Veritas™ Cluster Server Release Notes

Linux

6.0

February 2013

Symantec™

Veritas™ Cluster Server Release Notes

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.0

Document version: 6.0.4

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Veritas Cluster Server Release Notes

This document includes the following topics:

About this document

This document provides important information about Veritas Cluster Server (VCS) version 6.0 for Linux. Review this entire document before you install or upgrade VCS.

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

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

https://sort.symantec.com/documents

Component product release notes

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

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

`/product_name/docs`

Symantec recommends copying the `docs` directory on the software media that contains the product guides to the `/opt/VRTS` directory on your system.

This release includes the following component product release notes:

- *Veritas Storage Foundation Release Notes* (6.0)

About Veritas Cluster Server

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

About VCS agents

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

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

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

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

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

About Symantec Operations Readiness Tools

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

SORT can help you do the following:

Prepare for your next installation or upgrade	■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture. ■ Analyze systems to determine if they are ready to install or upgrade Symantec products. ■ Download the latest patches, documentation, and high availability agents from a central repository. ■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems.
Manage risks	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository. ■ Identify and mitigate system and environmental risks. ■ Display descriptions and solutions for hundreds of Symantec error codes.

Improve efficiency	■ Find and download patches based on product version and platform. ■ List installed Symantec products and license keys. ■ Tune and optimize your environment.

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

To access SORT, go to:

https://sort.symantec.com

Important release information

- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:
  http://www.symantec.com/docs/TECH164885
- For the latest patches available for this release, go to:
  http://sort.symantec.com/
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:
  http://www.symantec.com/docs/TECH170013
  Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

Changes introduced in 6.0

This section lists the changes in Veritas Cluster Server 6.0.

Changes related to installation and upgrades

The product installer includes the following changes in 6.0.

Support for product installation using yum on Linux

You can now install any of the Veritas products with yum. Yum installation is supported for Red Hat Enterprise Linux 5 and 6.

See the *Installation Guide* for more information.

The installer can now detect duplicate VCS cluster IDs and can automatically generate cluster IDs

The installer can now detect duplicate VCS cluster IDs and prompt you to select an unused one. It can also generate an unused ID during installation.

The installer can check product versions and hotfixes

You can check the existing product versions using the installer command with the `-version` option before or after you install. After you have installed the current version of the product, you can use the `showversion` script in the /opt/VRTS/install directory to find version information.

You can discover the following information with these commands:

- The installed version of all released Strorage Foundation and High Availability Suite of products
- The missing required RPMs or patches as applicable for platform
- The available updates (including patches or hotfixes) from SORT for the installed products

Depending on the product, the script can identify versions from 4.0 onward.

Using the installer's postcheck option

You can use the installer's postcheck option to diagnose installation-related problems and to provide troubleshooting information.

Rolling upgrade improvements

The rolling upgrade procedure has been streamlined and simplified.

Packaging updates

The following lists the package changes in this release.

- New `VRTSsfcpi60` RPM for product installer scripts

  The `VRTSsfcpi60` RPM is introduced in this release. The VRTSsfcpi60 RPM contains the installer scripts and libraries that the installer uses to install, configure and upgrade Veritas products.

- New `VRTSvbs` RPM

  The `VRTSvbs` RPM enables the VBS command line interface on a Veritas Operations Manager managed host in a Virtual Business Services configuration. For more information, see the *Virtual Business Service–Availability User's Guide*.

- `VRTSvcsag` RPM change
  The `/etc/VRTSvcs/conf/types.cf` and various sample configurations are packaged with the `VRTSvcsag` RPM instead if the `VRTSvcs` RPM. This is to ease hotfix installation for bundled agents that require attribute changes.

For more information, see the *Installation Guide*.

Support for Kernel-based Virtual Machines (KVM) on Linux

Storage Foundation High and Availability Solutions provide configurations to enhance the Kernel-based Virtual Machine (KVM) environment. Storage Foundation High and Availability Solutions 6.0 products are supported on the Red Hat Enterprise Linux (RHEL) 6.1 distribution.

Storage Foundation and High Availability Solutions products provide the following functionality for KVM guest virtual machines:

- Storage visibility
- Storage management
- High availability
- Cluster failover
- Replication support

For implementation information:

See the *Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide for Linux*.

Changes to the VCS engine

Support multiple children in service group dependency

A service group can have dependency on multiple child service groups. All child dependencies must be satisfied for the parent service group to go online. The dependency types which are not supported in a multiple child configuration are "online local hard" and "offline local". The `-propagate` option cannot be used if the dependency tree contains global and/or remote dependency.

For more details, refer to the *Administrator's Guide*.

Single command line option to online or offline service groups in a service group dependency

If you have a parent service group in a service group dependency, you can online the entire dependency tree bottom-up with a single command, without having to online each group manually. For online operation, the command starts to online the service groups from the lowermost service group in dependency tree. Similarly, If you have a child service group in a service group dependency, you can offline the entire dependency tree top-down with a single command, without having to offline each group manually. For offline operation, the command offlines the service group from the top of the dependency tree.

You can use the following commands to online or(and) offline the service groups respectively:

- `hagrp -online -propagate <grp_name> -sys <sys_name>`
- `hagrp -offline -propagate <grp_name> -sys <sys_name>`

Refer to the *manual pages* associated with `hagrp` comand for more information on the `-propagate` option.

Note: The `-propagate` option cannot be used if the dependancy tree contains global and/or remote dependency.

Using the `-propagate` option with `hagrp -online`, all required child service groups are automatically brought online by VCS. Similarly using the `-propagate` option with `hagrp -offline`, all the dependent parent groups are automatically brought offline by VCS.

Ability to send notifications to a wider audience

You can configure users, other than the owners of resources, resource types, service groups, systems, or clusters, as recipients of notifications about events related to a resource, resource type, service group, system, or cluster.

Use the following attributes to configure recipients of notifications:

- ResourceRecipients
- TypeRecipients
- GroupRecipients
- SystemRecipients
- ClusterRecipients

The registered recipients get notifications about the events that have a severity level that is equal to or greater than the level specified in the attribute. For more information, see *Veritas Cluster Server Administrator's Guide*.

Ability to specify a single user across all nodes of VCS cluster

You can add a user to the VCS configuration without specifying the host name, such as **admin** and assign administrator privileges. The admin user can log in using the host-specific credentials and perform any administrative operations.

Thus, you need not add the same user, 'admin', multiple times in the VCS configuration as admin@host1, admin@host2, and so on. Once added, the user can operate from any node of the cluster.

Changes related to IMF

This release includes the following changes to Intelligent Monitoring Framework (IMF):

IMF is enabled by default and has new agent support

The Intelligent Monitoring Framework (IMF) functionality is enabled by default for all agents that can leverage IMF.

Following are the new IMF-aware agents in VCS 6.0:

- DB2udb (provides only PRON IMF support)
- Sybase
- SybaseBk

Enable and disable IMF by automated script for agents

VCS provides the `/opt/VRTSvcs/bin/haimfconfig` script to enable and disable IMF for agents when VCS is either running or stopped. You can use this script to disable IMF for all IMF-aware agents, including bundled agents, enterprise agents, and custom agents. You must run the script once on each node of the cluster.

Note: The automated script restarts the agent, if it is running on the current node, to enable the IMF after confirming with the user.

Prevention of Concurrency Violation (PCV) using IMF

With the new IMF-based Proactive PCV feature, VCS can proactively prevent the same VCS failover service group from coming online on more than one node in

the cluster. Typically, VCS detects such a concurrency violation after it has occurred. This feature is available only for application resources and is disabled by default.

