Veritas™ Cluster Server Release Notes

Linux

5.1 Service Pack 1 Platform Release 2

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: 5.1 SP1 PR2

Document version: 5.1SP1PR2.0

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 Web site at the following URL:

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

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

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

docs@symantec.com

About Symantec Connect

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

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

Release Notes

This document includes the following topics:

About this document

This document provides important information about Veritas Cluster Server (VCS) version 5.1 SP1 PR2 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: 5.1SP1PR2.0 of the *Veritas Cluster Server Release Notes*. Before you start, ensure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

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

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 in PDF formats:

`/product_name/docs`

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

This release includes the following component product release notes:

- *Veritas Storage Foundation Release Notes* (5.1 SP1 PR2)

About Veritas Cluster Server

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

About VCS agents

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

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

The Veritas High Availability Agent Pack gives you access to agents that provide high availability for various applications, databases, and third-party storage solutions. The Agent Pack is available through Symantec™ Operations Readiness Tools (SORT). For more information about SORT, See "About Symantec Operations Readiness Tools" on page 9.. For information about agents under development and agents that are available through Symantec consulting services, contact your Symantec sales representative.

Intelligent Monitoring Framework (IMF) is an extension to the VCS Agent Framework, that allows the VCS agents to register the resources to be monitored so as to receive immediate notification of resource state changes without having to periodically poll the resources.

The following VCS agents are IMF-aware:

- Mount
- Process
- Application

Note: Intelligent Monitoring Framework (IMF) is supported for VxFS and CFS mounts only.

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

- Collect, analyze, and report on server configurations across UNIX or Windows environments. You can use this data to do the following:
  - Assess whether your systems are ready to install or upgrade Symantec enterprise products
  - Tune environmental parameters so you can increase performance, availability, and use
  - Analyze your current deployment and identify the Symantec products and licenses you are using
- Upload configuration data to the SORT Web site, so you can share information with coworkers, managers, and Symantec Technical Support

- Compare your configurations to one another or to a standard build, so you can determine if a configuration has "drifted"
- Search for and download the latest product patches
- Get notifications about the latest updates for:
  - Patches
  - Hardware compatibility lists (HCLs)
  - Array Support Libraries (ASLs)
  - Array Policy Modules (APMs)
  - High availability agents
- Determine whether your Symantec enterprise product configurations conform to best practices
- Search and browse the latest product documentation
- Look up error code descriptions and solutions

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

To access SORT, go to:

http://sort.symantec.com

Important release information

- The latest product documentation is available on the Symantec Web site at: http://www.symantec.com/business/support/overview.jsp?pid=15107
- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website: http://entsupport.symantec.com/docs/335001
- For the latest patches available for this release, go to: http://sort.symantec.com/

Configuration changes specific to RHEL6

This section describes configuration changes specific to Red hat Enterprise Linux (RHEL) 6.

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 Veritas Cluster Server 5.1SP1

This section lists the changes in Veritas Cluster Server 5.1SP1.

Changes related to the installation and upgrades

The product installer includes the following changes in 5.1 SP1 PR2.

Rolling upgrade support

To reduce downtime, the installer supports rolling upgrades. A rolling upgrade requires little or no downtime. A rolling upgrade has two main phases. In phase 1, the installer upgrades kernel packages on a subcluster. In phase 2, non-kernel packages are upgraded.

All high availability products support a rolling upgrade. You can perform a rolling upgrade from 5.1 or from any RPs to the current release.

You can perform a rolling upgrade using the script-based or Web-based installer.

See the *Veritas Cluster Server Installation Guide*.

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

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

- The installer detects and lists the aggregated links that you can choose to configure as private heartbeat links.
- The installer provides an option to detect NICs on each system and network links, and sets link priority to configure LLT over Ethernet.
- The installer provides an option to configure LLT over UDP.
- The installer now supports VCS cluster configuration up to 64 nodes.

See the *Veritas Cluster Server Installation Guide*.

Installer supports configuration of non-SCSI3 based fencing

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

See the *Veritas Cluster Server Installation Guide*.

Web-based installer supports configuring VCS cluster in secure mode

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

See the *Veritas Cluster Server Installation Guide*.

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

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

See the *Veritas Cluster Server Installation Guide*.

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

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

See the *Veritas Cluster Server Installation Guide*.

The Web-based installer supports adding nodes

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

The installer provides automated, password-less SSH configuration

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

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

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

The installer can check product versions

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

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

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

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

Installer changes related to I/O fencing configuration

In VCS 5.1 or SFHA 5.1, the installer started I/O fencing in disabled mode even if you had not chosen to configure I/O fencing. In 5.1 SP1, if you did not choose to configure I/O fencing during the product configuration, then the installer does not start I/O fencing in disabled mode in VCS and SFHA clusters.

