

Veritas™ Cluster Server Release Notes

Linux for IBM Power

5.0 Release Update 3



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: 5.0 RU3

Document version: 5.0RU3.0

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the *Veritas Cluster Server 5.0 Release Notes*.

The *Veritas Cluster Server 5.0 Release Notes* can be viewed at the following URL:

<http://entsupport.symantec.com/docs/283850>

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

www.symantec.com/techsupp/

Contacting Technical Support

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

www.symantec.com/business/support/assistance_care.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/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

Customer Service is available to assist with 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 maintenance 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

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance 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

Additional enterprise services

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

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

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

www.symantec.com

Select your country or language from the site index.

Release Notes

This document includes the following topics:

- [Introduction](#)
- [About Veritas Cluster Server](#)
- [Changes introduced in VCS 5.0 RU3](#)
- [Changes introduced in VCS 5.0 MP3](#)
- [Changes introduced in VCS 5.0 MP1](#)
- [Features introduced in VCS 5.0](#)
- [VCS system requirements](#)
- [No longer supported](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation errata](#)
- [VCS documentation](#)

Introduction

This document provides important information about Veritas Cluster Server (VCS) version 5.0 Release Update 3 for Linux for IBM Power. 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 release adds support for VCS on Linux on IBM Power and rolls forward changes introduced in the previous releases of VCS. The changes introduced in the previous releases are relevant to the platforms and architectures supported in those releases.

For the latest information on updates, patches, and software issues for this release, use the following TechNote on the Symantec Enterprise Support website:

<http://entsupport.symantec.com/docs/281993>

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

About Veritas Cluster Server

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

About VCS agents

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

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

The Veritas High Availability Agent Pack gives you access to agents that provide high availability for third-party storage solutions. Contact your Symantec sales representative for information about agents included in the agent pack, agents under development, and agents that are available through Symantec consulting services.

VCS provides a framework that allows for the creation of custom agents. Create agents in situations where the Veritas High Availability Agent Pack, the bundled agents, or the agents for enterprise applications do not meet your needs. You can also request a custom agent through Symantec consulting services.

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

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

See "[Supported VCS agents](#)" on page 32.

Changes introduced in VCS 5.0 RU3

This section lists the changes introduced in this release of VCS.

Symantec Product Authentication Service supports LDAP

This release adds support for Symantec Product Authentication Service (VxAT) with LDAP.

For more information on manually adding a domain, refer to the *Symantec Product Authentication Service Administrator's Guide version 4.3*.

Changes to Cluster Management Console

The VCS 5.0 RU3 installer does not include the Cluster Management Console 5.0 cluster connector and single cluster management console components as part of the optional VCS RPMs.

However, if you already have a management server of VCS Management Console 5.5 or later, you can use the direct connection to manage the newly configured VCS clusters. Symantec recommends that you install and configure the latest version 5.5 SP1 of VCS Management Console.

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

Change to the JRE package

The VTRSjre package is not shipped with this release. Symantec recommends using the native JRE package version 1.5 or later.

Persistent interface naming for SLES10

VCS requires the names of the NICs to be persistent across system reboots. For SLES10, a special configuration is needed, which is explained as follows:

To configure persistent interface naming on SLES10

- 1 If you already have the PERSISTENT_NAME directive in any of the NIC configuration files, remove those lines from those configuration files.
- 2 For each such PERSTSTENT_NAME directive, add a line to `/etc/udev/rules.d/30-net_persistent_names.rules`.

For example, if `/etc/sysconfig/network/ifcfg-eth-id-00:06:5b:38:bf:f0` contains the line “PERSISTENT_NAME=link02”, then delete this line from this file and add the following line to `/etc/udev/rules.d/30-net_persistent_names.rules`:

```
SUBSYSTEM=="net", ACTION=="add",  
SYSFS{address}=="00:06:5b:38:bf:f0",  
IMPORT="/lib/udev/rename_netiface %k link02"
```

- 3 Make similar entries in `/etc/udev/rules.d/30-net_persistent_names.rules` for each interface whose name you want to make persistent across reboots

Read the file

`/usr/share/doc/packages/sysconfig/README.Persistent_Interface_Names` for more information.

Changes introduced in VCS 5.0 MP3

This section lists the changes introduced VCS 5.0 MP3.

Support for additional Linux distributions

This release adds support for the following Linux distributions:

- Red Hat Enterprise Linux 5.0 Update 1 and Update 2
- SUSE Linux Enterprise Server 10 SP1 and SP2

See “[Supported operating systems](#)” on page 29.

Support for VLAN interfaces

The NIC and MultiNICA agents now support VLAN interfaces. The agents do not configure the NICs, but can monitor them.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Changes to VCS agent for Oracle

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

- Support for Oracle ASM
- Support to detect intentional offline
- Support to choose CUSTOM as one of the start up options when you use Hot backup feature of Oracle
- Support for csh and tcsh shell when Health check monitoring is enabled
- Support for a new action entry point pfile.vfd

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for more information.