For more information, refer to the *Administrator's Guide*.

Support plug-in for AMF support in custom agents

Script-based custom agents can now leverage IMF functionality by following the steps documented in the *Veritas Cluster Server Agent Developer's Guide*.

Enhanced amfstat utility

The `amfstat` utility is enhanced to display the new event types that AMF driver can support. For more details, refer to the `amfstat` man pages.

Changes related to VCS triggers

This release includes the following changes related to VCS triggers:

- VCS can execute trigger scripts specific to a service group and/or resource. Thus, trigger scripts for multiple objects need not be merged into single trigger script.
- You can enable the `nofailover`, `postonline` and `postoffline` triggers for each system by using the TriggersEnabled attribute.
- You can execute multiple scripts for a trigger. The trigger scripts must be installed inside the trigger directory using the `T<num>` nomenclature. VCS executes the trigger scripts in `T<num>` order.

  For example: If the preonline directory contains the scripts `T00preonline`, `T01preonline`, `T02preonline`, then the script `T00preonline` is executed first, then `T01preonline` and finally, `T02preonline` is executed.
- The agent restarts a faulted resource if the RestartLimit is set. Whenever the agent restarts a resource, VCS invokes the `resrestart` trigger if the TriggerResRestart attribute is set to 1 or if RESRESTART is specified in the TriggersEnabled attribute. Otherwise, VCS invokes the `resstatechange` trigger. See the *Veritas Cluster Server Administrator's Guide* for more information.

Caution: Use of `resstatechange` to indicate restart is being deprecated. In later releases, you must use only the `resrestart` trigger to indicate restarting of resources.

New attributes

The following sections describe the attributes introduced in VCS 6.0, 5.1SP1, VCS 5.1, and VCS 5.0MP3.

Attributes introduced in VCS 6.0

DNS agent attributes

- UseGSSAPI: Use this attribute if the DNS server that is configured is a Windows DNS server. The agent uses "-g" option with `nsupdate` command if this attribute is set to 1.
- RefreshInterval: This attribute represents the time interval in seconds after which the DNS agent attempts to refresh the resource records (RRs) on the DNS servers.
- CleanRRKeys: Use this attribute to direct the online agent function to clean up all the existing DNS records for the configured keys before adding new records. The default value (0) disables this behavior.

Cluster-level attribute

- EnableVMAutoDiscovery: Enables or disables auto discovery of virtual machines. By default, auto discovery of virtual machines is disabled.
- SystemRebootAction: Determines whether frozen service groups are ignored on system reboot.

Service group attributes

- OnlineClearParent: When this attribute is enabled for a service group and the service group comes online or is detected online, VCS clears the faults on all online type parent groups, such as online local, online global, and online remote.
- ProPCV: Indicates whether the service group is proactively prevented from concurrency violation for ProPCV-enabled resources.
- TriggerPath: Enables you to customize the trigger path.
- TriggerResRestart: Determines whether or not to invoke the restart trigger if resource restarts.
- TriggersEnabled: Determines if a specific trigger is enabled on a node or not.

Resource type attributes:

- AdvDbg: Enables activation of advanced debugging.

Sybase agent attributes:

- Quorum_dev: The quorum device manages the cluster membership, stores cluster configuration data and contains information shared among server

instances and nodes. It must be a disk accessible to all nodes in the cluster. Specify fully qualified quorum device name. Note: This attribute should be specified only for cluster edition. For example: /dev/vx/rdsk/Sybase_install_dg/quorum_vol.

- interfaces_File: Specifies the location of interfaces file for the Sybase instance. If this attribute is configured, [-I interfaces file] option is used when connecting to the isql session. If this attribute is not configured, the agent does not use the -I option.
- ShutdownWaitLimit: Specifies the maximum number of seconds for which the agent waits for the Sybase instance to stop after issuing the `shutdown with wait` command, and before attempting to issue the `kill -15 <data server-pid>` command, if required.
- DelayAfterOnline: Specifies the number of seconds that elapse after the Online entry point is complete and before the next monitor cycle is invoked.
- DelayAfterOffline: Specifies the number of seconds that elapse after the Offline entry point is complete and before the next monitor cycle is invoked.

SybaseBk agent attribute:

- interfaces_File: Specifies the location of the interfaces file for the Sybase instance. If this attribute is configured, [-I inerfaces file] option is used when connecting to the isql session. If this attribute is not configured, the agent does not use the -I option.

Oracle agent attributes

- DBName: Specifies the database name. This attribute is required only for a policy-managed database. The value of this attribute must be set to the database name.
- ManagedBy: Specifies whether the database is administrator-managed or policy-managed. The default value of this attribute is ADMIN. You need not explicitly set its value. In a policy managed RAC database, this attribute must be set to **POLICY**.

DiskGroupSnap agent attributes:

- FDType: Specifies the configuration to be used for the firedrill. The possible values of this attribute are Bronze and Gold (default).

Attributes introduced in VCS 5.1SP1

Application Agent attributes

- EnvFile: This attribute specifies the environment file that must be sourced before running `StartProgram`, `StopProgram`, `MonitorProgram` or `CleanProgram`.

- UseSUDash: This attribute specifies that the agent must run `su - user -c <program>` or `su user -c <program>` while running `StartProgram`, `StopProgram`, `MonitorProgram` or `CleanProgram`.

RemoteGroup agent attribute

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

DiskGroup agent attribute

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

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

LVMVolumeGroup agent attribute

- EnableLVMTagging: This attributes enables the LVM Tagging if the value of this attribute is set to 1. By default, the value of this attribute is "0", hence LVMTagging is disabled.

NFSRestart agent attribute

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

MultiNICA agent attribute

- Mii: if this attribute is set to 1, the agent uses ethtool and Mii hardware registers to determine the health of the network card.

NotifierSourceIP agent attribute

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

SambaServer agent attributes

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

ASMInst agent attributes

- MonitorOption: Enables or disables health check monitoring.

NetBios agent attribute

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

Sybase agent attribute

- Run_ServerFile: The attribute specifies the location of the RUN_SERVER file for a Sybase instance. If this attribute is not specified, the default location of this file is accessed while starting Sybase server instances.

Cluster-level attributes

- AutoAddSystemToCSG: Indicates whether the newly joined or added systems in the cluster become a part of the SystemList of the ClusterService service group if the service group is confirmed. The value 1 (default) indicates that the new systems are added to SystemList of ClusterService. The value 0 indicates that the new systems are not added to SystemList of ClusterService.

- CounterMissTolerance: If GlobalCounter does not update in CounterMissTolerance intervals of CounterInterval, then VCS reports about this issue depending on the CounterMissAction (that is, CounterMissTolerance * CounterInterval) time has elapsed since last update of GlobalCounter then CounterMissAction is performed. The default value of CounterMisstolerance is 20.
- CounterMissAction: The action mentioned in CounterMissAction is performed whenever the GlobalCounter is not updated for CounterMissTolerance intervals of CounterInterval.

  The two possible values of CounterMissAction are LogOnly and Trigger. LogOnly logs the message in Engine Log and SysLog. Trigger invokes a trigger which has a default action of collecting the comms tar file. The Default value of Trigger is LogOnly.
- PreferredFencingPolicy: The I/O fencing race policy to determine the surviving subcluster in the event of a network partition. Valid values are Disabled, System, or Group.

  Disabled: Preferred fencing is disabled. The fencing driver favors the subcluster with maximum number of nodes during the race for coordination points.

  System: The fencing driver gives preference to the system that is more powerful than others in terms of architecture, number of CPUs, or memory during the race for coordination points. VCS uses the system-level attribute FencingWeight to calculate the node weight.

