vftrk/vcs`) and restart VCS.

RemoteGroup agent on versions lower than 6.2 reports service group status as UNKNOWN [3638347]

When the RemoteGroup agent running on a VCS version lower than 6.2 tries to monitor a service group on a 6.2 cluster, it reports the service group status as UNKNOWN.

Workaround: No workaround.

RemoteGroup agent and non-root users may fail to authenticate after a secure upgrade [3649457]

On upgrading a secure cluster to 6.2, the following issues may occur with unable to open a secure connection error:

- The RemoteGroup agent may fail to authenticate with remote cluster.
- Non-root users may fail to authenticate.

Workaround

1. Set `LC_ALL=C` on all nodes before upgrade or perform the following steps after the upgrade on all nodes of the cluster:
   - Stop HAD.
   - Set `LC_ALL=C`.
   - Start HAD using `hastart`.
2. Reset LC_ALL attribute to the previous value once the non-root users are validated.

Global Cluster Option (GCO) require NIC names in specific format [3641586]

The `gcoconfig` script requires the NIC names in the letters followed by numbers format. For example, NIC names can be `eth0`, `eth123`, `xyz111` and so on. The

script fails to configure GCO between NICs which do not comply with this naming format.

Workaround: Rename the NIC name and use the letters followed by numbers format to configure GCO.

Issues related to the bundled agents

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.

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.

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.

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, 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 concurrency violation mechanism handles this scenario appropriately.

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 *Symantec 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 CFSMount Errors [2564411]

When you try to bring the FireDrill service group online through Java Console or `hagrp -online` command, the CFSMount 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 einviroment.

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.

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 observed when host is moved to maintenance mode [2735283]

When a Red Hat Enterprise Virtualization host running a virtual machine is moved to maintenance state, the virtual machine migration is initiated by RHEV. VCS detects the migration according to virtual machine state, such as "migrating". Due to timing issue RHEV Manager occasionally sends the virtual machine state as "up" even if the migration is in progress. Due to this state, the resource is marked ONLINE on the node to which it migrates and may cause concurrency violation.

Workaround: No workaround.

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 the storage. This occurs because of LVM2 commands return success even if all the paths to storage are disabled. Moreover, the LVMVolumegroup and LVMLogicalVolume agents report the resource state as ONLINE.

Workaround: Verify the multi-pathing environment and make sure that all the read and write operations to the disk are blocked when all paths to the storage are disabled.

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.

Unexpected behavior in VCS observed while taking the disk online [3123872]

If the VMwareDisks resource is configured for a disk connected to another virtual machine outside of an ESX cluster and if you bring the disk online on the configured node, you may observe unexpected behavior of VCS (like LLT connection break). The behavior is due to a known issue in VMware.

Workaround: Remove the disk from the other virtual machine and try again.

LVMLogicalVolume agent clean entry point fails to stop logical volume if storage connectivity is lost [3118820]

If storage connectivity is lost on a system on which the LVM resources are in ONLINE state and a volume is mounted using the Mount resource, LVMVolumeGroup agent monitor entry point detects the loss of connectivity and returns the resource state as offline. This causes agent framework to call clean entry point of LVMVolumeGroup agent; however, the state of the resource stays online. Agent framework waits for the clean entry point to return success so that the resource can be moved to the offline|faulted state. At this stage, the clean entry point fails as it is not able deactivate and export the volume group because the logical volume is mounted. There is no option available to forcefully deactivate and export the volume group. Hence, the service groups get stuck in this state. Even if the storage connectivity is restored, the problem does not resolve because the logical volume remains mounted. If the logical volume is unmounted, then the LVMVolumeGroup resource goes into FAULTED state and service group fails over.

Workaround: Manually unmount the logical volume.

VM goes into paused state if the source node loses storage connectivity during migration [3085214]

During virtual machine migrations in a RHEV environment, the VM may freeze in paused state if the source host loses storage connectivity.This issue is specific to RHEV environment.

Workaround: No workaround.

Virtual machine goes to paused state during migration if the public network cable is pulled on the destination node [3080930]

The virtual machine goes into paused state during migration if the public network cable is pulled on the destination node. This behavior depends on the stage at which the migration is disrupted. The virtual machine rolls back to the source node if the network cable is pulled during migration. Resource on the source node reports this as an online virtual machine that is in running state. On the destination node, the virtual machine goes into shut-off state.

If the virtual machine migration gets disrupted during the transfer from source to destination, it may happen that the virtual machine remains in paused state on the source node. In such a case, you must manually clear the state of the virtual machine and bring the it online on any one node.

This operational issue is a behavior of the technology and has no dependency on VCS. This behavior is observed even if the migration is invoked outside VCS control. Due to the disruption in virtual machine migration, it may happen that the locking mechanism does not allow the virtual machine to run on any host, but again, this is a virtualization technology issue.

Workaround: No workaround. Refer to the virtualization documentation.

NFS resource faults on the node enabled with SELinux and where rpc.statd process may terminate when access is denied to the PID file [3248903]

If SELinux is enabled on a system, it blocks `rpc.statd` process from accessing the PID file at `/var/run/rpc.statd.pid`. This may cause `rpc.statd` process to terminate on the system. NFS resource, which monitors the NFS services on the node, detects this and returns unexpected OFFLINE as `statd` process is not running. This is because SELinux does not allow `statd` process to access the PID file and may occur with VCS monitoring the NFS resources.