However, if you upgrade VCS or SFHA from 5.1 to 5.1SP1, the installer retains the I/O fencing configuration from the previous version.

Upgrade changes

The following lists the upgrade changes in this release.

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

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

Upgrade matrix:

- 5.0MP3 to 5.1SP1
- 5.0MP4 to 5.1SP1

- 5.1 to 5.1 SP1
- 5.1RP1 to 5.1SP1
- 5.1RP2 to 5 .1SP1

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

Upgrades that follow any other upgrade paths require a reboot.

Packaging updates

The following lists package changes in this release.

New VRTSamf package

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

VRTSacclib package is no longer shipped with VCS 5.1SP1

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

Changes to the VCS engine

The HAD can exchange messages up to 64KB size

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

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

1. Maximum message size = 64KB
2. Maximum object name size = 1KB
3. Maximum attribute name size = 1KB
4. Maximum scalar attribute value size = 4KB

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

Note: Points 2 through 10 were already supported in 5.1 release.

New attributes

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

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.

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 5.1SP1

Application Agent attributes

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

RemoteGroup agent attribute

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

DiskGroup agent attribute

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

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

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.

RVGPrimary agent attribute

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

NotifierSourceIP agent attribute

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

SambaServer agent attributes

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

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.
- ResourceSet: A resource set is used to define a subset of processors in the system. If a resource set is specified for a workload partition, it can use the processors within the specified resource set only. The value of the ResourceSet attribute is the name of the resource set created using the `mkrset` command. If set, the agent configures the WPAR to use only the resource set specified by this attribute.
- WorkLoad: Allows modification of resource control attributes of WPAR - shares_CPU and shares_memory. This attribute has two keys, CPU and MEM. The key CPU is used to specify the number of processor shares that are available to the workload partition. The key MEM is used to specify the number of memory shares that are available to the workload partition.

Changes to bundled agents

This section describes changes to the bundled agents for VCS.

New bundled agent

- VolumeSet agent: The VolumeSet agent brings online, takes offline, and monitors a Veritas Volume Manager (VxVM) volume set. Use this agent to make a volume set highly available.

Support for Veritas dynamic multi-pathing

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

- LVMVolumeGroup agent

About the ReturnIntOffline attribute of RemoteGroup agent

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

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

About the RemotePartial option

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

About the RemoteOffline option

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

About the RemoteFaulted option

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

Configuring RemoteGroup resources in parallel service groups

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

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

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

- The RemoteGroup resource continues to monitor the remote service group even when the resource is offline.
- The RemoteGroup resource does not take the remote service group offline if the resource is online anywhere in the cluster.

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

Changes to SambaServer and NetBios agents

The SambaServer and NetBios agents contain the following changes:

- VCS SambaServer and NetBios agents now support multiple resource instances on a node.
- These agents now support non default Samba configuration file and pid file names.

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

Changes to database agents

Changes to the Oracle agent

- Oracle agent now supports IMF monitoring.

Changes to the Sybase agent

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

Changes to LLT

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

- LLT startup time through the LLT init script is now optimized to use a constant time. LLT takes less than 16 seconds to start irrespective of the number of links specified in /etc/llttab file.
  In the previous releases, LLT took around (5 * number_of_links_specified _in_the_/etc/llttab_file) seconds to start.
- The lltstat command includes the following new options:
  - lltstat -nv active
    This command filters the output of `lltstat -nv` to display the status of only the active nodes in the cluster.
  - lltstat -nv configured
    This command filters the output of `lltstat -nv` to display the status of only the configured nodes in the cluster. Configured nodes include active nodes and any additional nodes which are listed in the `/etc/llthosts` file.

  See the `lltstat` manual page for more information.
- Support for different link speeds for LLT links
  LLT now removes the restriction to use private NICs with same media speed. You can now use different media speed for the private NICs and configure the NICs with lesser speed as low-priority links to enhance LLT performance.
- Support for destination-based load balancing
  LLT now also provides destination-based load balancing where the LLT link is chosen based on the destination node id and the port. With destination-based load balancing, LLT sends all the packets of a particular destination on a link.

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

Changes to GAB

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

- GAB logging daemon
  GAB implements a distributed network protocol. For situations when GAB decides to take the drastic action of killing its userland client process or panicking a node to resolve an issue, data from the affected node alone may not suffice for a meaningful support analysis. The new gablogd daemon attempts to address this issue. GAB starts this daemon by default at GAB configuration time.
  See the *Veritas Cluster Server Administrator's Guide* for more information.