  Group: The fencing driver gives preference to the node with higher priority service groups during the race for coordination points. VCS uses the group-level attribute Priority to determine the node weight.

Resource type attributes

- IMF: Determines whether the IMF-aware agent must perform intelligent resource monitoring.

  It is an association attribute with three keys Mode, MonitorFreq, and RegisterRetryLimit.
  - Mode: Defines whether to perform IMF monitoring based on the state of the resource. Mode can take values 0, 1, 2, or 3. Default is 0.
  - MonitorFreq: Specifies the frequency at which the agent invokes the monitor agent function. Default is 1.
  - RegisterRetryLimit: Defines the maximum number of times the agent attempts to register a resource. Default is 3.
- IMFRegList: Contains a list of attributes. The values of these attributes are registered with the IMF module for notification. If an attribute defined in IMFRegList attribute is changed then the resource, if already registered, is

unregistered from IMF. If IMFRegList is not defined and if any attribute defined in ArgList is changed the resource is unregistered from IMF.

- AlertOnMonitorTimeouts: Indicates the number of consecutive monitor failures after which VCS sends an SNMP notification to the user.

Attributes introduced in VCS 5.1

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

Resource type attributes:

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

Cluster level attributes:

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

Attributes introduced in VCS 5.0 MP3

VCS 5.0MP3 introduced the following attributes.

Resource type attributes:

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

- AgentDirectory: Complete path of the directory in which the agent binary and scripts are located. Use when the agent binaries are not installed at their default locations.

Cluster level attributes:

- DeleteOnlineResource: Defines whether you can delete online resources.
- HostMonLogLvl: Controls the behavior of the HostMonitor daemon. Configure this attribute when you start the cluster. You cannot modify this attribute in a running cluster.
- EngineShutdown: Provides finer control over the hastop command.
- BackupInterval: Time period in minutes after which VCS backs up configuration files.
- OperatorGroups: List of operating system user account groups that have Operator privileges on the cluster.
- AdministratorGroups: List of operating system user account groups that have administrative privileges on the cluster.
- Guests: List of users that have Guest privileges on the cluster.

System level attributes:

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

Service group level attributes:

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

Changes to VCS bundled agents

This section describes changes to the bundled agents for VCS.

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

Support for Windows DNS server

The DNS agent now supports Windowcs DNS server in its configuration. A new attribute UseGSSAPI is added to DNS agent configuration for this functionality.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on using this attribute and additional requirements for configuring DNS agent with Windows DNS server.

DNS agent supports DNS scavenging for Windows DNS servers

DNS agent can now be configured to send periodic refresh request to configured Windows DNS servers to avoid aging and scavenging of resource records.

RVGPrimary agent starts replication from new primary post role change to the other secondaries in the RDS

After a successful migration or takeover of a Secondary RVG, the RVGPrimary agent ensures to automatically start the replication from the new Primary to the additional Secondary(s) that exists in the RDS, if any.

Refer to the *SF Replication Administrator's Guide* for information about how the RVGPrimary agent works in a multiple secondary setup.

Support for ethtool in NIC agent

The Linux NIC and MultiNICA agents support ethtool based device status monitoring since VCS 5.1 release.

Campus Cluster firedrill made easy

Campus Cluster firedrill configuration has been made easier by introducing a new attribute FDType. You can control the firedrill flavor by setting just this attribute instead of specifying specific values for other attributes.

You need to install the Global Cluster Option (GCO) license to run firedrill in a Campus Cluster configuration.

Change in behavior of DiskGroup agent

The following changes have been affected to the DiskGroup agents:

- The type of the attribute PanicSystemOnDGLoss is changed from boolean to integer. Attribute PanicSystemOnDGLoss now accepts the following values:
  0: Do not halt the system
  1: Halt the system if either disk groups go into disabled state or the disk group resource faults due to monitor timeout.

2: Halt the system if disk group goes into disabled state.
3: Halt the system if disk group resource faults due to monitor timeout.

- The monitor agent function takes the service group containing the DiskGroup resource offline if the MonitorReservation attribute is set to 0, the value of the cluster wide attribute UseFence is set to SCSI3, and if the disk group is found imported without SCSI reservation.

Application agent enhancements

Application agent has undergone the following enhancements:

- Enhanced agent to support shared disk for StartProgram, StopProgram, MonitorProgram, and CleanProgram.
- Enhanced agent to understand Unix style for return codes for MonitorProgram, 0 (ONLINE) and 1 (OFFLINE).

Mount agent enhancements

Mount agent now supports ext4 and xfs file system.

Apache agent enhancements

The following are the enhancements to the Apache agent

- The httpDir attribute is enhanced such that you can specify the full path of the binary (including the binary name). If you specify only the directory name, the agent assumes the default binary name httpd.
- If the Apache Benchmarking binary in the httpDir directory does not use the default name, then the agent recognizes the alternative binary name ab2, and performs detail monitoring.
- Enhanced Apache agent version parsing to accommodate IBM HTTP server 7.0.

First failure data capture for VCS agent entry point times out

To debug entry point timeout, a new attribute AdvDbg is introduced. The attribute helps VCS are some information like process stack, process tree and core on entry point timeout. This is helpful in troubleshooting entry point timeout scenarios.

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

Changes in network agents

The NetMask attribute is a required attribute for the following agents:

- IP
- IPMultiNIC
- MultiNICA

Changes to database agents

Changes to DB2 agent

- The VCS agent for DB2 now supports intelligent resource monitoring for online Db2 processes in PRON mode for non-MPP and MPP configuration mode.
- Prior to VCS 6.0 release, in case of partition mobility from source node to the target node having high speed inter-connect/switch configuration, the switch name entry would not get updated in the db2nodes.cfg configuration file. In VCS 6.0 release, the DB2 agent ensures that the switch name gets updated correctly in the configuration file.
- Added IMF support for DB2 agent: The DB2 agents now leverages the IMF functionality for instantaneous detection of resource state change. This also reduces the CPU footprint of the agent by significantly reducing the periodic monitoring for resource states.
  See the *Administrator's Guide* for more information.

Changes to the Oracle agent

- The VCS agent for Oracle has two new attributes: DBName and ManagedBy.
- The VCS agent for Oracle has additional options to the StartUpOpt and ShutDownOpt attributes:
  - The StartUpOpt attribute introduces SRVCTLSTART_RO as additional startup options.
  - The ShutDownOpt attribute introduces SRVCTLSTOP_TRANSACT, SRVCTLSTOP_ABORT, and SRVCTLSTOP_IMMEDIATE as additional shut down options.
- The VCS agent for Oracle introduces support for policy managed database.
- The VCS agent for Oracle ASM instance introduces the following additional Startup options:
  - STARTUP_MOUNT
  - STARTUP_OPEN
  - SRVCTLSTART_MOUNT

- SRVCTLSTART_OPEN

- The VCS agent for Oracle ASM instance introduces the following additional Shutdown option:
  - SRVCTLSTOP
- With Oracle version 11.2.0.2, the Oracle agent for VCS supports Startup and Shutdown options that use `srvctl` utility for Oracle restart configuration.