Workaround: There is no prescribed solution available from Red Hat. You can perform the following steps as a workaround for this issue:

1. Disable SELinux.
2. Use `audit2allow` utility to create a policy to allow access to `rpc.statd` process.
3. Run `semodule -i <policy_module_name>.pp` to install the policy module generated by `audit2allow` utility.

NFS client reports I/O error because of network split brain [3257399]

When network split brain occurs, the failing node may take some time to panic. As a result, the service group on the failover node may fail to come online as some of the resources (such as IP resource) are still online on the failing node. The disk group on the failing node may also get disabled but IP resource on the same node continues to be online.

Workaround: Configure the preonline trigger for the service groups containing DiskGroup resource with reservation on each system in the service group:

1 Copy the preonline_ipc trigger from `/opt/VRTSvcs/bin/sample_triggers/VRTSvcs` to `/opt/VRTSvcs/bin/triggers/preonline/` as T0preonline_ipc:

```
# cp /opt/VRTSvcs/bin/sample_triggers/VRTSvcs/preonline_ipc
/opt/VRTSvcs/bin/triggers/preonline/T0preonline_ipc
```

2 Enable the preonline trigger for the service group.

```
# hagrp -modify <group_name> TriggersEnabled
PREONLINE -sys <node_name>
```

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.

The CoordPoint agent faults after you detach or reattach one or more coordination disks from a storage array (3317123)

After you detach or reattach a coordination disk from a storage array, the CoordPoint agent may fault because it reads an older value stored in the I/O fencing kernel module.

Workaround: Run the `vxfenswap` utility to refresh the registration keys on the coordination points for both server-based I/O fencing and disk-based I/O fencing. But, even if the registrations keys are not lost, you must run the `vxfenswap` utility to refresh the coordination point information stored in the I/O fencing kernel module.

For more information on refreshing registration keys on the coordination points for server-based and disk-based I/O fencing, refer to the *Symantec Cluster Server Administrator's Guide*.

Mount resource does not support spaces in the MountPoint and BlockDevice attribute values [3335304]

Mount resource does not handle intermediate spaces in the configured MountPoint or BlockDevice attribute values.

Workaround: No workaround.

LVMVolumeGroup agent does not report UNKNOWN if invalid volume group is configured [3246748]

VCS LVMVolumeGroup agent can identify whether configured volume group is valid or not based on the error code returned by the native LVM2 commands. If a DiskReservation resource is configured along with LVMVolumeGroup resource and is in online state on any of the node, then LVM2 commands from the node on which the resource is offline fail to read the disk. Hence, LVMVolumeGroup agent cannot validate the volume group. Currently, if LVM2 command returns failure, the resource state is returned as OFFLINE.

Workaround: No workaround.

Manual configuration of RHEVMInfo attribute of KVMGuest agent requires all its keys to be configured [3277994]

The RHEVMInfo attribute of KVMGuest agent has 6 keys associated with it. When you edit main.cf to configure RHEVMInfo attribute manually, you must make sure that all the keys of this attribute are configured in main.cf. If any of its keys is left unconfigured, the key gets deleted from the attribute and agent does not receive the complete attribute. Hence, it logs a Perl error `Use of uninitialized value` in the engine log. This is due to the VCS engine behavior of handling the attribute with key-value pair.

Workaround: Use `ha` commands to add or modify RHEVMInfo attribute of KVMGuest resource.

NFS lock failover is not supported on Linux [3331646]

If a file is locked from an NFS client, the other client may also get the lock on the same file after failover of the NFS share service group. This is because of the changes in the format of the lock files and the lock failover mechanism.

Workaround: No workaround.

SambaServer agent may generate core on Linux if LockDir attribute is changed to empty value while agent is running [3339231]

If LockDir attribute is changed to an empty value while agent is running and debugging is enabled, the logging function may access invalid memory address resulting in SambaServer agent to generate core dump.

Workaround: When LockDir attribute is changed while agent is running, ensure that its new value is set to a non-empty valid value.

Independent Persistent disk setting is not preserved during failover of virtual disks in VMware environment [3338702]

VMwareDisks agent supports Persistent disks only. Hence, Independent disk settings are not preserved during failover of virtual disk.

Workaround: No workaround.

LVMLogicalVolume resource goes in `UNABLE TO OFFLINE` state if native LVM volume group is exported outside VCS control [3606516]

If you export the LVM volume group without stopping LVM logical volumes, the LVMLogicalVolume resource falsely reports online. If offline is initiated for LVMLogicalVolume resource, it fails as the volume group was not exported cleanly and LVMLogicalVolume Agent fails to deactivate the logical volume causing LVMLogicalVolume to go in `UNABLE TO OFFLINE` state.

Workaround: Make sure volume group is deactivated and exported using VCS or manually deactivate the LVM logical volumes.

DiskGroup resource online may take time if it is configured along with VMwareDisks resource [3638242]

If a service group is configured with VMwareDisks and DiskGroup resource, the DiskGroup resource may take time to come online during the service group online. This is because VxVM takes time to recognize a new disk that is attached by VMwareDisks resource. A VMwareDisks resource attaches a disk to the virtual machine when the resource comes online and a DiskGroup resource, which depends on VMwareDisks resource, tries to import the disk group. If `vxconfigd` does not detect the new disk attached to the virtual machine, the DiskGroup resource online fails with the following error message because the resource is not up even after the resource online is complete.

```
VCS ERROR V-16-2-13066 ... Agent is calling clean for resource(...)
```

Workaround: Configure OnlineRetryLimit to appropriate value.