Changes to I/O fencing

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

Support for preferred fencing

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

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

Support for Non-SCSI3 fencing

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

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

Enhancements to server-based fencing

This release includes the following enhancements and new features related to server-based fencing:

- Single CP-server based fencing
  Support to use a single highly available CP server that is configured on an SFHA cluster to provide server-based fencing support for multiple application clusters
- Support for CP server on AIX and HP-UX
  CP server now supports AIX and HP-UX in addition to Linux and Solaris operating systems.

Support to migrate between fencing modes when the cluster is running

The vxfenswap utility now supports migrating between disk-based and server-based fencing configurations in a cluster that is running.

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

Changes to global clustering

VCS global clustering monitors and manages the replication jobs and clusters at each site. In the event of a site outage, global clustering controls the shift of replication roles to the Secondary site, bring up the critical applications and redirects client traffic from one cluster to the other.

Before Release 5.1SP1, if there was a disaster at the Primary site or a network disruption, the applications were taken offline on the original Primary and failed over to the Secondary. When the original Primary returned or the network disruption was corrected, you had the following options:

- Manually resynchronize the original Primary with the data from the new Primary, once the original Primary comes back up. The applications are only active on the new Primary site.
- Automatically resynchronize the original Primary with the data from the new Primary, once the original Primary comes back up. The applications are only active on the new Primary site.

Beginning in Release 5.1SP1, you have a third option. Applications can be active on both the original Primary and Secondary sites. After the original Primary returns or the network disruption is corrected, you have the option of specifying which site is the Primary going forward. This option is called the primary-elect feature, and it is enabled through the VCS global clustering.

The key difference between the primary-elect feature and the other options is that if a network disruption occurs, applications continue to run on the Primary site and they are also failed over to the Secondary. This feature lets you maintain application availability on both sites while the network is down.

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 26.

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://entsupport.symantec.com/docs/330441

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

Supported Linux operating systems

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

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

http://entsupport.symantec.com/docs/335001

The Veritas 5.1 SP1 PR2 release supports only the following 64-bit operating systems and hardware:

- Red Hat Enterprise Linux 6 (RHEL 6) (2.6.32-71.el6 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64).

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

Symantec supports only Red Hat distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported Red Hat errata and updates is available in the following TechNote. Read this TechNote before you install Symantec products.

http://entsupport.symantec.com/docs/335001

Required Linux RPMs for VCS

Make sure you installed 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 1-1 lists the RPMs that VCS requires for a given Linux operating system.

Table 1-1 Required RPMs

Operating system	Required RPMs
RHEL 6	compat-libstdc++-33-3.2.3-69.el6.x86_64.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.i686.rpm pam-1.1.1-4.el6.x86_64.rpm

Supported software for VCS

VCS supports the following volume managers and file systems:

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

  VCS 5.1 SP1 PR2 supports the following versions of SF:

  - SF 5.1 SP1 PR2
    - VxVM 5.1 SP1 PR2 with VxFS 5.1 SP1 PR2

Supported VCS agents

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

Features no longer supported

No longer supported agents and components

VCS no longer supports the following:

- Configuration wizards
- Disk agent
- CampusCluster agent
- SANVolume agent
- VRTSWebApp
- Apache agent configuration wizard
- Oracle 8.0.x, Oracle 8.1.x, and Oracle 9i - not supported by the updated Oracle agent.
- VCS documentation package (VRTSvcsdc)
  The VCS documentation package (VRTSvcsdc) is deprecated. The software disc contains the documentation for VCS in Portable Document Format (PDF) in the *cluster_server/docs* directory.
  Symantec recommends copying pertinent documents from the disc to your system directory /opt/VRTS/docs for reference.
- hahbsetup tool. This tool is removed as no supported feature requires this tool.
- VRTScutil package. This package is no longer supported.

Fixed issues

Issues fixed in 5.1 SP1 PR2 release

This section covers the incidetns that were fixed in 5.1 SP1 PR2 release.

Table 1-2 Issues fixed in release

Incident	Description
2220674	RemoteGroup agent crashes if VCSAPI log level is set to a non-zero value.
2187918	HAD dumps core while overriding the static attribute to resource level.

Table 1-2 Issues fixed in release *(continued)*

Incident	Description
2211333	Parent service group does not fail over in case of online local firm dependency with child service group.
2198682	Warning message must be logged in vxfend_A.log when fencing is configured with 3 disks even when single_cp=1.
2208017	The ResourceInfo display is restricted to 20 characters and does not display all key-values if they exceed the 20-character limit. Therefore, only the complete key-values falling within the 20-character limit are displayed.

Issues fixed in 5.1SP1 release

This section covers the incidents that are fixed in 5.1SP1 release.

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

See "Documentation" on page 60.