Changes to the Sybase agent

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

- The VCS agents for Sybase and SybaseBk now support intelligent resource monitoring.
- Intelligent monitoring framework (IMF) is enabled by default in VCS 6.0 release. The `haimfconfig` script can be used to enable/disable IMF for Sybase and SybaseBk agents. The Sybase agent now leverages the IMF functionality for instantaneous detection of resource state change. This also reduces the CPU footprint of the agent by significantly reducing the periodic monitoring for resource states.
  See the *Administrator's Guide* and *Agent for Sybase Installation and Configuration Guide* for more information.
- In addition to Sybase ASE Enterprise edition support, Sybase agent is enhanced to support Sybase ASE cluster edition. For Sybase ASE cluster edition, the VCS agent for Sybase makes the Sybase adaptive server highly available in a VCS cluster.
- The Sybase agent introduces the following new attributes:
  - Quorum_dev
  - interfaces_File
  - ShutdownWaitLimit (default value 60)
  - DelayAfterOnline (default value 10)
  - DelayAfterOffline (default value 2)
- The SybaseBk agent introduces the following new attribute:
  - interfaces_File
- The default value of ToleranceLimit attribute is set to 1 (one) for Sybase agent.

- The DetailMonitor attribute is deprecated in VCS 6.0. Instead, LevelTwoMonitorFreq attribute of Sybase agent may be used. The default value of LevelTwoMonitorFreq attribute is 0 (zero).
- The long pathname limitation for $SYBASE is resolved.
- With VCS 6.0 release using VCS Cluster Manager (Java Console), Sybase and SybaseBk agents encrypt the password by default. Sybase and SybaseBk agents supports both plain text and encrypted password. If required, the plain text value can be specified for agent attributes using the command line or by editing the configuration file.
- Sybase agent uses new timeout option during shutdown of Sybase dataserver used instead of "shutdown with nowait". For Sybase ASE Enterprise edition the timeout option for shutdown command is supported for versions 12.5.4 and 15.0.2 onwards.
  For Sybase ASE Cluster Edition the timeout option for shutdown command is supported from versions 15.5 ESD #1 onwards.

Changes to VCS clusters in secure mode

In this release, the installation and configuration experience of secure cluster is considerably simplified. You can easily convert the cluster into secure cluster with this simplified secure cluster configuration model.

The new architecture is based on embedded VxAT, where the security components are installed as a part of the VCS package. The root broker is no longer a single-point-of-failure in the new architecture. There is no dependency on a separate VRTSat package. Non-root users who are already logged on VCS hosts are now not prompted for password. Additionally, a cluster-level user feature is introduced to simplify user administration in secure clusters.

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

Changes to LLT

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

- LLT now supports VLAN tagging (IEEE 802.1Q).
- The `lltconfig` command includes the following new options:
  - -N
    You can use this option to list all the used cluster IDs.
  - -M
    You can use this option to display the currently loaded LLT module version information.

See the `lltconfig` manual page for more information.

See the `lltstab` manual page for more information.

- Link utilization statistics are enhanced that help in the root cause analysis of performance related issues.
- Periodic flushing of ARP cache is disabled.
- When MAC address of a NIC changes, LLT immediately relearns the new MAC address and also updates the peer nodes about the change.

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

Changes to GAB

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

Better GAB and I/O fencing integration to ensure application availability

In the event of a split-brain situation before VxFEN module implements the decision, sometimes GAB proceeds with attempting to resolve the join after the split-brain. GAB removes all but one joining subcluster. This behavior can cause the entire cluster to shut down. To avoid this scenario, GAB now gives priority to the fencing module.

With the GAB and I/O fencing integration in this release, if the I/O fencing module's decision is still pending before GAB initiates a join of the subcluster, GAB delays the iofence message. GAB wait depends on the value of the VxFEN tunable parameter *panic_timeout_offst* based on which VxFEN computes the delay value and passes to GAB.

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

GAB can now recognize clients with names in addition to ports

When kernel clients initialize GAB API, they can now define a client name string. GAB now adds a client name which enables GAB to track the client even before GAB port is registered. GAB also passes the client name information to LLT when registering the LLT port. The `lltstat -p` command also displays the GAB client names when providing the status details of the ports in use.

This feature is applicable only to GAB kernel clients, and not applicable for user-land GAB clients such as HAD.

The gabconfig command has new -C option

The -C option of the `gabconfig` command lists the names of the GAB clients that have registered with GAB. The `-C` option when used with `-a` option lists the client names along with the port membership details.

Changes to I/O fencing

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

Installer support to migrate between fencing configurations in an online cluster

You can now use the installer to migrate between disk-based and server-based fencing configurations. You can also replace the coordination points for any I/O fencing configuration in an online cluster using the same installer option. The installer uses the `vxfenswap` script internally.

You can also use response files to perform these I/O fencing reconfiguration operations.

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

Support for racer node re-election during I/O fencing race

At the time of a network partition, the VxFEN module elects the lowest node in each sub-cluster as the racer node to race for the coordination points on behalf of the sub-cluster. The other spectator nodes wait on the racer node to do the fencing.

In the previous releases, the I/O fencing race was entirely dependent on the single racer node as follows:

- If the racer node is not able to reach a majority of coordination points, then the VxFEN module on the racer node sends a LOST_RACE message and all nodes in the subcluster also panic when they receive the LOST_RACE message.
- If the racer node panics during the arbitration, then the spectator nodes in the sub-cluster assume that the racer node lost the race and the spectator nodes also panic.

With the new racer node re-election feature, the VxFEN module re-elects the node with the next lowest node id in the sub-cluster as the racer node. This feature optimizes the chances for the sub-cluster to continue with the race for coordination points.

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

Support for multiple virtual IP addresses in CP servers

You can now configure multiple network paths (virtual IP addresses) to access a CP server. CP server listens on multiple virtual IP addresses. If a network path fails, CP server does not require a restart and continues to listen on one of the other available virtual IP addresses.

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

Support for Quorum agent in CP servers

With the support for multiple virtual IP addresses, you can now use the Quorum agent to configure CP server service group failover policies. You can specify the minimum number of IP resources that must be online for the Quorum resource to remain online.

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

With fencing enabled, GAB can now automatically seed the cluster when some cluster nodes are unavailable

In the earlier releases, if some of the nodes are not up and running in a cluster, then GAB port does not come up to avoid any risks of preexisting split-brain. In such cases, you can manually seed GAB using the command `gabconfig -x` to bring the GAB port up. However, if you have enabled I/O fencing in the cluster, then I/O fencing can handle any preexisting split-brain in the cluster.

In this release, I/O fencing has extended this functionality to be able to automatically seed GAB as follows:

- If a number of nodes in a cluster are not up, GAB port (port a) still comes up in all the member-nodes in the cluster.
- If the coordination points do not have keys from any non-member nodes, I/O fencing (GAB port b) also comes up.

This new functionality is disabled by default. You must manually enable this automatic seeding feature of GAB in clusters where I/O fencing is configured in enabled mode.

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

You can still use the `gabconfig -x` command to manually seed the cluster.

Graceful shutdown of a node no longer triggers I/O fencing race condition on peer nodes

In the earlier releases, a gracefully leaving node clears its I/O fencing keys from coordination points. But the remaining sub-cluster races against the gracefully leaving node to remove its registrations from the data disks. During this operation, if the sub-cluster loses access to the coordination points, the entire cluster may panic if the racer loses the race for coordination points.

In this release, this behavior has changed. When a node leaves gracefully, the CVM or other clients on that node are stopped before the VxFEN module is unconfigured. Hence, data disks are already clear of its keys. The remaining sub-cluster tries to clear the gracefully leaving node's keys from the coordination points but does not panic if it is not able to clear the keys.

Changes related to virtualization support

This section lists virtualization changes for this release.

New KVMGuest agent on Linux

The KVMGuest agent monitors the Linux Kernel-based Virtual Machines (KVM guest) and brings the KVM guest online and offline. KVMGuest agent uses virsh commands.

You can use this agent to make the KVM guests highly available and to monitor them. This agent is added as a part of virtualization support.