VCS agent for DB2 supports DB2 9.5

The Veritas Cluster Server Agent for DB2 supports DB2 9.5.

See “[Supported VCS agents](#)” on page 32.

See the *Veritas Cluster Server Agent for DB2 Installation and Configuration Guide* for more information.

Change in attributes

This release has the following changes for VCS attributes:

- AYATimeout - VCS heartbeat attribute
The default value of the heartbeat attribute AYATimeout is changed from 300 seconds to 30 seconds. [622413]
- Preonline - VCS service group attribute
You can now localize the Preonline attribute for the nodes in a cluster. [530440]
- AutoFailOver - VCS service group attribute
If you have configured system zones in campus clusters, you can fail over that service group manually across system zones.

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

New attributes

This release introduces the following new system attributes:

- HostMonitor—Monitors the usage of resources on the host.
- HostUtilization—Indicates the usage percentages of resources on the host.

This release introduces the following new service group attributes:

- PreSwitch—Indicates whether the VCS engine should switch a service group in response to a manual group switch.
- PreSwitching—Indicates that the VCS engine invoked the PreSwitch action function for the agent; however, the action function is not yet complete.

This release introduces the following new resource type level attribute:

- OfflineWaitLimit—Indicates the number of monitor intervals to wait for the resource to go offline after completing the offline procedure. Increase the value of this attribute if the resource is likely to take a longer time to go offline.

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

Support for intentional offline

Certain VCS agents can identify when an application has been intentionally shut down outside of VCS control. If an administrator intentionally shuts down an application outside of VCS control, VCS does not treat it as a fault. VCS sets the service group state as offline or partial, depending on the state of other resources in the service group. This feature allows administrators to stop applications without causing a failover. The intentional offline feature is available for agents registered as V5.1 or later.

See the *Veritas Cluster Server Agent Developer's Guide* for information on agents that support this feature.

VCS process to monitor resource utilization

VCS uses the HostMonitor daemon to monitor the resource utilization of CPU and Swap. VCS reports to the engine log if the resources cross the threshold limits that are defined for the resources.

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

Support for VxFS file system lock

If the VxFS file system has "mntlock=key" in its mount options, then you cannot unmount the file system without specifying the key. This prevents accidental unmounts when a resource is online.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information.

Changes to bundled agents

VCS 5.0 MP3 introduces the following new agents:

- **DiskGroupSnap**—Verifies the configuration and data integrity in a campus cluster environment.

See “**No longer supported**” on page 32.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for details.

Change in behavior: DiskGroup agent

If you set the value of the `PanicSystemOnDGLoss` attribute to 1, VCS panics the system when the disk group becomes unavailable irrespective of whether you use I/O fencing.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information.

DNS agent updates and monitors the hostname to IP address mapping

The DNS agent now updates and monitors the mapping of host names to IP addresses (A, AAAA, and PTR records), in addition to the canonical name (CNAME) mapping for a DNS zone when failing over nodes across subnets.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

New and modified attributes for the bundled agents in VCS 5.0 MP3

The new and modified attributes for the bundled agents in VCS 5.0 MP3 are as follows:

Table 1-1 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 5.0

Agent	New and modified attributes	Default Value
Apache		
New attributes		
	PidFile	
	IntentionalOffline	0
DiskGroup		
New attributes		
	UmountVolumes	0

Table 1-1 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 5.0 (*continued*)

Agent	New and modified attributes	Default Value
Modified attributes		
	SupportedActions	{ "license.vfd", "disk.vfd", "udid.vfd", "verifyplex.vfd", checkudid, numdisks, campusplex, joindg, splitdg, getvxvminfo, volinuse }
DNS		
New attributes		
	SupportedActions	{ "dig.vfd", "keyfile.vfd", "master.vfd" }
	ResRecord	
	CreatePTR	
	OffDelRR	
LVMVolumeGroup		
New attributes		
	SupportedActions	{ volinuse }
Mount		
New attributes		
	RegList	{ VxFSMountLock }
	VxFSMountLock	0
Modified attributes		
	SupportedActions	{ "mountpoint.vfd", "mounted.vfd", "vxfslic.vfd", "chgmtntlock", "mountentry.vfd" }
NFSRestart		

Table 1-1 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 5.0 (*continued*)

Agent	New and modified attributes	Default Value
New attributes		
	SupportedActions	{ "lockdir.vfd", "nfsconf.vfd" }
Share		
New attributes		
	SupportedActions	{ "direxists.vfd" }

Enhancements to VCS campus clusters

VCS 5.0 MP3 includes the following enhancements to campus cluster configurations:

- Support for campus clusters using VxVM remote mirror configuration
 In a campus cluster setup, you can configure VxVM diskgroups for remote mirroring.
- Support for fire drill in campus clusters using the DiskGroupSnap agent
 A fire drill tests the disaster-readiness of a configuration by mimicking a failover without stopping the application and disrupting user access. The DiskGroupSnap agent is used to perform the fire drill in campus clusters. You can bring online, take offline, and monitor the fire drill versions of disk groups for testing.
 You can configure this agent to either use the disks at the secondary site (Bronze configuration) or use the disks configured on a dummy site at the secondary site (Gold configuration). Symantec recommends that you use the Gold configuration for the DiskGroupSnap resource.
 See [“Limitations with DiskGroupSnap agent”](#) on page 57.
- Support for manual failover of service groups across system zones in campus clusters
 The AutoFailOver attribute controls the service group behavior in response to service group and system faults. For campus clusters, you can set the value of the AutoFailOver attribute as 2 to manually fail over the service group across the system zones that you defined in the SystemZones attribute.
 The manual failover functionality requires that you enable the HA/DR license and that the service group is a non-hybrid service group.

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

Support for secure communication in global clusters

In global clusters, VCS provides the option of making the following communications secure:

- Communication between the wide-area connectors
- Communication between the wide-area connectors and the Steward process

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

LLT supports NIC bonding

You can configure NIC bonds (aggregated interfaces) as private links under LLT. LLT treats each aggregated interface as a single link. So, you must configure these NICs that form the bond in such a way that the NICs are connected to the same switch or hub.

Note: If the NICs are connected to different switches or hubs, you must establish connection between the switches or hubs.

See the *Veritas Cluster Server Installation Guide* for instructions to configure private heartbeats that use aggregated interfaces.

I/O fencing supports iSCSI devices

You can now use iSCSI devices as coordinator disks for I/O fencing. However, I/O fencing supports iSCSI devices only when you use DMP disk policy. Make sure that the `/etc/vxfenmode` file has the disk policy set to DMP before you use iSCSI devices as coordinator disks.

For the latest information on supported hardware visit the following URL:

<http://entsupport.symantec.com/docs/283161>

See the *Veritas Cluster Server Installation Guide* for instructions to configure I/O fencing using DMP policy.

Changes to I/O fencing commands

Following are the changes to the I/O fencing commands:

- The `vxfcntlsthdw` command now supports `-d` option for dmp devices.
- The `vxflenconfig -l` command lists the coordinator disks that the `vxfen` driver uses.

Refer to the corresponding manual pages for more information on the commands.

Support to replace coordinator disks in a running cluster

You can now replace the coordinator disks in a running cluster using the `vxfsenwap` utility.

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

Enhanced encryption mechanism for agent passwords

The `vcencrypt` utility now supports stronger encryption mechanism for agent passwords. Use the `vcencrypt` utility to generate a security key that you can use to create a more secure password for the agent.

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

Support for Security-Enhanced Linux on Redhat Enterprise Linux 5

VCS is enhanced to run on Security-Enhanced Linux (SE Linux) in the enabled and the enforced modes. VCS supports SE Linux only when the security context is within the "unconfined_t" environment.

Support to add a node to a cluster running in secure mode

You can use the `installvcs` to add a node to a cluster running in secure mode.

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

Multi-tiered application support using the RemoteGroup agent for global groups

VCS supports the RemoteGroup agent when it points to a global group. The RemoteGroup agent must map the state of the global group in the local cluster.

See the *Veritas Cluster Server User's Guide* for more information on the functionality of this agent.

VCS documentation is available on the software disc

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.

Changes introduced in VCS 5.0 MP1

This section lists the changes introduced in the VCS 5.0 MP1 release.

Change in string size for some attribute values

For group name, resource name, attribute name, type name, and VCS username and password, the string size is limited to 1024 characters.

Support dropped for SANVolume agent

This release of VCS does not support the SANVolume agent that was shipped with VCS 5.0.

VCS FEN messages are now VxFEN messages

Error messages that are related to the fencing module, VCS FEN, are now read as VxFEN.

Campus cluster support

You can configure a campus cluster using functionality provided by Veritas Volume Manager.

To set up a campus cluster, make sure the disk group contains mirrored volumes. The mirrors must be on separate storage at different sites. Use site tags to distinguish between mirrors located at different sites. You could also use enclosure-based naming.

See the *Veritas Volume Manager Administrator's Guide* for detailed instructions.

Symantec recommends using I/O fencing in campus clusters.

Change in behavior: hstop command

VCS ignores the value of the cluster-level attribute EngineShutdown while the system is shutting down.

Change in behavior: BrokerIP attribute of the RemoteGroup agent

The BrokerIP attribute now requires only the IP address. Do not include the port number when you configure the attribute. [789878]

For a secure remote cluster only, if you need the RemoteGroup agent to communicate to a specific authentication broker, then set this attribute.

Type: string-scalar

Example: "128.11.245.51"

Fire drill support in Veritas Cluster Management Console

Veritas Cluster Management Console adds support for fire drills. The console lets you run fire drills and displays the status of the last fire drill.