For example, if the DiskGroup resource name is `res_rawdg`:

```
# hares -override res_rawdg OnlineRetryLimit
# hares -modify res_rawdg  OnlineRetryLimit 2
```

SFCache Agent fails to enable caching if cache area is offline [3644424]

SFCache agent cannot enable caching if cache area associate with this particular object is in offline state. User need to manually online the cache area to make sure that caching can be enabled/disabled.

Workaround: Online the cache area using `sfcache` command

```
# sfcache online <cache_area_name>
```

RemoteGroup agent may stop working on upgrading the remote cluster in secure mode [3648886]

RemoteGroup agent may report the resource state as UNKNOWN if the remote cluster is upgraded to VCS 6.2 in secure mode.

Workaround: Restart the RemoteGroup agent.

VMwareDisks agent may fail to start or storage discovery may fail if SELinux is running in enforcing mode [3106376]

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'
```

Issues related to the VCS database agents

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 Symantec 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 *Symantec 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.

Oracle agent fails to offline pluggable database (PDB) resource with PDB in backup mode [3592142]

If the PDB is in backup mode and if you attempt to offline the corresponding PDB resource, this will cause PDB resource to go into "Unable to Offline" state.

Workaround: Manually remove the PDB from the backup mode before attempting to take the PDB resource offline.

Clean succeeds for PDB even as PDB staus is UNABLE to OFFLINE [3609351]

Oracle does not allow any operation on a PDB when the PDB is in backup mode. This is an expected behavior of Oracle. Therefore, a shutdown fails when it is initiated on a PDB in backup mode and returns an UNABLE TO OFFLINE status for the PDB. If PDB is removed from the backup mode using the SQL script, the agent framework is unable to change the UNABLE TO OFFLINE status of the PDB as clean is called. Since Oracle does not differntiate between clean and offline for PDB, clean succeeds for the PDB in spite of being in UNABLE TO OFFLINE state.

Workaround: No workaround.

Second level monitoring fails if user and table names are identical [3594962]

If the table inside CDB has same name as the user name, second level monitoring fails and Oracle agent fails to update the table. For example, if user name is `c##pdbuser1` and table is created as `c##pdbuser1.vcs`, then Oracle agent is unable to update it.

Workaround: Avoid having identical user and CDB table names.

Monitor entry point times out for Oracle PDB resources when CDB is moved to suspended state in Oracle 12.1.0.2 [3643582]

In Oracle-12.1.0.2.0, when CDB is in SUSPENDED mode, then the SQL command for PDB view (`v$pdbs`) hangs. Due to this, the monitor entry point in PDB gets timed out and there is no issue found in oracle-12.1.0.1.0 .

Workaround: No workaround.

Oracle agent fails to online and monitor Oracle instance if threaded_execution parameter is set to true [3644425]

In Oracle 12c, the threaded execution feature is enabled. The multithreaded Oracle Database model enables Oracle processes to execute as operating system threads in separate address spaces. If Oracle Database 12c is installed, the database runs in the process mode. If you set a parameter to run the database in threaded mode, some background processes on UNIX and Linux run with each process containing one thread, whereas the remaining Oracle processes run as threads within the processes.

When you enable this parameter, Oracle agent is unable to check smon (mandatory process check) and lgwr (optional process check) processes which where tradtionally used for monitoring and which now run as threads.

Workaround: Disable the threaded execution feature as it is no supported on Oracle 12C.

Issues related to the agent framework

Agent may fail to heartbeat under heavy load [2073018]

An agent may fail to heart beat 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.

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.

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

For a resource registered with AMF, if you run `hagrp -offline` or `hagrp -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.

Delayed response to VCS commands observed on nodes with several resources and system has high CPU usage or high swap usage [3208239]

You may experience a delay of several minutes in the VCS response to commands if you configure large number of resources for monitoring on a VCS node and if the CPU usage is close to 100 percent or swap usage is very high.

Some of the commands are mentioned below:

- `# hares -online`
- `# hares -offline`
- `# hagrp -online`
- `# hagrp -offline`
- `# hares -switch`

The delay occurs as the related VCS agent does not get enough CPU bandwidth to process your command. The agent may also be busy processing large number of pending internal commands (such as periodic monitoring of each resource).

Workaround: Change the values of some VCS agent type attributes which are facing the issue and restore the original attribute values after the system returns to the normal CPU load.

1. Back up the original values of attributes such as MonitorInterval, OfflineMonitorInterval, and MonitorFreq of IMF attribute.
2. If the agent does not support Intelligent Monitoring Framework (IMF), increase the value of MonitorInterval and OfflineMonitorInterval attributes.

   ```
   # haconf -makerw
   # hatype -modify <TypeName> MonitorInterval <value>
   # hatype -modify <TypeName> OfflineMonitorInterval <value>
   # haconf -dump -makero
   ```

   Where <TypeName> is the name of the agent with which you are facing delays and <value> is any numerical value appropriate for your environment.

3 If the agent supports IMF, increase the value of MonitorFreq attribute of IMF.