Virtualization support

VCS can be installed and run inside a virtual machine or guest created using Red hat KVM (kernel-based virtual machine). The following clustering configurations are supported:

- VCS cluster across VM Guests (VM-VM) on the same or different physical hosts - for application availability
- VCS cluster across physical machines (PM-PM) without resource monitoring inside VM Guests - for virtual machine availability
- VCS cluster across physical machines (PM-PM) with resource monitoring inside VM Guests - for both virtual machine and application availability
- VCS cluster across physical machine and VM Guest.(additional configuration to KVM and Veritas Cluster Server clustering configurations)

The following table shows the system requirements for the KVM-supported configurations.

Table 1-1 System requirements for the KVM-supported configurations

Supported OS version in host	RHEL 6 Update 1
Supported OS in VM Guest	RHEL 6 Update 1
Hardware requirement	Full virtualization-enabled CPU

Licensing changes in the SFHA Solutions 6.0 release

Storage Foundation and High Availability Solutions 6.0 introduces the following licensing changes:

- The Cluster File System license is deprecated. CFS customers are entitled to the Storage Foundation Cluster File System High Availability (SFCFS HA) functionality.
- The VVR Option is renamed as Veritas Replicator Option. This option includes VVR (volume-based replication) and the new file-based replication solution.
- The VVR Enterprise license is deprecated; you can use Storage Foundation Enterprise and add Veritas Replicator Option to get this functionality. VVR Enterprise customers are entitled to Storage Foundation Enterprise with Replicator Option.
- The VCS license enables full cluster functionality as well as the limited start/stop functionality.
- Storage Foundation Enterprise CFS for Oracle RAC (Linux/x64) customers are entitled to Storage Foundation Enterprise for Oracle RAC (Linux/x64.)

The following functionality is included in the Standard and Enterprise licenses:

- The Compression feature is available with the Standard license.
- The SmartTier feature is now available with the Standard license.
- The Deduplication feature is available with the Enterprise license.

The following products are included in this release:

- Dynamic Multi-Pathing
- VirtualStore
- Storage Foundation Basic
- Storage Foundation Standard
- Storage Foundation Enterprise
- Veritas Cluster Server

- Veritas Cluster Server HA/DR
- Storage Foundation Standard HA: Storage Foundation Standard plus Veritas Cluster Server
- Storage Foundation Enterprise HA: Storage Foundation Enterprise plus Veritas Cluster Server
- Storage Foundation Enterprise HA/DR
- Storage Foundation Enterprise Cluster File System HA
- Storage Foundation Enterprise Cluster File System HA/DR
- Storage Foundation Enterprise for Oracle RAC
- Storage Foundation Enterprise HA/DR for Oracle RAC
- Storage Foundation Enterprise for Sybase ASE CE
- Storage Foundation Enterprise HA/DR for Sybase CE

HA: High Availability

HA/DR: High Availability and Disaster Recovery

Veritas Replicator Option can be added to all Storage Foundation and High Availability products, except Dynamic Multi-Pathing and Veritas Cluster Server.

Note that products, features, and options may differ by operating system and platform. Please see the product documentation for information on supported platforms.

Enhancements to collecting a VxExplorer troubleshooting archive

The Symantec Operations Readiness Tools (SORT) data collector contains functionality to collect and submit a VxExplorer archive. You can send this archive to Symantec Technical Support for problem diagnosis and troubleshooting. VxExplorer does not collect customer data.

The legacy `VxExplorer` script now works differently. When you run the script, it launches the SORT data collector on the specified local host with the `-vxexplorer` option.

To learn more about using the data collector to collect a VxExplorer archive, see:

www.symantec.com/docs/HOWTO32575

Changes related to product documentation

The Storage Foundation and High Availability Solutions 6.0 release includes the following changes to the product documentation.

Table 1-2 lists the documents introduced in this release.

Table 1-2 New documents

New documents	Notes
Veritas Storage Foundation Installation Guide	Installation and upgrade information for Storage Veritas Foundation.
Veritas Storage Foundation Administrator's Guide	Administration information for Veritas Storage Foundation.
Veritas Storage Foundation and High Availability Release Notes	Release-specific information for Veritas Storage Foundation and High Availability users.
Veritas Storage Foundation and High Availability Solutions Solutions Guide	Solutions and use cases for Veritas Storage Foundation and High Availability Solutions.
Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide	Troubleshooting information for Veritas Storage Foundation and High Availability Solutions.
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	Virtualization-related information for Veritas Storage Foundation and High Availability Solutions.
Symantec VirtualStore Release Notes	Release-specific information Symantec VirtualStore.
Veritas Storage Foundation for Sybase ASE CE Release Notes	Release-specific information for Veritas Storage Foundation for Sybase ASE CE.
Veritas Storage Foundation for Sybase ASE CE Installation Guide	Installation information for Veritas Storage Foundation for Sybase ASE CE.
Veritas Storage Foundation for Sybase ASE CE Administrator's Guide	Administration information for Veritas Storage Foundation for Sybase ASE CE.
Virtual Business Services–Availability User's Guide	Information about Virtual Business Services. This document is available online.

Table 1-3 lists the documents that are deprecated in this release.

Table 1-3 Deprecated documents

Deprecated documents	Notes
Veritas File System Administrator's Guide	Content now appears in the *Veritas Storage Foundation Administrator's Guide* and in the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide.*
Veritas Volume Manager Administrator's Guide	Content now appears in the *Veritas Storage Foundation Administrator's Guide* and in the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide.*
Veritas Storage Foundation Advanced Features Administrator's Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Solutions Guide.*
Veritas Volume Manager Troubleshooting Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide.*
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	Content now appears in the *Veritas Cluster Server Bundled Agents Reference Guide.*
Veritas Volume Replicator Planning and Tuning Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide.*
Veritas Volume Replicator Advisor User's Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide.*

Table 1-4 lists documents that are no longer bundled with the binaries. These documents are now available online.

Table 1-4 Online documents

Document
Veritas Cluster Server Agent Developer's Guide
Veritas File System Programmer's Reference Guide

Changes introduced in VCS 5.1SP1PR2

This section introduces the changes introduced in VCS 5.1SP1PR2.

Added support for Red Hat Enterprise Linux 6

Added support for Red Hat Enterprise Linux (RHEL6) in this release.

Changes related to NFSv4 exports

Prior to RHEL6 for NFSv4 exports, the OS did not assign the root of the pseudo file system exported to NFS clients. Hence, it was required to put fsid=0 option in one of the Share resources to make the Share path as a root. For RHEL6, this is not mandatory. By default, '/' is the root of the pseudo file system exported to NFS clients.

Changes introduced in VCS 5.1SP1PR3

This section introduces the changes introduced in VCS 5.1SP1PR3.

Simplifying Install and Configuration

The VirtualStore installer has been simplified to allow a typical mode of install. Also it is no longer required to install the VMware PERL SDK on all VirtualStore nodes.

VMware View Integration

The Cloning Wizard has the option to automatically import the virtual machine clones into a VMware View pool.

Ability to power on virtual machines after cloning

The Cloning wizard has the option to power on the virtual machine clones after they have been created.

Support for Multiple VirtualStore clusters

Multiple VirtualStore Plug-ins, one for each VirtualStore cluster, are not supported with a single vCenter Server.

Multiple VirtualStore clusters registered with a single vCenter Server is supported.

VCS system requirements

This section describes system requirements for VCS.

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

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

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

See "Hardware compatibility list" on page 37.

See "Supported Linux operating systems" on page 37.