LLT, GAB, and I/O fencing fixed issues

Table 1-3 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-3 LLT, GAB, and I/O fencing fixed issues

Incident	Description
1908938	[GAB] In a large cluster, cascaded lowest node failures result in GAB panic during sequence space recovery.
1840826	[GAB] Prevent 'gabconfig -c' while port 'a' is in the middle of iofence processing.
1861439 1849527	[LLT] Removing the LLT links from a single node in a four-node cluster causes other nodes to panic.
2066020	[LLT] The `dlpiping` utility exits with an error similar to "dlpiping: send ECHO_REQ failed."

Table 1-3 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
2005045	[LLT] The `hastart` command fails to start HAD on one of the nodes with message "GabHandle::open failed errno = 16" in syslog after HAD is stopped on all the nodes in the cluster simultaneously.
1859023	[LLT] The `lltconfig -T query` command displays a partially incorrect output
1846387 2084121	[Fencing] The `vxfenswap` and the `vxfentsthdw` utilities fail when rsh or ssh communication is not set to the same node.
1922413	[Fencing] The `vxfentsthdw` utility should detect storage arrays which interpret NULL keys as valid for registrations/reservations.
1847517	[Fencing] The `vxfenswap` utility has an incorrect usage message for -n option
1992560	[Fencing] The `vxfentsthdw` utility uses scp to communicate with the local host.
2098065	[Fencing] The `vxfenclearpre` utility cannot clear keys from coordinator disks and data disks when there is a preexisting split brain.
1512956	[Fencing] The `vxfenclearpre` utility displays error messages
2143933	[VxCPS] For a four-node cluster, the installer fails to configure server-based fencing which uses three CP servers as its coordination points. The process fails while registering the CP clients on the third CP server.
2097935	[VxCPS] Need strict host name matching in coordination point installer.

Bundled agents fixed issues

Table 1-4 lists the fixed issues for bundled agents.

Table 1-4 Bundled agents fixed issues

Incident	Description
252354	Application agent does not pass the value of CleanReason to the script that the Agent executes as per the value of the CleanProgram attribute.
1969191	If you restart HAD, online Share resources may go offline.

Table 1-4 Bundled agents fixed issues *(continued)*

Incident	Description
2001039	In a frozen local service group, if a resource goes offline outside VCS control, a RemoteGroup resource also goes offline.
2019904	The OS determines the interface that the notifier uses to send packets. The user must be optionally able to specify the interface. This is an enhancement.
2045972	If you upgrade to VCS 5.1, Application resources may go to the FAULTED state.
2101577	When a MultiNICA resource fails over, the agent deletes the default route for the active device. For more information, See "While upgrading the VCS stack from a version prior to VCS 5.1, reconfiguration of MultiNICA IPv4RouteOptions attribute is required" on page 40.

VCS engine fixed issues

Table 1-5 lists the fixed issues for VCS engine.

Table 1-5 VCS engine fixed issues

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

Table 1-5 VCS engine fixed issues *(continued)*

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

Table 1-5 VCS engine fixed issues *(continued)*

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

Enterprise agents fixed issues

Table 1-6 lists the fixed issues for enterprise agents.

Table 1-6 Enterprise agents fixed issues

Incident	Description
776230	The Sybase agent offline script must support the databases that require a longer time to shut down. This is an enhancement.
1859080	If the Owner attribute of the Sybase agent is nine characters or more in length, then the agent does not detect a resource that comes online.
2117378	When the monitor entry function encounters an invalid process ID (PID) in a PID file, the function fails to detect the exit status of the process. This behavior occurs if the PID has an empty line below it. The function also accordingly fails to detect the correct state of a failed DB2 Connect instance.

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.

Known issues in VCS 5.1SP1PR2

This section covers the known issues in 5.2SP1PR2 release.

Upgrades from previous versions of VCS not supported

This release does not support upgrades of previous versions of VCS running on RHEL 5. You must uninstall VCS, upgrade the operating system, then reinstall VCS.

Upgrading the Linux kernel when the root volume is under DMP control

This section includes the procedures for upgrading the Linux kernel when the root volume is under DMP control.

Linux kernel can be upgraded on RHEL6 systems without turning off the DMP native support. Only one reboot is required to bring system LVM volume on DMP after kernel upgrade.

To update the kernel on a RHEL6 system

1. Update kernel with the rpm command.

   ```
   # rpm -Uvh kernel_rpm --force
   ```

2. Turn on the dmp_native_support tunable:

   ```
   # vxdmpadm settune dmp_native_support=on
   ```

   This enables booting with new kernel with LVM devices with DMP.

3. Reboot.

SELinux error during installation of VRTSvcsag package

Description: During the installation of VRTSvcsag package on RHEL 6 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, if you observe this issue.