```
# haconf -makerw
# hatype -modify <TypeName> IMF -update MonitorFreq <value>
# haconf -dump -makero
```

Where <value> is any numerical value appropriate for your environment.

4 Wait for several minutes to ensure that VCS has executed all pending commands, and then execute any new VCS command.

5 If the delay persists, repeat step 2 or 3 as appropriate.

6 If the CPU usage returns to normal limits, revert the attribute changes to the backed up values to avoid the delay in detecting the resource fault.

CFSMount agent may fail to heartbeat with VCS engine and logs an error message in the engine log on systems with high memory load [3060779]

On a system with high memory load, CFSMount agent may fail to heartbeat with VCS engine resulting into V-16-1-53030 error message in the engine log.

VCS engine must receive periodic heartbeat from CFSMount agent to ensure that it is running properly on the system. The heartbeat is decided by AgentReplyTimeout attribute. Due to high CPU usage or memory workload (for example, swap usage greater than 85%), agent may not get enough CPU cycles to schedule. This causes heartbeat loss with VCS engine and as a result VCS engine terminates the agent and starts the new agent. This can be identified with the following error message in the engine log:

```
V-16-1-53030 Termination request sent to CFSMount
agent process with pid %d
```

Workaround: Increase the AgentReplyTimeout value and see if CFSMount agent becomes stable. If this does not resolve the issue then try the following workaround. Set value of attribute NumThreads to 1 for CFSMount agent by running following command:

```
# hatype -modify CFSMount NumThreads 1
```

Even after the above command if CFSMount agent keeps on terminating, report this to Symantec support team.

Logs from the script executed other than the agent entry point goes into the engine logs [3547329]

The agent logs of C-based and script-based entry points get logged in the agent log when the attribute value of LogViaHalog is set to 1 (one). To restore to the older logging behavior in which C-based entry point logs were logged in agent logs and script-based entry point logs were logged in engine logs, you can set the LogViaHalog value as 0 (zero). However, it is observed that some C-based entry point logs continue to appear in the engine logs even when LogViaHalog is set to 1 (one). This issue is observed on all the database agents.

Workaround: No workaround.

Issues related to global clusters

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.

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.

LLT known issues

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

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.

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 `lltab` 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 `lltab` file, and update the file before you restart LLT.

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.

If you manually re-plumb (change) the IP address on a network interface card (NIC) which is used by LLT, then LLT may experience heartbeat loss and the node may panic (3188950)

With the LLT interfaces up, if you manually re-plumb the IP address on the NIC, then the LLT link goes down and LLT may experience heartbeat loss. This situation may cause the node to panic.

Workaround: Do not re-plumb the IP address on the NIC that is currently used for LLT operations. Take down the stack before you re-plumb the IP address for the LLT interface.

A network restart of the network interfaces may cause heartbeat loss for the NIC interfaces used by LLT

A network restart may cause heartbeat loss of the network interfaces configured LLT. LLT configured for UDP or LLT configured for RDMA may experience loss of heartbeat between the interfaces, which may cause the node to panic.

Workaround: Recommendations before you restart the network:

- Assess the effect of a network restart on a running cluster that is using LLT over RDMA or LLT over UDP.
- Do not use the network restart functionality to add or configure a new NIC to the system.
- If you are using the network restart functionality, make sure that the LLT interfaces are not affected.
- Increase the `llt-peerinact` time to a higher value to allow network restart to complete within that time.
  Run the `# lltconfig -T peerinact:6000` command to increase the peerinact time to 1 minute.

When you execute the /etc/init.d/llt start script to load the LLT module, the syslog file may record messages related to kernel symbols associated with Infiniband (3136418)

When you execute /etc/init.d/llt start to start the LLT module on some Linux kernel versions, the syslog file may display the following messages for multiple such symbols:

```
kernel: llt: disagrees about version of symbol ib_create_cq
kernel: llt: Unknown symbol ib_create_cq
```

The LLT module is shipped with multiple module *.ko files, which are built against different kernel versions. If the kernel version on the node does not match the kernel version against which the LLT module is built, the LLT module fails to load and logs RDMA-related messages in the syslog file. In this case, the kernel logs these messages. The modinst script loads the compatible module on the system and starts LLT without any issues.

Workaround: Rearrange the kernel versions in the `/opt/VRTSllt/kvers.lst` file such that the first line displays the kernel version that is most likely to be compatible with the kernel version on the node. This rearrangement allows the modinst script to load the best possible kernel module first. Therefore, the warning message is less likely to appear.

Dependency may get overruled when uninstalling multiple RPMs in a single command [3563254]

When performing uninstallation of multiple RPMs through a single comment, the system identifies and follows the specified dependency among the RPMs as the uninstallation progresses. However, if the pre-uninstallation script fails for any of the RPMs, the system does not abort the task but instead uninstalls the remaining RPMs.

For example, if you run `rpm -e VRTSllt VRTSgab VRTSvxfen` where the RPMs have a dependency between each other, the system bypasses the dependency if the pre-uninstallation script fails for any RPM.

Workaround: Uninstall the RPMs independently.

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.

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 *Symantec Cluster Server Administrator's Guide* for more details.

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.

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.

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 vxfend 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 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).

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 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 VRTSat 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
```

Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later [3226290]

The issue exists because the VCS 5.1SP1 release version does not support separate directories for truststores. However, VCS version 6.0 and later support separate directories for truststores. Because of this mismatch in support for truststores, you cannot set up trust between client systems and servers.

Workaround: Set up trust manually between the coordination point server and client systems using the `cpsat` or `vcsat` command so that the servers and client systems can communicate in a secure mode.

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.

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

When you configure fencing in customized mode and do no 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 vxcpserv 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.