Hardware compatibility list

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

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

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

Supported Linux operating systems

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

Table 1-5 shows the supported Linux operating systems for this release.

Table 1-5 Supported Linux operating systems

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 6	Update 1, 2	2.6.32-131.0.15.el6 2.6.32-220.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only

Table 1-5 Supported Linux operating systems *(continued)*

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 5	Update 5, 6, 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 11	SP1	2.6.32.12-0.7	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 10	SP4	2.6.16.60-0.85.1	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Linux 6	**6.1	2.6.32-131.0.15.el6	64-bit x86, EMT*/Opteron
Oracle Linux 5	**Update 5, 6, 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron

* Extended Memory Technology

** RHEL-compatible mode only.

Note: Only 64-bit operating systems are supported.

If your system is running an older version of either Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Oracle Linux, upgrade it before attempting to install the Veritas software. Consult the Red Hat, SUSE, or Oracle documentation for more information on upgrading or reinstalling your operating system.

For DMP, SF, SFHA, SFCFSHA, SFRAC, VCS, and VirtualStore, Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

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

Required Linux RPMs for VCS

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

Table 1-6 lists the RPMs that VCS requires for a given Linux operating system.

Table 1-6 Required RPMs

Operating system	Required RPMs
RHEL 5	compat-libstdc++-33-3.2.3-61.x86_64.rpm glibc-2.5-49.i686.rpm glibc-2.5-49.x86_64.rpm ksh-20100202-1.el5.x86_64.rpm libgcc-4.1.2-48.el5.x86_64.rpm libgcc-4.1.2-48.el5.i386.rpm libstdc++-4.1.2-48.el5.i386.rpm pam-0.99.6.2-6.el5_4.1.x86_64.rpm
RHEL 6	compat-libstdc++-33-3.2.3-69.el6.x86_64.rpm compat-libstdc++-296-2.96-144.el6.i686.rpm glibc-2.12-1.7.el6.x86_64.rpm glibc-2.12-1.7.el6.i686.rpm ksh-20100621-2.el6.x86_64.rpm libgcc-4.4.4-13.el6.i686.rpm libgcc-4.4.4-13.el6.x86_64.rpm libstdc++-4.4.4-13.el6.x86_64.rpm pam-1.1.1-4.el6.x86_64.rpm
SLES 10	glibc-2.4-31.81.11.x86_64.rpm glibc-32bit-2.4-31.81.11.x86_64.rpm ksh-93t-13.17.19.x86_64.rpm libgcc-4.1.2_20070115-0.32.53.x86_64.rpm libstdc++-4.1.2_20070115-0.32.53.x86_64.rpm pam-0.99.6.3-28.23.15.x86_64.rpm

Table 1-6 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 11	glibc-2.11.1-0.17.4.x86_64.rpm glibc-32bit-2.11.1-0.17.4.x86_64.rpm ksh-93t-9.9.8.x86_64.rpm libgcc43-32bit-4.3.4_20091019-0.7.35.x86_64.rpm libgcc43-4.3.4_20091019-0.7.35.x86_64.rpm libstdc++33-3.3.3-11.9.x86_64.rpm libstdc++43-32bit-4.3.4_20091019-0.7.35.x86_64.rpm

Supported software for VCS

VCS supports the following volume managers and file systems:

- ext2, ext3, reiserfs, NFS, and bind on LVM2, raw disks, and VxVM.
- ext4 and xfs on LVM2 and raw disks
- Veritas Storage Foundation (SF): Veritas Volume Manager (VxVM) with Veritas File System (VxFS)

  VCS 6.0 supports the following versions of SF:

  - SF 6.0
    - VxVM 6.0 with VxFS 6.0
  - SF5.1SP1
    - VxVM 5.1SP1 with VxFS 5.1SP1

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

Supported VCS agents

Table 1-7 lists the agents for enterprise applications and the software that the agents support.

Table 1-7 Supported software for the VCS agents for enterprise applications

Agent	Application	Application version	Linux version
DB2	DB2 Enterprise Server Edition	9.1, 9.5, 9.7	RHEL5, OLE5, SLES10
		9.5, 9.7	SLES11
		9.7	RHEL6, OLE6
Oracle	Oracle	10gR2, 11gR1, 11gR2	RHEL 5, SLES10, SLES 11, OEL5
Sybase	Sybase Adaptive Server Enterprise	12.5.x, 15.x	RHEL5, RHEL6 SLES10, SLES11, OEL5, OEL6

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

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

No longer supported

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

- Several documents are deprecated in this release.
  See "Changes related to product documentation" on page 33.

No longer supported agents and components

VCS no longer supports the following:

- Configuration wizards
- CampusCluster agent
- SANVolume agent
- VRTSWebApp
- Oracle 8.0.x, Oracle 8.1.x, and Oracle 9i - not supported by the Oracle agent.
- VCS documentation package (VRTSvcsdc)

The VCS documentation package (VRTSvcsdc) is deprecated. The software disc contains the documentation for VCS in Portable Document Format (PDF) in the *cluster_server/docs* directory.
Symantec recommends copying pertinent documents from the disc to your system directory /opt/VRTS/docs for reference.

- The *Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide* is deprecated and its content is accommodated in the *Veritas Storage Foundation and HIgh Availability Solutions Replication Administrator's Guide* and *Veritas Cluster Server Bundled Agents Reference Guide*.
- hahbsetup tool. This tool is removed as no supported feature requires this tool.
- VRTScutil RPM. This RPM is no longer supported.

Deprecated attributes

Deprecated Oracle agent attributes:

- AgentDebug
- DetailMonitor

Deprecated Mount agent attributes:

- SecondLevelMonitor
- SecondLevelTimeout

Deprecated Host Monitor attribute:

- CPUUsageMonitoring: The attribute can no longer be used to disable CPU usage monitoring by Host Monitor agent.

Sybase agent attribute:

- DetailMonitor

Known issues

This section covers the known issues in this release.

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

See "Documentation" on page 79.

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

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

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.

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 `libvird` process manually.

Issues related to installation

This section describes the known issues during installation and upgrade.

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups (2591399)

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 (2115662)

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

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

For example: `rpm -Uvh --nopreun VRTSvlic-3.02.60.007-0.x86_64.rpm`

Manual upgrade of VRTSvlic RPM loses keyless product levels [2115662]

If you upgrade the `VRTSvlic` RPM manually, the product levels that were set using `vxkeyless` may be lost. The output of the `vxkeyless display` command will not display correctly. To prevent this, perform the following steps while manually upgrading the VRTSvlic RPM.

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

   ```
   # vxkeyless display
   ```

2. Set the product level to NONE.

   ```
   # vxkeyless set NONE
   ```

3. Upgrade the VRTSvlic RPM.

   ```
   # rpm -Uvh --nopreun VRTSvlic-3.02.60.007-0.x86_64.rpm
   ```

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

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

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

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 1-8 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 .

Table 1-8 Whether attributes are configured and required actions that you need to perform during upgrade *(continued)*

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
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"

Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

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

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

To resolve this issue, perform the following steps:

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

```
# vxkeyless set NONE
```

3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic`:
   - Verify if the key has `VXKEYLESS` feature Enabled using the following command:
     ```
     # vxlicrep -k <license_key> | grep VXKEYLESS
     ```
   - Delete the key if and only if `VXKEYLESS` feature is Enabled.

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

4. Restore the previous list of products with the command:
   ```
   # vxkeyless set product1[|,product]
   ```

SELinux error during installation of VRTSvcsag RPM

During the installation of VRTSvcsag RPM on RHEL 5 SELinux enabled machine, you may observe following SELinux error:

```
/usr/sbin/semodule: Failed on /opt/VRTSvcs/bin/selinux/vcsag.pp!
```

This error occurs due to improper installation of the SELinux package. As a result, SELinux commands may not function properly.