Trace messages from the gablogd daemon on the console for RHEL6 or later

On RHEL6 or later, the `gablogd` daemon prints informational and trace messages similar to the following [2139883]:

```
INFO: task gablogd:22812 blocked for more than 120 seconds.
"echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this
 message.
gablogd  D ffff81004100daa0   0 22812   1     23106 22809 (NOTLB)
 ffff810faf539e38 0000000000000082 000000000000084c 0000000000000001
 ffff810faf539de8 0000000000000007 ffff810fc2a130c0 ffff810138ee8100
 000019f130082599 0000000000018572 ffff810fc2a132a8 00000001f76c3d63
Call Trace:
 [<ffffffff88ee3690>] :gab:gab_linux_sv_wait+0x53/0x68
 [<ffffffff8008e68d>] default_wake_function+0x0/0xe
 [<ffffffff88ecd4c8>] :gab:gab_daemonlog+0xae1/0xc52
 [<ffffffff88ee326c>] :gab:gab_linux_ioctl+0x10e/0x1a3
 [<ffffffff88ee331d>] :gab:gab_linux_compat_ioctl+0x1c/0x20
 [<ffffffff800fbe53>] compat_sys_ioctl+0xc5/0x2b2
 [<ffffffff8006249d>] sysenter_do_call+0x1e/0x76
```

Workaround: As the operating system message indicates, set the following:

```
echo 0 > /proc/sys/kernel/hung_task_timeout_secs
```

Error messages in syslog (2213651)

If you install or uninstall a product on a node, you may see the following warnings in syslog: /var/log/messages. These warnings are harmless and can be ignored.

When installing, the log may display:

```
Dec  3 17:21:26 cdc-d2950130 kernel: type=1400 audit(1291368086.666:20): avc:
denied  { write } for  pid=16553 comm="semanage"
path="/var/tmp/installer-201012031718JNC/install.VRTSvxvm.cdc-d2950130"
dev=sdok1
ino=1443459 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file

Dec  3 17:21:32 cdc-d2950130 kernel: type=1400 audit(1291368092.123:22): avc:
denied  { write } for  pid=16556 comm="semanage"
path="/var/tmp/installer-201012031718JNC/install.VRTSvxvm.cdc-d2950130"
dev=sdok1
ino=1443459 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file
```

```
Dec  3 17:21:55 cdc-d2950l30 kernel: type=1400 audit(1291368115.245:24): avc:
denied  { write } for  pid=16950 comm="semodule"
path="/var/tmp/installer-201012031718JNC/install.VRTSvxfs.cdc-d2950l30"
dev=sdok1
ino=1443463 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file

Dec  3 17:21:55 cdc-d2950l30 kernel: type=1400 audit(1291368115.312:25): avc:
denied  { write } for  pid=16950 comm="semodule"
path="/var/tmp/installer-201012031718JNC/install.VRTSvxfs.cdc-d2950l30"
dev=sdok1
ino=1443463 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file
```

When uninstalling, the log may display:

```
Dec  3 17:29:00 cdc-d2950l30 kernel: type=1400 audit(1291368540.794:27): avc:
denied  { write } for  pid=19151 comm="semodule"
path="/var/tmp/uninstallsfcfsrac-201012031725oCf/uninstall.VRTSvxfs.cdc-
d2950l30"
dev=sdok1 ino=1186738 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file

Dec  3 17:29:00 cdc-d2950l30 kernel: type=1400 audit(1291368540.866:28): avc:
denied  { write } for  pid=19151 comm="semodule"
path="/var/tmp/uninstallsfcfsrac-201012031725oCf/uninstall.VRTSvxfs.cdc-
d2950l30"
dev=sdok1 ino=1186738 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file

Dec  3 17:29:45 cdc-d2950l30 kernel: type=1400 audit(1291368585.473:30): avc:
denied  { write } for  pid=19683 comm="semanage"
path="/var/tmp/uninstallsfcfsrac-201012031725oCf/uninstall.VRTSvxvm.cdc-
d2950l30"
dev=sdok1 ino=1186621 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file

Dec  3 17:29:50 cdc-d2950l30 kernel: type=1400 audit(1291368589.975:32): avc:
denied  { write } for  pid=19687 comm="semanage"
path="/var/tmp/uninstallsfcfsrac-201012031725oCf/uninstall.VRTSvxvm.cdc-
d2950l30"
dev=sdok1 ino=1186621 scontext=unconfined_u:system_r:semanage_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:user_tmp_t:s0 tclass=file
```

The syslog gets flooded with SELinux messages

The SELinux audit feature logs lot of SELinux messages in the syslog. [2220265]

The following is an example of the SELinux messages:

```
Dec 13 22:29:12 cdc-d2950133 setroubleshoot: SELinux is preventing
 /sbin/rpc.statd access to a leaked /var/VRTSvcs/log/NFSRestart_A.log file
 descriptor. For complete SELinux messages,  run
sealert -l 8cdfb88d-7ab2-4815-b5e5-ea6696bfa89e
```

This does not affect the VCS functionality. This is a SELinux behavior. Such messages occur when SELinux dontaudit rules are enabled for policies, as dontaudit feature of SELinux does an access check of the applications.