- Viewing the status of the last fire drill—The service group listing tables display a column for the Physical Fire Drill Status, which indicates the results of the last fire drill.
- Running a fire drill.
 - Verify that replication for an application is working correctly.
 - Verify that a secondary disaster recovery (DR) application service group can be brought online successfully.
- Viewing fire drill logs—If a service group is configured with a physical fire drill group, a tab labelled Fire Drill Logs appears on the secondary tab bar in the Group:Summary view. Click this tab to view the VCS log messages about the fire drill group on the remote cluster and the resources that belong to it.

See the *Veritas Cluster Server User's Guide* for information about fire drills.

Viewing the status of the last fire drill

The column Fire Drill Status has been added to service group listing tables. A service group listing table is on the Cluster:Groups view.

For VCS global service groups that are configured with a fire drill group, this column indicates the results of the most recently run fire drill. The following are the possible states:

unknown	No fire drill has been run or the Cluster Management Console has come online after the most recent fire drill
running	Fire drill in progress
passed	Fire drill group came online on the secondary cluster
failed	Fire drill group did not come online on the secondary cluster

If multiple management servers are connected to the global cluster that contains the primary global group, the table does not show fire drill status for that group.

Running a fire drill

The Cluster Management Console supports fire drills in multi-cluster mode only. Before you run a fire drill, you must do the following:

- Configure the local (primary) and remote (secondary) global groups
- Set up the replication for the storage at the primary and secondary sites
- Configure the fire drill group using the FDSETUP command line wizard.

To run a fire drill from the Cluster Management Console

- 1 On the navigation bar, click **Home**.
- 2 On the secondary tab bar, click **Clusters**.
- 3 In the Home:Clusters view, in the Clusters Listing table, click the name of the primary global cluster.
- 4 On the secondary tab bar, click **Groups**.
- 5 In the Cluster:Groups view, in the Groups Listing table, click the name of the primary global group.
- 6 In the Group:Summary view, in the Remote Operations task panel, click **Run fire drill**.

You can view results of the fire drill in the Cluster:Groups view, the Group:Summary view, and in the Group:Fire Drill Logs view.

Viewing fire drill logs

Running a fire drill creates fire drill logs. If a service group is configured with a fire drill group, a tab labeled Fire Drill Logs appears on the secondary tab bar in the Group:Summary view.

To view fire drill logs

- 1 On the navigation bar, click **Home**.
- 2 On the secondary tab bar, click **Clusters**.
- 3 In the Home:Clusters view, in the Clusters Listing table, click the name of a VCS global cluster.

The global cluster must contain a global service group (primary group) that is configured with a fire drill group at a secondary location.

- 4 On the secondary tab bar, click **Groups**.

- 5 In the Cluster:Groups view, in the Groups Listing table, click the name of the primary global group.
- 6 In the Group:Summary view, on the secondary tab bar, click **Fire Drill Logs**.
This tab contains VCS log messages about the fire drill group on the remote (secondary) cluster and the resources that belong to it.

Features introduced in VCS 5.0

This section lists the features introduced in the VCS 5.0 release.

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

Cluster Management Console

The new Cluster Management Console replaces Cluster Manager (Web Console) and CommandCentral Availability.

Cluster Management Console enables administration and analysis for VCS clusters in your enterprise from a single console. You can install Cluster Management Console on a stand-alone system to manage multiple clusters or you can install the console on cluster nodes to manage a local cluster. When installed to manage a local cluster, the console is configured as part of the ClusterService group and the AppName attribute is set to cmc.

Cluster Monitor is now called Cluster Connector

CommandCentral Availability installed a component called Cluster Monitor on cluster nodes. The updated component is called Cluster Connector.

VCS privileges for operating system user groups

VCS 5.0 lets you assign VCS privileges to native users at an operating system (OS) user group level in secure clusters.

Assigning a VCS role to a user group assigns the same VCS privileges to all members of the user group, unless you specifically exclude individual users from those privileges.

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

Five levels of service group dependencies

VCS now supports configuring up to five levels of service group dependencies. The exception is the online local hard dependency, for which only two levels are supported.

New RemoteGroup agent to monitor service groups in remote clusters

The new RemoteGroup agent monitors and manages service groups in a remote cluster. See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the agent.

Enhancements to the hastop command

You can customize the behavior of the hastop command by configuring the new EngineShutdown attribute for the cluster.

[Table 1-2](#) lists the EngineShutdown attribute's values and description.

Table 1-2 EngineShutdown attribute values

Attribute value	Description
Enable	Process all hastop commands. This is the default behavior.
Disable	Reject all hastop commands.
DisableClusStop	Do not process the hastop -all command; process all other hastop commands.
PromptClusStop	Prompt for user confirmation before running the hastop -all command; process all other hastop commands.
PromptLocal	Prompt for user confirmation before running the hastop -local command; reject all other hastop commands.
PromptAlways	Prompt for user confirmation before running any hastop command.