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 cooridnator disk group

Workaround: There is no workaround for this issue.

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.

The vxfenswap utility deletes comment lines from the `/etc/vxfemode` file, if you run the utility with hacli option (3318449)

The vxfenswap utility uses RSH, SSH, or hacli protocol to communicate with peer nodes in the cluster. When you use vxfenswap to replace coordination disk(s) in disk-based fencing, vxfenswap copies `/etc/vxfenmode` (local node) to `/etc/vxfenmode` (remote node).

With the hacli option, the utility removes the comment lines from the remote `/etc/vxfenmode` file, but, it retains comments in the local `/etc/vxfenmode` file.

Workaround: Copy the comments manually from local `/etc/vxfenmode` to remote nodes.

When you configure CP server only for HTTPS-based communication, the `engine_A.log` displays a misleading message (3321101)

The `engine_A.log` file displays the following message when you configure CP server only for HTTPS-based communication but not for IPM-based communication.

```
No VIP for IPM specified in /etc/vxcps.conf
```

Workaround: Ignore the message.

The `vxfentsthdw` utility may not run on systems installed with partial SFHA stack [3333914]

The `vxfentsthdw` utility runs if the SFHA stack and VCS are fully installed with properly configured SF and VxVM. It also runs if the entire SFHA stack and VCS are not installed. However, partial installs where SF is installed and configured but VCS is not installed is not supported. The utility will display an error with the `-g` or `-c` options.

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

When a client node goes down, for reasons such as node panic, I/O fencing does not come up on that client node after node restart (3341322)

This issue happens when one of the following conditions is true:

- Any of the CP servers configured for HTTPS communication goes down.
- The CP server service group in any of the CP servers configured for HTTPS communication goes down.
- Any of the VIPs in any of the CP servers configured for HTTPS communication goes down.

When you restart the client node, fencing configuration starts on the node. The fencing daemon, vxfend, invokes some of the fencing scripts on the node. Each of these scripts has a timeout value of 120 seconds. If any of these scripts fails, fencing configuration fails on that node.

Some of these scripts use `cpsadm` commands to communicate with CP servers. When the node comes up, `cpsadm` commands try to connect to the CP server using VIPs for a timeout value of 60 seconds. So, if the multiple `cpsadm` commands that are run within a single script exceed the timeout value, then the total timeout value exceeds 120 seconds, which causes one of the scripts to time out. Hence, I/O fencing does not come up on the client node.

Note that this issue does not occur with IPM-based communication between CP server and client clusters.

Workaround: Fix the CP server.

VCS fails to take virtual machines offline while restarting a physical host in RHEV and KVM environments (3320988)

In RHEV and KVM environments, the virtualization daemons `vdsmd` and `libvirtd`required to operate virtual machines are stopped before VCS is stopped during a reboot of the physical host. In this scenario, VCS cannot take the virtual machine resource offline and therefore the resource fails to stop. As a result , LLT, GAB and fencing fail to stop. However, the virtual network bridge is removed leading to the loss of cluster interconnects and causing a split-brain situation.

Workaround: If the virtual network bridge is not assigned to any virtual machine, remove the virtual bridge and configure LLT to use the physical interface. Alternatively, before initiating a reboot of the physical host, stop VCS by issuing the `hastop -local` command. The `-evacuate` option can be used to evacuate the virtual machines to another physical host.

Fencing may panic the node while shut down or restart when LLT network interfaces are under Network Manager control [3627749]

When the LLT network interfaces are under Network Manager control, then shutting down or restarting a node may cause fencing race resulting in a panic. On RHEL, VCS requires that LLT network interfaces are not put under Network Manager control, as it might cause problems when a node is shut down or restarted. During shutdown, the Network Manager service might stop before the VCS shutdown scripts are called. As a result, fencing race is triggered and the losing sub-cluster panics.

Workaround: Either exclude the network interfaces to be used by LLT from Network Manager control or disable the Network Manager service before configuring LLT. Please refer to the Red Hat documentation to do the same.

The `vxfenconfig -l` command output does not list Coordinator disks that are removed using the `vxdmpadm exclude dmpnodename=<dmp_disk/node>` command [3644431]

After you remove a Coordinator disk used by fencing or fencing disk group by running the `vxdmpadm exclude dmpnodename=<dmp_disk/node>` command, the removed disk is not listed in the `vxfenconfig -l` command output.

In case of a split brain, the `vxfen` program cannot use the removed disk as a coordination point in the subsequent fencing race.

Workaround: Run the `vxdmpadm include dmpnodename=<dmp_disk/node>` command to again enable the dmp disk. This disk will show up in subsequent `vxfencondig -l` output.

Symantec Cluster Server agents for Volume Replicator known issues in 6.2.1

The following are new additional Symantec Cluster Server agents for Volume Replicator known issues in 6.2.1 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
        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 vxfsckd
        //          {
        //           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 vxfsckd
        //          {
        //           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.

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 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.

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.