Workaround: Reinstall the SELinux package and relabel filesystem by either init or fixfiles method.

Secure WAC communication needs to be disabled explicitly [2392568]

If you have WACs communicating securely where VCS is configured in secure mode and if you disable the VCS security, the WAC where VCS security is disabled continues attempting to communicate securely without success. Therefore, you need to explicitly disable WAC security when you disable VCS security.

Workaround: No workaround. Secure WAC communication needs to be disabled explicitly.

Web installer has no option to remove node from a cluster

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

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

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

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

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

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

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

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

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

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

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

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

sfmh discovery issue when you upgrade your Veritas product to 6.0 (2622987)

If a host is not reporting to any management server but sfmh discovery is running before you upgrade to 6.0, sfmh-discovery may fail to start after the upgrade.

Workaround:

If the host is not reporting to VOM, stop sfmh-discovery manually before upgrading to 6.0 by executing the following command on the managed host:

```
/opt/VRTSsfmh/adm/vxvmdiscovery-ctrl.sh stop
```

Incorrect server names sometimes display if there is a clock synchronization issue (2627076)

When you install a cluster with the Web-based installer, you choose to synchronize your systems with an NTP server due to a clock synchronization issue, you may see the NTP server name in messages instead of your server names.

Workaround:

Ignore the messages. The product is still installed on the correct servers.

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.

Operational issues for VCS

LVMLogicalVolume online entry point stops responding and times out for mirrored volumes on SLES10 [2077294]

`LVMLogicalVolume` uses `lvchange` command to activate the logical volumes. In case of mirrored volumes, the `lvchange` command itself stops responding when it is invoked through a script. This causes online entry point to time out and the online entry point of `LVMLogicalVolume` resource stops responding. This is an issue with the SLES10.

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

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

Workaround: You need to restore at least one path.

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

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

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

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

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

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

Workaround:

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

Issues related to the VCS engine

Extremely high CPU utilization may cause HAD to fail to heartbeat to GAB

When CPU utilization is very close to 100%, HAD may fail to heartbeat to GAB. [1818687]

Agent framework can reject `hares -action` command

When a probed resource is disabled and later enabled then, the agent framework can reject `hares -action` command till the agent successfully monitors the resource.

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 [2397532]

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 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 [2490404]

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: Forcefully stop the agent process.

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

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

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

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

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

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

If secure and non-secure WAC are connected the engine_A.log receives logs every 5 seconds [1539646]

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

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

Oracle group fails to come online if Fire Drill group is online on secondary cluster [2556835]

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.

POSTONLINE and POSTOFFLINE triggers are not enabled by default [2567387]

Before VCS 6.0, POSTONLINE and POSTOFFLINE triggers were enabled by default, so the triggers got executed whenever a service group came online. In VCS 6.0,

you must explicitly enable the POSTONLINE and POSTOFFLINE triggers whenever you upgrade to VCS 6.0.

Alternatively, if you want the triggers to execute after the upgrade:

1. Before upgrade, set vcs_start = 0 in /etc/default/vcs

   so that HAD does not start after the upgrade.
2. Upgrade the existing VCS to VCS 6.0.
3. Set vcs_start = 1 in /etc/default/vcs
4. Start VCS on each node using `hastart`.
5. Set TriggersEnabled in main.cf for required groups as ollows:

   ```
   TriggersEnabled @<systemname>={POSTONLINE, POSTOFFLINE}
   ```

Example of trigger behavior:

```
group scriptfileonoff (
        SystemList = { vcssx235 = 0, vcssx236 = 1 }
        AutoStartList = { vcssx235, vcssx236 }
        TriggersEnabled @vcssx235 = { POSTONLINE }
        )
        MyFileOnOff MFileOnOff (
                PathName = "/tmp/mf1"
                )
        MyFileOnOff MFileOnOff1 (
                PathName = "/tmp/mf2"
```

Two CmdServer instances seen running on a node [2399292]

You may see two instances of CmdServer running on a node. One of these using IPv4 and the other IPv6.

This does not impact functionality in any way.

Workaround: No workaround.

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.

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.

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

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

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

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

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

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

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

In case the operating system is not installed inside the guest or the installed operating system is corrupted, `virsh` utility still reports the guest state as running.

Thus, VCS also reports the resource state as ONLINE. Since RedHat KVM does not provide the state of the operating system inside guest, VCS cannot detect the guest state based on the state of the operating system.

Workaround: No workaround for this known issue.

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

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

This issue may be observed:

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

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

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

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

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

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

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

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

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

Application agent cannot handle a case with user as root, envfile set and shell as csh [2584285]

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

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

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

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

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

Workaround: No workaround.

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

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

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

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

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

A service group contains DiskGroup, Volume and Mount resources and this service group is online. Volume is mounted by Mount resource with VxFSMountLock enabled. An attempt to manually unmount the volume using `umount -l` system command causes the mount point to go away; however, the file system lock remains as it is. The volume cannot be stopped as it is mount locked and hence the disk group cannot be imported. This causes the disk group resource to go into UNABLE to OFFLINE state. Also, any attempt to again mount the file system fails, because it is already mount locked. This issue is due to file system behavior on Linux.

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

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

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

Workaround:

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

Concurrency violation in the service group [2555306]

Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and PanicSystemOnDGLoss is set to 0

This happens when:

- In a cluster environment/configuration, if cluster wide UseFence attribute is set to SCSI3 and service group contains Volume resource and DiskGroup resource with the PanicSystemOnDGLoss attribute set to 0 (zero).

- If storage connectivity is lost or all paths under VxDMP are disabled, VCS fails over the service group. If storage connectivity is restored on the node on which the service group was faulted and DG is not deported manually, then volume may get started if disk group is not deported during the service group failover. So volume resource shows state as online on both the nodes and thus cause concurrency violation. This may lead to data corruption.

Workaround: Ensure that the disk group is deported soon after storage connectivity is restored.

You are recommended to always configure Volume resource whenever Disk group resources is configured and set the attribute PanicSystemOnDGLoss to 1 or 2 as per requirement.

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 [2555191]

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

Workaround: Clear the fault and 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.

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

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

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

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

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. [2615089]

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

Issues related to the VCS database agents

Health check monitoring does not work with VCS agent for Oracle [2101570, 1985055]

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

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

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

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

Make sure that the ohasd has an entry in the init scripts [1985093]

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

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

No health check monitoring for Oracle agent on SLES11 platform [1938167]

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

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

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

Workaround:

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

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

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

NOFAILOVER action specified for certain Oracle errors

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

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

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

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

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

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

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

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

Concurrency violation due to process startup on failover node is not detected when detail monitoring is set for Oracle resources [2917558]

Inside a failover service group, when the administrator starts an Oracle resource on a node and if the Oracle instance is online on any other node within the cluster, the instance would come up. However, the database does not get mounted. In such circumstances, this startup attempt is detected by basic monitoring. If detail monitoring is enabled, this startup attempt does not get detected.

Workaround: No workaround.

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

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 [1511211]

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.

Issues related to global clusters

The engine log file receives too many log messages on the secure site in global cluster environments [1539646]

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.

Issues related to LLT

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 `llttab` file while configuring LLT if you want to configure a vlan later. If you have already specified the MAC address of a NIC, then delete the MAC address from the `llttab` file, and update the file before you restart LLT.

LLT port stats sometimes shows recvcnt larger than recvbytes (1788315)

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 (1809827)

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.

Issues related to GAB

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.

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.