Simulator supports deleting simulated clusters

VCS Simulator now supports deleting simulated clusters.

Symantec recommends using the same tool (command line or Java Console) to create and delete a cluster. For example, if you created the cluster from the Java Console, delete the cluster from the Java Console.

I/O fencing updates: DMP support

Dynamic multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. You can configure coordinator disks to use Veritas Volume Manager DMP feature.

You can set the coordinator disks to use either raw or DMP as the hardware path to a drive.

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

Minimal downtime upgrade to VCS 5.0

See the *Veritas Cluster Server Installation Guide* for a strategy on upgrading to VCS 5.0 while ensuring a minimal downtime for your applications.

Backup of VCS configuration files

VCS backs up all configuration files (*config.cf*) including *main.cf* and *types.cf* to *config.cf.autobackup*. The configuration is backed up only if the *BackupInterval* attribute is set and the configuration is writable.

When you save a configuration, VCS saves the running configuration to the actual configuration file (i.e. *config.cf*) and removes all autobackup files. This does away with the VCS behavior of creating stale files.

If you do not configure the *BackupInterval* attribute, VCS does not save the running configuration automatically.

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

Support for security services

VCS 5.0 uses the Symantec Product Authentication Service to provide secure communication between cluster nodes and clients, including the Java and the Web consoles. VCS uses digital certificates for authentication and uses SSL to encrypt communication over the public network.

Separate logger thread for HAD

The VCS engine, HAD, runs as a high-priority process to send heartbeats to kernel components and to respond quickly to failures. In VCS 5.0, HAD runs logging activities in a separate thread to reduce the performance impact on the engine due to logging.

Enhanced NFS lock failover

The new NFSRestart agent provides high availability to NFS locks. Use the agent in conjunction with the NFS agent.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information.

Support for VLAN interfaces

The NIC and MultiNICA agents now support VLAN interfaces. The agents do not configure the NICs, but can monitor them.

See the OS vendor's documentation on how to configure VLAN on your host, and ensure that the switch or router connected to such an interface is compatible with your configuration. Both server-side and switch-side VLAN configurations are supported.

Virtual fire drill

VCS supports a virtual fire drill capability that lets you test whether a resource can fail over to another node in the cluster. Virtual fire drills detect discrepancies between the VCS configuration and the underlying infrastructure on a node; discrepancies that might prevent a service group from going online on a specific node.

See the *Veritas Cluster Server User's Guide* for more information on running virtual fire drills.

New term: Daemon Down Node Alive (DDNA)

Daemon Down Node Alive (DDNA) is a condition in which the VCS high availability daemon (HAD) on a node fails, but the node is running. When HAD fails, the hashadow process tries to bring HAD up again. If the hashadow process succeeds in bringing HAD up, the system leaves the DDNA membership and joins the regular membership.

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

Change in behavior: Use comma or semicolon as delimiter

VCS 5.0 does not support using spaces as delimiters to separate vector, association, or keylist values. You must use a comma or a semicolon as a delimiter.

Change in behavior: New format for engine version

The new `EngineVersion` attribute replaces the `MajorVersion` and `MinorVersion` attributes. VCS stores version information in the following format:

major.minor.maintenance_patch_num.point_patch_num

For example:

5.0.30.0

Change in behavior for the resfault trigger

VCS now provides finer control over the resfault trigger. The resfault trigger is now invoked if the `TriggerResFault` attribute is set to 1.

Change in behavior: New location for enterprise agents

VCS enterprise agents are now installed in the `/opt/VRTSagents/ha/bin` directory.

The *agentTypes.cf* files are now located at `/etc/VRTSagents/ha/conf/agent`.

Change in behavior: New location of message catalogs and attribute pools

VCS stores binary message catalogs (BMCs) at the following location:

/opt/VRTS/messages/language/module_name

The variable *language* represents a two-letter abbreviation.

The attribute pools also move from `/var` to `/opt`.

Change in behavior: New option for the hastart and had commands

Use the `-v` option to retrieve concise information about the VCS version. Use the `-version` option to get verbose information.

Changes to bundled agents

VCS introduces the following new agents:

- `NFSRestart`—Provides high availability for NFS record locks.
- `RemoteGroup`—Monitors and manages a service group on another system.
- `SANVolume`—Monitors volumes in a SAN environment managed using Storage Foundation Volume Server.

- Apache (now bundled on all platforms)—Provides high availability to an Apache Web server.

See “[No longer supported](#)” on page 32.

Changes to licensing for VCS

[Table 1-3](#) describes the licensing scheme that VCS now follows.

Table 1-3 VCS licensing scheme

License	What's included
VCS	<p>This license includes the following:</p> <ul style="list-style-type: none"> ■ VCS ■ Cluster Management Console ■ Database agents ■ Application agents ■ Virtual fire drill support
VCS HA/DR	<p>This license includes the following:</p> <ul style="list-style-type: none"> ■ VCS ■ Cluster Management Console ■ Database agents ■ Application agents ■ Replication agents ■ Global clustering ■ Fire drill support

Note:

Database agents are included on the VCS 5.0 disc. The replication and application agents are available via the Veritas High Availability Agent Pack.