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>
```

VCS engine shows error for cancellation of reaper when Apache agent is disabled [3043533]

When `haimfconfig` script is used to disable IMF for one or more agents, the VCS engine logs the following message in the engine log:

```
AMF imf_getnotification ERROR V-292-2-193
Notification(s) canceled for this reaper.
```

This is an expected behavior and not an issue.

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.

Agent cannot become IMF-aware with agent directory and agent file configured [2858160]

Agent cannot become IMF-aware if Agent Directory and Agent File are configured for that agent.

Workaround: No workaround.

ProPCV fails to prevent a script from running if it is run with relative path [3617014]

If the absolute path is registered with AMF for prevention and the script is run with the relative path, AMF fails to prevent the script from running.

Workaround: No workaround.

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/VRTSvcs/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

Configuring application for high availability with storage using VCS wizard may fail on a VMware virtual machine which is configured with more than two storage controllers [3640956]

Application configuration from VCS wizard may fail on VMware virtual machine which is configured with multiple SCSI controllers.

Workaround: There is no workaround available.

Agent kill on source during migration may lead to resource concurrency violation (3042499)

In the case of a migration initiated outside Symantec Cluster Server (VCS) control, there is a very small window in which the agent restart might not be able to recognize the migration event. As this is initiated outside VCS, there is no way to synchronize the agent restart and the migration., Also, there is no intermediate state in KVM that can indicate that the event was a migration. This problem does not occur in Red Hat Enterprise Virtualization (RHEV), as there are clear states visible that can specify the virtual machine events. This is applicable to KVM environment only.

Workaround: There is no workaround for this issue.

Host fails to reboot when the resource gets stuck in ONLINE|STATE UNKNOWN state [2738864]

In a Red Hat Enterprise Virtualization environment, if a host reboot is performed on which the KVMGuest resource monitoring the virtual machine is ONLINE, then the host reboot fails. This is because the VDSM is getting stopped before VCS could shutdown the virtual machine. In this case, the virtual machine state remains ONLINE|STATE UNKNOWN, and hence VCS stop fails eventually failing the host reboot as well.

Workaround: Switch the service group to other node before initiating a host reboot.

VM state is in PAUSED state when storage domain is inactive [2747163]

If the storage domain associated with the running virtual machine becomes inactive, the virtual machine may go to **paused** state.

Workaround: Make sure that the storage domain is always active when running virtual machine.

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.

Workaound: Restart the `libvirtd` process and run:

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

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.

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.

Increased memory usage observed even with no VM running [2734970]

Increased memory usage was observed on the hosts even when VMs were either not running or had stopped. This is due to the RHEV behavior.

Workaround: No workaround.

Resource faults when it fails to ONLINE VM beacuse 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.

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.

Virtual machine may return the not-responding state when the storage domain is inactive and the data center is down (2848003)

In a Red Hat Enterprise Virtualization Environment, if the storage domain is in an inactive state and the data center is in down state, the virtual machine may return a not-responding state and the KVMGuest resource in OFFLINE state.

Workaround: To resolve this issue:

1. Activate the storage domain in RHEV-M.
2. Check that the data center is in the up state.

Guest virtual machine may fail on RHEL 6.1 if KVM guest image resides on CVM-CFS [2659944]

If a KVM guest image file resides on CVM-CFS, the migration of that guest virtual machine may fail with "Permission Denied" error on RHEL 6.1. This causes guest virtual machine to go in "shut-off" state on both source and destination node, and the associated VCS KVMGuest.

Workaround: Make sure that the guest image file is having 777 permission.

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
```

CD ROM with empty file vmPayload found inside the guest when resource comes online [3060910]

When you unset the DROpts attribute on a KVMGuest resource and online the resource on the host, a CD ROM with an empty file vmPayload is available inside the guest.

The KVMGuest agent adds a CD ROM to the virtual machine configuration when you online a KVMGuest resource with the DROpts attribute set. The CD ROM carries some site-specific parameters to be used inside the guest. When you offline the same resource, the agent removes the CD ROM, but for some reason, the CD ROM does not get removed completely. If you unset the DROpts attribute and online the resource later, a CD ROM with an empty file vmPayload continues to be available inside the guest.

Workaround: This does not impact the functionality of the virtual machine in any way and can be ignored.

VCS fails to start virtual machine on another node if the first node panics [3042806]

In the KVM environment, if a node on which a virtual machine is running panics, then VCS fails to start that virtual machine on another node. This issue occurs because KVM Hypervisor is not able to acquire lock on the virtual machine. This issue is due to KVM Hypervisor behavior and is very rarely observed.

Workaround: Restart libvirtd process to resolve this issue. Command to restart libvirtd:

```
# service libvirtd restart
```

VM fails to start on the target node if the source node panics or restarts during migration [3042786]

If a virtual machine (VM) migration is initiated and the source node (node on which VM was running) panics or is restarted forcefully, VM fails to start on any other node in a KVM environment. This issue is due to the KVM locking mechanism. The VM start fails with the following error:

```
error: Failed to start domain VM1
error: Timed out during operation: cannot acquire state change lock
```

Workaround: Restart (kill and start) the libvirtd daemon on the second node using the following command:

```
# service libvirtd restart
```

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.

Cluster communication breaks when you revert a snapshot in VMware environment [3409586]

If VCS is running on the guest operating system when a VMware virtual machine snapshot is taken, the virtual machine snapshot contains the run-time state of the cluster. When you restore the snapshot, the state of the cluster which is restored can be inconsistent with other nodes of the cluster. Due to the inconsistent state, VCS is unable to communicate with other nodes of the cluster.

Workaround: Before you take a snapshot of the virtual machine, Symantec recommends that you stop VCS services running inside the virtual machine.