Issues related to I/O fencing

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 *Veritas 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 nodes' and users' information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster nodes' and users' 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@galaxy,
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, 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 in secure mode (2478502)

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 steps on all the nodes of the CP server:

- Rename cpsadm to cpsadmbin.

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

- 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 "$@"
```

- Provide the following permissions to the new file:

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

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

When you configure fencing in customized mode and do 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 atleast 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: No workaround.

NIC resource gets created with incorrect name while configuring CPSSG with the configure_cps.pl script (2585229)

The name of the NIC resource created by the `configure_cps.pl` script does not come out correct when, for example, m^{th} VIP is mapped to n^{th} NIC and every m is not equal to n. In this case, although CPSSG continues to function without any problem, when you unconfigure CPSSG using `configure_cps.pl`, it fails.

Workaround: To unconfigure CPSSG, you must remove the CPSSG configuration from the VCS configuration.

Veritas Cluster Server agents for Veritas Volume Replicator known issues in 6.0

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in 6.0 release.

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

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

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.

Pearl errors seen while using haimfconfig command

Pearl errors seen while using `haimfconfig` command:

```
Pearl errors seen while using haimfconfig command
```

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

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

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

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

It should be replaced as follows in main.cf:

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

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.

VCS Cluster Manager (Java Console) does not encrypt Sybase and SybaseBk agent passwords [2379510]

If `isvcsagencrypt` flag is set to True in Sybase.xml and SybaseBk.xml files, the attribute values get encrypted. However, the password attributes of Sybase and SybaseBk agents do not have the `isvcsagentcrypt` flag set to True in Sybase.xml and SybaseBk.xml files.

Workaround: Sybase and SybaseBk agents are modified to encrypt the password by default. As a result, you need not encrypt passwords if you use the VCS Cluster Manager (Java Console) to configure attributes.

Issues related to Virtual Business Services (VBS)

Fault propagation for Virtual Business Services with shared service groups and different controllers [2407832]

Fault propagation may not work for certain configurations having shared service groups and distinct controllers.

Workaround: No workaround.

Virtual Business Services fail to start if a participating service group has multiple children with the LOCAL FIRM dependency type [2490098]

Virtual Business Services fail to start if a participating service group has multiple children with the LOCAL FIRM dependency type. This occurs because Vertias Cluster Server (VCS) does not support propagating dependencies.

Workaround: Pull the dependent VCS groups into the Virtual Business Services without any dependencies. The Virtual Business Services will recognize the VCS dependencies and treat them as soft Virtual Business Services dependencies.

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 79.

Limitations related to installing and upgrading VCS

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

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

Limitations related to VCS engine

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

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

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

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 SLES11SP1
  - ext4 filesystem configured on VxVM
  - xfs filesystem configured on VxVM

Share agent limitations

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

Driver requirements for DiskReservation agent

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

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

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

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

Limitations related to IMF

- IMF registration on Linux for “bind” file system type is not supported.
- In case of SLES11 SP1 and RHEL6.1:
  - IMF should not be enabled for the resources where the BlockDevice can get mounted on multiple MountPoints.
  - If FSType attribute value is nfs, then IMF registration for “nfs” file system type is not supported.

Limitations related to the VCS database agents

DB2 RestartLimit value

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

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

Limitation with intentional offline functionality of VCS agent for Oracle

The Oracle resource never faults after an intentional offline.

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

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.

Security-Enhanced Linux is not supported on SLES distributions

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

Systems in a cluster must have same system locale setting

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

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

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

The DiskGroupSnap agent has the following limitations:

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

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

System reboot after panic

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

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

Cluster Manager (Java console) limitations

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

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

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

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

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

VCS Cluster Manager fails to connect to the VCS engine if the /etc/hosts file contains IPv6 entries.

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

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the UseFence attribute is set to the default, “None”.

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 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 *Veritas Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

In a cluster using SCSI-3 based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on both the data disks and coordinator disks. In a cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordination points and data disks to prevent possible difficulties with subsequent cluster startup.

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

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

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.

Documentation errata

The following sections cover additions or corrections for Document version: 6.0.4 of the product documentation. These additions or corrections may be included in

later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).

See the corresponding Release Notes for documentation errata related to that component or product.

See "Documentation" on page 79.

See "About Symantec Operations Readiness Tools" on page 9.

Veritas Cluster Server Bundled Agents Reference Guide

Topic: MultiNICA agent > Attributes

The Device attribute description also needs to mention the following:

Device attribute must be localized per system and must have different base IP addresses (as explained in *Sample Configuration: IPMultiNIC and MultiNICA* of *IPMultiNIC agent*).

Veritas Cluster Server Installation Guide

The note in the *Supported software for VCS* section must read as VCS supports the previous and the next versions of SF to facilitate product upgrades.

Documentation

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

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

http://sort.symantec.com/documents

Documentation set

Table 1-9 lists the documents for Veritas Cluster Server.

Table 1-9 Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_60_lin.pdf

Table 1-9 Veritas Cluster Server documentation *(continued)*

Title	File name
Veritas Cluster Server Release Notes	vcs_notes_60_lin.pdf
Veritas Cluster Server Administrator's Guide	vcs_admin_60_lin.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_60_lin.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev_60_unix.pdf
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_60_lin.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_60_lin.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_60_lin.pdf

Table 1-10 lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-10 Veritas Storage Foundation and High Availability Solutions products documentation

Document title	File name
Veritas Storage Foundation and High Availability Solutions Solutions Guide	sfha_solutions_60_lin.pdf
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	sfha_virtualization_60_lin.pdf

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

http://sort.symantec.com/documents

Manual pages

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

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

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

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

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

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

See the `man`(1) manual page.

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

To edit the man(1) configuration file

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

```
export LC_ALL=C
```

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

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

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

where other man paths are specified in the configuration file.

3 Add new section numbers. Change the line:

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

to

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

New features related to Virtual Business Services (VBS)

Virtual Business Services (VBS) feature extends the VOM Business Entity (BE) functionality available since VOM 3.x release. VOM BE support Application Entity type which is hence forth referred to as Virtual Business Services (VBS) in VOM 4.1 and SFHA 6.1 releases. Virtual Business Services allow users to define and manage heterogeneous, inter-cluster, multi-tier applications. Each tier is

represented by a Service Group which may be configured on separate VCS Clusters or ApplicationHA nodes. VBS builds on top of VCS HA/DR and ApplicationHA to provide business service availability across physical and virtual environments. The application tiers (Service Group) can optionally be linked with configurable dependency types and fault actions.

Ordered Start/Stop operations on a VBS

VBS allows ordered start/stop of the entire business service via a single click or through a single command-line interface. If applications are hosted on VMWare virtual machines, you can configure the virtual machines to be automatically started or stopped when you start or stop the Virtual Business Service.

Ability to perform VBS operations via CLI

You can work on Virtual Business Services using CLI from any node of any participating tier clusters. Thus VOM CS is optional after you have configured VBS.

DR support for VBS

VBS provides a comprehensive DR solution that builds on top of VCS DR. The DR support is based on having one or more GCO Service Groups in the VBS.

VBS support for robust fault management

VBS provides robust fault management with configurable actions like stopping, starting, restarting a Service Group (application tier) when the child Service Group faults or recovers. Three dependency types (viz. Soft, Firm, Restart) are supported between Service Groups configured within a VBS.

Secure access control for Virtual Business Services

Operations on a VBS via CLI are permitted only for root users in the participating tiers. Operations through VOM support role based access control (RBAC).

Audit trail of operations performed on a VBS

All user actions performed on a VBS are easily traceable via audit trail and logs.