New attributes

VCS 5.0 introduces the following new attributes. See the *Veritas Cluster Server User's Guide* for more information.

Resource type attributes:

- AgentFile—Complete name and path of the binary for an agent. Use when the agent binaries are not installed at their default locations.

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

Cluster attributes:

- **EngineShutdown**—Provides finer control over the `hastop` command.
- **BackupInterval**—Time period in minutes after which VCS backs up configuration files.
- **OperatorGroups**—List of operating system user account groups that have Operator privileges on the cluster.
- **AdministratorGroups**—List of operating system user account groups that have administrative privileges on the cluster.
- **Guests**—List of users that have Guest privileges on the cluster.

System attributes:

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

Service group attributes:

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

Removed attributes

VCS 5.0 does not use the following attributes:

- **DiskHbStatus**—Deprecated. This release does not support disk heartbeats. Symantec recommends using I/O fencing.
- **MajorVersion**—The **EngineVersion** attribute provides information about the VCS version.
- **MinorVersion**—The **EngineVersion** attribute provides information about the VCS version.

Updates to the DB2 agent

The Veritas High Availability Agent for DB2 introduces the following changes:

- The attributes StartUpOpt and ShutDownOpt provide new start up and shut down options. Using the StartUpOpt attribute, you can start the instance or partition, activate database commands after processes start, or create customized start up sequences. Using the ShutDownOpt attribute, you can perform a normal stop or customize your shut down sequence.
- In previous releases when you enabled in-depth monitoring (IndepthMonitor=1), it executed a default SQL query. The in-depth monitor now allows you to classify actions for DB2 errors according to their severity. You can associate predefined actions with each error code with a monitoring script that you can customize. You can find a sample of in-depth monitoring script in the following directory:
`/etc/VRTSagents/ha/conf/Db2udb/sample_db2udb.`
You must install the custom script in the `/opt/VRTSagents/ha/bin/Db2udb` directory to enable indepth monitoring.
- You can enable the AgentDebug attribute to get more debugging information from the agent and the database.

Updates to the Oracle agent

The Veritas High Availability Agent for Oracle introduces the following changes:

- New monitoring option—The basic monitoring option of the Oracle agent now allows health check monitoring in addition to the process check monitoring. You can choose the health check monitoring option for Oracle 10g and later.
- Support for virtual fire drills—VCS requires you to keep the configurations in sync with the underlying infrastructure on a cluster node. Virtual fire drills detect such discrepancies that prevent a service group from going online on a specific system.

Refer to the *Veritas Cluster Server User's Guide* for more information.

The agent uses the Action entry point to support the virtual fire drill functionality.

Updates to the Sybase agent

The Veritas High Availability Agent for Sybase introduces the following changes:

- The agent supports Sybase ASE 12.5.x and 15.
- The agent supports encrypted passwords.

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 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 “[Supported hardware](#)” on page 29.

See “[Supported operating systems](#)” on page 29.

Veritas Installation Assessment Service

The Veritas Installation Assessment Service (VIAS) utility assists you in getting ready for a Veritas Storage Foundation and High Availability Solutions installation or upgrade. The VIAS utility allows the preinstallation evaluation of a configuration, to validate it prior to starting an installation or upgrade.

<https://vias.symantec.com/>

Supported hardware

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

<http://entsupport.symantec.com/docs/283161>

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

Supported operating systems

VCS operates on the Linux operating systems and kernels distributed by Red Hat and SUSE.

Table 1-4 lists the supported operating system versions for Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server (SLES). The table also lists the supported kernel versions and the architecture.

Table 1-4 Supported Linux operating system and kernel versions

Operating System	Kernel	Architecture
RHEL 5 Update 2 and later	2.6.18-92.el5	ppc64
SLES 10 with SP2	2.6.16.60-0.21-default	ppc64
	2.6.16.60-0.21-smp	

Note: If your system runs an older version of SUSE Linux Enterprise Server, you must upgrade the operating system before you attempt to install the VCS software. Refer to the SUSE documentation for more information on upgrading your system.

Symantec supports only Red Hat and SUSE 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 erratas and updates and SUSE service packs is available in the following TechNote. The TechNote also includes any updates to the supported operating systems and software. Read this TechNote before you install Symantec products.

<http://entsupport.symantec.com/docs/285834>

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 1-5](#) lists the RPMs that VCS requires for a given Linux operating system.