VCS Cluster Configuration wizard issues

VCS Cluster Configuration wizard does not automatically close in Mozilla Firefox [3281450]

You can use the `haappwizard` utility to launch the Symantec High Availability wizard to configure application monitoring with Symantec Cluster Server (VCS) on Linux systems. If you configure the utility to launch the wizard in Mozilla Firefox browser, the browser session does not automatically close after the VCS configuration is complete.

Workaround: Use one of the following workarounds:

- Close the Mozilla Firefox browser session once the wizard-based configuration steps are complete.
- Specify a different browser while configuring the `haappwizard` utility.

Configuration inputs page of VCS Cluster Configuration wizard shows multiple cluster systems for the same virtual machine [3237023]

The **Configuration inputs** panel of the VCS Cluster Configuration wizard shows multiple cluster systems for the same virtual machine. This occurs because the value specified for the node in the SystemList attribute is different than the one returned by the `hostname` command.

Workaround: Ensure that the value specified for the node in the SystemList attribute and the one returned by the `hostname` command is the same.

VCS Cluster Configuration wizard fails to display mount points on native LVM if volume groups are exported [3341937]

On storage selection page of application wizard, mount points mounted on native LVM devices are not shown. If you have one or more native LVM volume groups, and one of them is exported, the application wizard fails to detect mount points configured on these devices.

Workaround: Ensure that you do not have any native volumes exported if you want to configure application that uses native LVM storage.

IPv6 verification fails while configuring generic application using VCS Cluster Configuration wizard [3614680]

The VCS Cluster Configuration wizard fails to check whether IPv6 IP is already plumbed while configuring a generic application through the Virtual IP page. The wizard does neither displays a warning if IPv6 IP is already plumbed elsewhere nor indicates whether it is reachable through a ping.

Workaround: Manually ensure that IPv6 is not plumbed elsewhere on the network before configuring the generic application through the wizard.

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 88.

Limitations related to bundled agents

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 on SLES 11, SLES 11SP2
- ext4 filesystem configured on VxVM
- xfs filesystem configured on VxVM
- btrfs 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.

Volumes in a disk group start automatically irrespective of the value of the StartVolumes attribute in VCS [2162929]

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.

Application agent limitations

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

Campus cluster fire drill does not work when DSM sites are used to mark site boundaries [3073907]

The campus cluster FireDrill agent currently uses the SystemZones attribute to identify site boundaries. Hence, campus cluster FireDrill is not supported in DSM enabled environment.

Workaround: Disable DSM and configure the SystemZones attribute on the application service group to perform the fire drill.

Mount agent reports resource state as OFFLINE if the configured mount point does not exist [3435266]

If a configured mount point does not exist on a node, then the Mount agent reports the resource state as OFFLINE instead of UNKNOWN on that particular node. If an attempt is made for onlining the resource, it fails on that node as the mount point does not exist.

Workaround: Make sure that configured mount point exists on all nodes of the cluster or alternatively set the CreateMntPt attribute value of Mount agent to 1. This will ensure that if a mount point does not exist then it will create while onlining the resource.

Limitation of VMwareDisks agent to communicate with the vCenter Server [3528649]

If VMHA is not enabled and the host ESX faults then even after the disks are attached to the target virtual machine, they remain attached to the failed virtual machine. This issue occurs because the request to detach the disks fails since the host ESX itself has faulted. The agent then sends the disk attach request to the vCenter Server and attaches the disks to the target virtual machine. Even though the application availability is not impacted, the subsequent restart of the faulted virtual machine fails. This issue occurs because of the stale link between the virtual machine and the disks attached. Even though the disks are now attached to the target virtual machine the stale link with the failed virtual machine still exists.

Workaround: Detach the disks from the failed virtual machine and then restart the virtual machine.

Limitations related to VCS engine

Loads fail to consolidate and optimize when multiple groups fault [3074299]

When multiple groups fault and fail over at the same time, the loads are not consolidated and optimized to choose the target systems.

Workaround: No workaround.

Preferred fencing ignores the forecasted available capacity [3077242]

Preferred fencing in VCS does not consider the forecasted available capacity for fencing decision. The fencing decision is based on the system weight configured.

Workaround: No workaround.

Failover occurs within the SystemZone or site when BiggestAvailable policy is set [3083757]

Failover always occurs within the SytemZone or site when the BiggestAvailable failover policy is configured. The target system for failover is always selected based on the biggest available system within the SystemZone.

Workaround: No workaround.

Load for Priority groups is ignored in groups with BiggestAvailable and Priority in the same group[3074314]

When there are groups with both BiggestAvailable and Priority as the failover policy in the same cluster, the load for Priority groups are not considered.

Workaround: No workaround.

Symantec cluster configuration 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.

Environment variable used to change log directory cannot redefine the log path of the wizard [3609791]

By default, the Symantec cluster configuration wizard writes the logs in `/var/VRTSvcs/log` directory. VCS provides a way to change the log directory through environment variable VCS_LOG, but this does not apply to the logs of VCS wizards.

Workaround: No workaround.

Limitations related to IMF

- IMF registration on Linux for “bind” file system type is not supported.
- In case of SLES11 and RHEL6:
  - IMF must not be enabled for the resources where the BlockDevice can get mounted on multiple MountPoints.

- If the FSType attribute value is nfs, then IMF registration for “nfs” file system type is not supported.

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.

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.

Pluggable database (PDB) online may timeout when started after container database (CDB) [3549506]

PDB may take long time to start when it is started for the first time after starting CDB. As a result, the PDB online initiated using VCS may cause ONLINE timeout and the PDB online process may get cancelled.