Workaround: Run the following command to disable the SELinux messages.

```
semodule -D -B
```

Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 SP1 PR2

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 SP1 PR2 release.

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

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

RVGPrimary online script does not function correctly (1949293)

The RVGPrimary online script does not function correctly.

Issues with bunker replay

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

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

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

```
unreachable...

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

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

The nfsconf.vfd action entry point does not work

When you run nfsconf.vfd action entry point, it displays the following error message. [2223678]

```
Undefined subroutine &main:: uname called at
/opt/VRTSvcs/bin/NFSRestart/actions/nfsconf.vfd line 123.
```

Workaround: Please ensure that the nfslock service is disabled. Run following command to ensure that the service is disabled using:

```
# /sbin/chkconfig --list nfslock
```

Known issues in VCS 5.1SP1

This section covers the known issues in 5.1SP1 release.

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.

Issues related to installation

This section describes the known issues during installation and upgrade.

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

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

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

Errors observed during partial upgrade of SFHA

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

Workaround

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

Manual upgrade of VRTSvlic package loses keyless product levels

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

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

To manually upgrade the VRTSvlic package

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

   ```
   # vxkeyless display
   ```

2. Set the product level to NONE.

   ```
   # vxkeyless set NONE
   ```

3. Upgrade the VRTSvlic package.

   ```
   # rpm -Uvh --nopreun VRTSvlic-3.02.51.010-0.x86_64.rpm
   ```

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

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

Issue with soft links getting deleted in a manual upgrade

While performing a manual upgrade (from 5.1 to 5.1SP1) of the VRTSvlic package, 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-3.02.51.010-0.x86_64.rpm`

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

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 1-7 Whether attributes are configured and required actions that you need to perform during upgrade

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

Issues with keyless licensing reminders after upgrading VRTSvlic

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

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

To resolve this issue, perform the following steps:

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

   ```
   # vxkeyless set NONE
   ```

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

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

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

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

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

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

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

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

Splitting a cluster that uses server-based fencing is currently not supported. [2110148]

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

However, if you use the installer to reconfigure the VCS, the installer retains the same cluster UUID of *clus1* in both *clus1A* and *clus1B*. If both *clus1A* and *clus1B* use the same CP servers for I/O fencing, then the CP server allows registration

only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: None.

The installer enters an infinite loop because the cluster is already running in secure mode while configuring server-based fencing (2166599)

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

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

System moves to low system resources while bringing large number process resources online (2061338)

When large number of resource like Process resources (1500 in number) is configured on a Linux system (kernel version 2.6.18-92.el5)than agent gets killed. The reason for is fork failure and fork return with an error number 513. A bug against Red Hat "Issue #822253" exits. Red Hat has confirmed that this is their issue. This issue does't appear on kernel version 2.6.18-194.el5 which they claim to have the fix for this issue.

Operational issues for VCS

Issues with configuration of resource values

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

```
VCS WARNING V-16-2-13806 Thread(1437547408) ArgListValues overflow;

Cannot append values more than upper limit of (425).
```

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

LVM SG transition fails in all paths disabled status

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 LVM VG 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 online manually, or deactivate it and export the volume group before starting VCS.

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

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

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

```
::1         localhost
```

Issues related to the VCS engine

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

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

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

VCS Engine logs messages when it eventually connects to a remote cluster

Description: In a global cluster, if a local cluster fails to connect with a remote cluster in the first attempt but succeeds eventually, then you may see the following warning messages in the engine logs. [2110108]

```
VCS WARNING V-16-1-10510 IpmHandle: pen Bind Failed.
unable to bind to source address 10.209.125.125. errno = 67
```

Workaround: There is currently no workaround for this issue. This issue has no impact on any functionality.

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.

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

The AutoAddSystemToCSG attribute determines whether the newly joined or added systems in a cluster become part of the SystemList of the ClusterService service group if the service group is configured. The value 1 (default) indicates that the new systems are added to SystemList of ClusterService. AutoAddSystemToCSG has an impact only when you execute the hasys –add command or when a new node joins the cluster. [2159139]

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

```
# hagrp –modify ClusterService SystemList –add newnode n
# hagrp –modify ClusterService AutoStartList –add newnode
```

Workaround

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

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