Table 1-5 Required RPMs

Operating system	Required RPMs
RHEL 5	glibc-2.5-34.ppc.rpm glibc-2.5-34.ppc64.rpm glibc-common-2.5-34.ppc.rpm libgcc-4.1.2-44.el5.ppc.rpm libgcc-4.1.2-44.el5.ppc64.rpm compat-libgcc-296-2.96-138.ppc.rpm libstdc++-4.1.2-44.el5.ppc.rpm libstdc++-4.1.2-44.el5.ppc64.rpm compat-libstdc++-296-2.96-138.ppc.rpm compat-libstdc++-33-3.2.3-61.ppc.rpm compat-libstdc++-33-3.2.3-61.ppc64.rpm java-1.4.2-gcj-compat-1.4.2.0-40jpp.115.ppc.rpm
SLES 10	glibc-2.4-31.54.ppc.rpm glibc-64bit-2.4-31.54.ppc.rpm compat-libstdc++-64bit-5.0.7-22.2.ppc.rpm compat-libstdc++-5.0.7-22.2.ppc.rpm compat-2006.1.25-11.2.ppc.rpm libgcc-4.1.2_20070115-0.21.ppc.rpm libgcc-64bit-4.1.2_20070115-0.21.ppc.rpm libstdc++-4.1.2_20070115-0.21.ppc.rpm libstdc++-64bit-4.1.2_20070115-0.21.ppc.rpm

Supported software

VCS supports the following volume managers and files systems:

- ext2, ext3, reiserfs, NFS, NFSv4, and bind on LVM2, Veritas Volume Manager (VxVM) 5.0, and raw disks.
- Veritas Volume Manager (VxVM) with Veritas File System (VxFS)
 - VxVM
 - VRTSvxvm-common-5.0.33.00-RU3_SLES10
 - VRTSvxvm-platform-5.0.33.00-RU3_SLES10

VRTSvxxvm-common-5.0.33.00-RU3_RHEL5
 VRTSvxxvm-platform-5.0.33.00-RU3_RHEL5

- VxFS
 - VRTSvxfs-common-5.0.33.00-RU3_SLES10
 - VRTSvxfs-platform-5.0.33.00-RU3_SLES10
 - VRTSvxfs-common-5.0.33.00-RU3_RHEL5
 - VRTSvxfs-platform-5.0.33.00-RU3_RHEL5

Supported VCS agents

Veritas agents support a specified application version on Linux for IBM Power if the application vendor supports that version on Linux for IBM Power.

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

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

Agent	Agent version	VCS version	Application	Application version	OS version
DB2	5.0 RU3	5.0 and later	DB2 Enterprise Server Edition	8.1, 8.2 9.1, 9.5	RHEL5, SLES 10
Oracle	5.0 RU3	5.0 and later	Oracle Database Enterprise Edition	9i 10g R1 10g R2	RHEL5, SLES10
Sybase	5.0 RU3	5.0 and later	Sybase Adaptive Server Enterprise	12.5.x, 15	RHEL5, SLES10

See the Installation and Configuration Guide for the agents for more details.

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

No longer supported

VCS no longer supports the following:

- CampusCluster agent
- Apache agent configuration wizard
- The updated Oracle agent does not support Oracle 8.0.x and Oracle 8.1.x.

- The updated DB2 Agent does not support DB2 7.2
- VCS documentation package (VRTSvcscd)
The VCS documentation package (VRTSvcscd) is deprecated. The software disc contains the documentation for VCS in Portable Document Format (PDF) in the *cluster_server/docs* directory.
Symantec recommends copying pertinent documents from the disc to your system directory */opt/VRTS/docs* for reference.

Known issues

The following issues are open for this release of VCS.

Operational issues for VCS

This section covers the operational issues for VCS.

Volumes outside of VCS control that are mount locked cannot be unmounted without specifying the key

If a VxFS file system has "mntlock=key" in its mount options, then you cannot unmount the file system without specifying the key. Groups having DiskGroup resources configured with UmountVolumes set, may fail to switch or failover if the volumes are mount locked. [1276594]

Saving large configuration results in very large file size for main.cf

If your service groups have a large number resources or resource dependencies, and if the PrintTree attribute is set to 1, saving the configuration may cause the configuration file to become excessively large in size and may impact performance. [616818]

Workaround: Disable printing of resource trees in regenerated configuration files by setting the PrintTree attribute to 0.

AutoStart may violate limits and prerequisites load policy

The load failover policy of Service Group Workload Management may be violated during AutoStart when all of the following conditions are met: [251660]

- More than one autostart group uses the same Prerequisites.
- One group, G2, is already online on a node outside of VCS control, and the other group, G1, is offline when VCS is started on the node.

- The offline group is probed before the online group is probed.

In this scenario, VCS may choose the node where group G2 is online as the AutoStart node for group G1 even though the Prerequisites load policy for group G1 is not satisfied on that node.

Workaround: Persistently freeze all groups that share the same Prerequisites before using the `hastop -force` command to stop the cluster or node where any such group is online. This workaround is not required if the cluster or node is stopped without the force option.