Workaround: Increase the OnlineTimeout attribute value of the Oracle type resource.

Security-Enhanced Linux is not supported on SLES distributions

VCS does not support Security-Enhanced Linux (SELinux) on 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. [1272012]

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
```

Host on RHEV-M and actual host must match [2827219]

You must configure the host in RHEV-M with the same name as in the hostname command on a particular host. This is mandatory for RHEV Manager to be able to search the host by hostname.

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 *Symantec 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”.

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.

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.

Limitations related to LLT

This section covers LLT-related software limitations.

Limitation of LLT support over UDP or RDMA using alias IP [3622175]

When configuring the VCS cluster, if alias IP addresses are configured on the LLT links as the IP addresses for LLT over UDP or RDMA, LLT may not work properly.

Workaround: Do not use alias IP addresses over UDP or RDMA.

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.

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 *Symantec 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.

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.

Node may panic if HAD process is stopped by force and then node is shut down or restarted [3640007]

A node may panic if the HAD process running on it is stopped by force and then it is shut down or restarted.This limitation is observed when you perform the following steps on a cluster node:

1 Stop the HAD process with the `force` flag.

```
# hastop -local -force
```

or

```
# hastop -all -force
```

2 Restart or shut down the node.

The node panics because forcefully stopping VCS on the node leaves all the applications, file systems, CVM, and other process online on that node. If the same node is restarted in this starte, VCS triggers a fencing race to avoid data curruption. However, the restarted node loses the fencing race and panics.

Workaround: No workaround.

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.

Virtualization software limitations

This section describes the virtualization software limitations in this release of Symantec Cluster Server (VCS).

Paths cannot be enabled inside a KVM guest if the devices have been previously removed and re-attached from the host

LUNs are exported to the KVM guest via virtio-scsi interface. When some physical link between the host and the SAN array fails for a certain time (45-60 seconds by default), the HBA driver in the host will remove the timed-out devices. When the link is restored, these devices will be re-attached to the host; however, the access from inside the KVM guest to these devices cannot be automatically restored too

without rebooting the system or manually re-attaching the devices. For DMP, these subpaths will remain in DISABLED state.

This is a known limitation of KVM.

Workaround:

From the KVM host, tune the `dev_loss_tmo` parameter of the Fibre Channel ports to a very large value, and set the `fast_io_fail_tmo` parameter to 15.

To restore access to the timed-out devices

1 Add the following lines into `/dev/udev/rules.d/40-kvm-device` file:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add", \
  RUN+="/bin/sh -c 'grep -q off \
  /sys/class/fc_remote_ports/%k/fast_io_fail_tmo;if [ $? -eq 0 ]; \
  then echo 15 > /sys/class/fc_remote_ports/%k/fast_io_fail_tmo 2> \
  /dev/null;fi;'"
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add", \
  RUN+="/bin/sh -c 'echo 8000000 > \
  /sys/class/fc_remote_ports/%k/dev_loss_tmo 2> /dev/null'"
```

2 Create the `/etc/modprobe.d/qla2xxx.conf` file with the following content:

```
options qla2xxx qlport_down_retry=8000000
```

3 Create the `/etc/modprobe.d/scsi_transport_fc.conf` with the following content:

```
options scsi_transport_fc dev_loss_tmo=8000000
```

4 Rebuild the `initrd` file and reboot.

Application component fails to come online [3489464]

In the KVM virtualization environment, if you try to bring an application resource online, the online operation fails. This behavior is observed both from the command line interface as well as the Symantec High Availability view of the Veritas Operations Manager Management Server.

Workaround: Perform the following steps:

1 Set the locale of the operating system (OS) to default value, and then retry the operation. For detailed steps, see OS vendor documentation.

2 Restart High Availability Daemon (HAD).

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

Each product in the Storage Foundation and High Availability Solutions product line includes release notes, an installation guide, and additional documents such as administration and agent guides. In most cases, you may also need to refer to the documentation for the product's components.

The SFHA Solutions documents describe functionality and solutions that apply across the product line. These documents are relevant whichever SFHA Solutions product you use.

Symantec Cluster Server documentation

Table 1-14 lists the documents for Symantec Cluster Server.

Table 1-14 Symantec Cluster Server documentation

Title	File name	Description
Symantec Cluster Server Release Notes	vcs_notes_621_lin.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.
Symantec Cluster Server Installation Guide	vcs_install_621_lin.pdf	Provides information required to install the product.
Symantec Cluster Server Administrator's Guide	vcs_admin_62_lin.pdf	Provides information required for administering the product.
Symantec Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_62_lin.pdf	Provides information about bundled agents, their resources and attributes, and more related information.

Table 1-14 Symantec Cluster Server documentation *(continued)*

Title	File name	Description
Symantec Cluster Server Generic Application Agent Configuration Guide	vcs_gen_agent_62_lin.pdf	Provides notes for installing and configuring the generic Application agent.
Symantec Cluster Server Agent Developer's Guide (This document is available online only.)	vcs_agent_dev_62_unix.pdf	Provides information about the various Symantec agents and procedures for developing custom agents.
Symantec Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_62_lin.pdf	Provides notes for installing and configuring the DB2 agent.
Symantec Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_62_lin.pdf	Provides notes for installing and configuring the Oracle agent.
Symantec Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_62_lin.pdf	Provides notes for installing and configuring the Sybase agent.

Manual pages

The manual pages for Symantec 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 Symantec 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