```
# hagrp –modify ClusterService SystemList –delete newnode
# hagrp –modify ClusterService AutoStartList –delete newnode
```

Issues related to the bundled agents

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

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. [2140342]

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

Issues related to the agent framework

Agent may fail to heartbeat under heavy load (2073018)

Description: 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 the following log when it stops and restarts the agent:

Resolution: 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

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

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

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

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

For example, you have resource (*rD*), which depends on a resource's attribute value (*rT*:*Attr_rt*). When you delete the target resource (*rT*), and re-add it (*rT*), the

dependent resource (*rD*) does not get the correct value for the attribute (*Attr_rt*). [1539927]

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

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

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

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

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

Agent performance and heartbeat issues

Depending on the system capacity and the number of resources configured under VCS, the agent may not get enough CPU cycles to function properly. This can prevent the agent from producing a heartbeat synchronously with the engine. If you notice poor agent performance and an agent's inability to heartbeat to the engine, check for the following symptoms.

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

```
FFDC_AGFWMain_729_agent_type.log  FFDC_AGFWTimer_729_agent_type.log core
FFDC_AGFWSvc_729_agent_type.log   agent_typeAgent_stack_729.txt
```

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

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

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

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

Issues related to global clusters

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

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

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

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. (2107386)

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 port stats sometimes shows recvcnt larger than recvbytes

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. [1788315]

This does not impact the LLT functionality.

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

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

Issues related to GAB

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

Trace messages from the gablogd daemon on the console for RHEL5 Update 5 or later

On RHEL5 Update 5 or later, the `gablogd` daemon prints informational and trace messages similar to the following [2139883]:

```
INFO: task gablogd:22812 blocked for more than 120 seconds.
"echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
gablogd     D ffff81004100daa0     0 22812     1      23106 22809 (NOTLB)
 ffff810faf539e38 0000000000000082 000000000000084c 0000000000000001
 ffff810faf539de8 0000000000000007 ffff810fc2a130c0 ffff810138ee8100
 000019f130082599 0000000000018572 ffff810fc2a132a8 00000001f76c3d63
Call Trace:
 [<ffffffff88ee3690>] :gab:gab_linux_sv_wait+0x53/0x68
 [<ffffffff8008e68d>] default_wake_function+0x0/0xe
 [<ffffffff88ecd4c8>] :gab:gab_daemonlog+0xae1/0xc52
 [<ffffffff88ee326c>] :gab:gab_linux_ioctl+0x10e/0x1a3
 [<ffffffff88ee331d>] :gab:gab_linux_compat_ioctl+0x1c/0x20
 [<ffffffff800fbe53>] compat_sys_ioctl+0xc5/0x2b2
 [<ffffffff8006249d>] sysenter_do_call+0x1e/0x76
```

Workaround: As the operating system message indicates, set the following:

```
echo 0 > /proc/sys/kernel/hung_task_timeout_secs
```

Issues related to I/O fencing

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

CP server cannot bind to multiple IPs (2085941)

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

Resolution: No known resolution for this issue.

Issues related to Symantec Product Authentication Service with VCS

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

The vcsat and cpsat commands may appear to be hung

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

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

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

Workaround:

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

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

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

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

Software limitations

Software limitation of 5.1 SP1 PR2 release

This section covers the software limitations of the 5.1 SP1 PR2 release.

Limitations related to IMF

- IMF registration on Linux for “bind” file system type is not supported.
- 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.

Limitation related to LLT

LLT may not start on the node which does not have 'BOOTPROTO=none' in the interface configuration file ifcfg-ethX. [2191023]

Limitation related to CPI

CPI may not work if the security configuration on the node has the iptables firewall running. [2195696]

Workaround: Modify the iptables chain rules to open port 2821.

Software limitations of VCS 5.1SP1 release

This section covers the software limitations of VCS 5.1SP1 release.

Limitations related to installing and upgrading VCS

Limitation when you use the installer from a remote system

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. [589334]

Limitations related to the 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 the 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

Limitations of the DiskGroup agent

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

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

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.

Share agent

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 DiskReservation agent has a reserver module in the kernel mode that reserves disks persistently. Any driver that works correctly with the scsiutil utility shipped with the VRTSvcsdr package is supported. Refer to the manual page for scsiutil functionality.

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 if the value of the system level attribute autostartvolumes in Veritas Volume Manager is set to On, irrespective of the value of the StartVolumes attribute in VCS.

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 global clusters

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

Systems in a cluster must have same system locale setting

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

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.

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.

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).

Use the VCS 5.1 Java Console to manage clusters

Cluster Manager (Java Console) from versions lower than VCS 5.1 cannot be used to manage VCS 5.1SP1 clusters. Symantec recommends using the latest version of Cluster Manager.

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

Run Java Console on a non-cluster system

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

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

VCS Cluster Manager 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”.

Limitations related to I/O fencing

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

Stopping systems in clusters with I/O fencing configured

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

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

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

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

Documentation errata

The following sections cover additions or corrections for Document version: 5.1SP1PR2.0 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 60.

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

Veritas Cluster Server Installation Guide

Updates for the procedure to set up non-SCSI3 fencing in virtual environments manually

Topic: Setting up non-SCSI3 fencing in virtual environments manually

Some information is missing in the procedure that is documented in the Installation Guide. Refer to the following procedure.

Setting up non-SCSI3 fencing in virtual environments manually

To manually set up I/O fencing in a non-SCSI-3 PR compliant setup

1. Configure I/O fencing in customized mode with only CP servers as coordination points.

2. Make sure that the VCS cluster is online and check that the fencing mode is customized.

   ```
   # vxfenadm -d
   ```

3. Make sure that the cluster attribute UseFence is set to SCSI3.

   ```
   # haclus -value UseFence
   ```

4. On each node, edit the /etc/vxenviron file as follows:

   ```
   data_disk_fencing=off
   ```

5. On each node, edit the /etc/sysconfig/vxfen file as follows:

   ```
   vxfen_vxfnd_tmt=25
   ```

6. On each node, edit the /etc/vxfenmode file as follows:

   ```
   loser_exit_delay=55
   vxfen_script_timeout=25
   ```

   Refer to the sample /etc/vxfenmode file.

7. On each node, set the value of the LLT sendhbcap timer parameter value as follows:

- Run the following command:

```
lltconfig -T sendhbcap:3000
```

- Add the following line to the /etc/llttab file so that the changes remain persistent after any reboot:

```
set-timer senhbcap:3000
```

8 On any one node, edit the VCS configuration file as follows:

- Make the VCS configuration file writable:

```
# haconf -makerw
```

- For each resource of the type DiskGroup, set the value of the MonitorReservation attribute to 0 and the value of the Reservation attribute to NONE.

```
# hares -modify <dg_resource> MonitorReservation 0

# hares -modify <dg_resource> Reservation "NONE"
```

- Run the following command to verify the value:

```
# hares -list Type=DiskGroup MonitorReservation!=0

# hares -list Type=DiskGroup Reservation!="NONE"
```

The command should not list any resources.

- Modify the default value of the Reservation attribute at type-level.

```
# haattr -default DiskGroup Reservation "NONE"
```

- Make the VCS configuration file read-only

```
# haconf -dump -makero
```

9 Make sure that the UseFence attribute in the VCS configuration file main.cf is set to SCSI3.

10 To make these VxFEN changes take effect, stop and restart VxFEN and the dependent modules

- On each node, run the following command to stop VCS:

```
# /etc/init.d/vcs stop
```

- After VCS takes all services offline, run the following command to stop VxFEN:

```
# /etc/init.d/vxfen stop
```

- On each node, run the following commands to restart VxFEN and VCS:

```
# /etc/init.d/vxfen start
# /etc/init.d/vcs start
```

Veritas Cluster Server Administrator's Guide

This section covers the documentation errata in the *Veritas Cluster Server Administrator's Guide*.

Correction for setting up a disaster recovery fire drill

Topic: Setting up a disaster recovery fire drill

Issue: The content below is incorrect:

After the fire drill service group is taken offline, reset the value of the ReuseMntPt attribute to 1 for all Mount resources.

Use the following corrected information:

After the fire drill service group is taken offline, reset the value of the ReuseMntPt attribute to 0 for all Mount resources.

Documentation

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

Documentation set

Table 1-8 lists the documents for Veritas Cluster Server.

Table 1-8 Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_51sp1pr2_lin.pdf
Veritas Cluster Server Release Notes	vcs_notes_51sp1pr2_lin.pdf

Table 1-8 Veritas Cluster Server documentation *(continued)*

Title	File name
Veritas Cluster Server Administrator's Guide	vcs_admin_51sp1pr2_lin.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_51sp1pr2_lin.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev_51sp1pr2.pdf
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	vcs_vvr_agent_51sp1pr2_lin.pdf
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_51sp1pr2_lin.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_51sp1pr2_lin.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_51sp1pr2_lin.pdf

Table 1-9 lists the documentation for Veritas Volume Replicator.

Table 1-9 Veritas Volume Replicator documentation

Document title	File name
Veritas Volume Replicator Administrator's Guide	vvr_admin_51sp1pr2_lin.pdf
Veritas Volume Replicator Planning and Tuning Guide	vvr_planning_51sp1pr2_lin.pdf
Veritas Volume Replicator Advisor User's Guide	vvr_advisor_users_51sp1pr2_lin.pdf

Table 1-10 lists the documentation for Symantec Product Authentication Service (AT).

Table 1-10 Symantec Product Authentication Service documentation

Title	File name
Symantec Product Authentication Service Release Notes	vxat_notes.pdf
Symantec Product Authentication Service Administrator's Guide	vxat_admin.pdf

Manual pages

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

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

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

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

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

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

See the `man`(1) manual page.

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
   ```