Trigger not invoked in REMOTE_BUILD state

In some situations, VCS does not invoke the injeopardy trigger if the system is a REMOTE_BUILD state. VCS fires the trigger when the system goes to the RUNNING state.

The hagetcf script reports an error

Running the hagetcf script to gather information about the VCS cluster generates the following error:

```
tar: cannot stat ./var/VRTSvcs/log/*.A.log. Not dumped.
```

Workaround: This message may be safely ignored.

Node cannot join cluster because port v is not ready for configuration

This behavior is observed when a node leaves a cluster and another node tries to join the cluster at the same time. If the GAB thread is stuck in another process, the new node cannot join the cluster and GAB logs the following warning:

```
GAB WARNING V-15-1-20126 Port v not ready  
for reconfiguration, will retry.
```

Using the coordinator attribute

This release contains an attribute for disk groups called coordinator, which configures disks as coordinator disks by the I/O fencing driver. Setting the attribute prevents the coordinator disks from being reassigned to other disk groups. See the Veritas Volume Manager documentation for additional information about the coordinator attribute.

The attribute requires that the disk group contain an odd number of disks. Symantec recommends that you use only three coordinator disks. Using more (five or seven) disks may result in different subclusters.

VCS controlled mount fails, while manual mount of volume succeeds

Security-enhanced Linux must be disabled, because the Security-enhanced (SE) Linux support is provided for evaluation purposes only and the Security policy files are not currently available for the Veritas product stack. Problems such as the mount issue in the subject title can result when Security-enhanced Linux is enabled.

Workaround: To disable SE Linux at boot time on both SLES 9 and RHEL 4, set the kernel boot parameter `selinux` to 0 (`selinux=0`) and reboot the machine. Assuming the system has been configured for booting from the machine `machine_name`, edit the file `/boot/machine_name/menu.lst` to include `selinux=0` on the kernel line. Then reboot the machine to ensure the setting takes effect.

Network interfaces change their names after reboot

On SUSE systems, network interfaces change their names after reboot even with `HOTPLUG_PCI_QUEUE_NIC_EVENTS=yes` and `MANDATORY_DEVICES="..."` set.

Workaround: Use `PERSISTENT_NAME= ethX` where `X` is the interface number for all interfaces.

Unloading DiskRes driver requires a reboot on RHEL 4

On systems running RHEL 4, you must reboot a system after if you are upgrading or replacing the DiskRes driver.

Unmount fails while taking service group offline

A known issue in Red Hat Enterprise Linux 4 could cause unmount to fail. When an NFS client does some heavy I/O, unmounting a resource in the NFS service group may fail while taking the service group offline. Refer to Red Hat's Bugzilla id 154387 for more information.

Some alert messages do not display correctly

The following alert messages do not display correctly [612268]:

- | | |
|-------|---|
| 51033 | Global group %s is unable to failover within cluster %s and AutoFailOver is %s. Administrative action is required. |
| 51032 | Parallel global group %s faulted on system %s and is unable to failover within cluster %s. However, group is still online/partial on one or more systems in the cluster |

51031	Unable to automatically fail over global group %s remotely because local cluster does not have Authority for the group.
51030	Unable to find a suitable remote failover target for global group %s. Administrative action is required
50916	Unable to automatically failover global group %s remotely due to inability to communicate with remote clusters. Please check WAN connection and state of wide area connector.
50914	Global group %s is unable to failover within cluster %s and ClusterFailOverPolicy is set to %s. Administrative action is required.
50913	Unable to automatically fail over global group %s remotely because clusters are disconnected and ClusterFailOverPolicy is set to %s. Administrative action is required.
50836	Remote cluster %s has faulted. Administrative action is required.
50761	Unable to automatically fail over global group %s remotely because ClusterList values for the group differ between the clusters. Administrative action is required.

Offlining one IP resource fails other IP resources on the same node

When you plumb a new IP address on any interface, the IP address is marked as secondary if another IP address exists on the interface in the same subnet. If you remove the primary IP address from the interface, all the secondary IP addresses in the subnet are also removed. So, if you offline one IP resource, the other IP resources on the same node fail. [1205382]

Workaround: Configure the base IP address in the same subnet as that of the primary IP address.

Issues related to the installer

This section covers the issues related to the installer.

Installer issue with aggregated interfaces

The product installer discovers all the aggregated interfaces but it cannot detect the speed of the interfaces. [1838370]

See *Veritas Cluster Server Installation Guide*.

Issues related to the VCS engine

This section covers the issues related to the VCS engine.

CPUUsage should be computed as 'long long' value

If System attribute CPUUsageMonitoring is set to monitor CPU Usage and the system is up and running for sufficiently long time, the CPUUsage is always computed as 100% even though there is no load on the system. [1365195]

Workaround: Freeze the system so that it is not considered as a target for failover. If there are groups online on this node, switch all the groups to another node in the cluster. Reboot the system. Unfreeze the system.

Resources in a parent service group may fail to come online if the AutoStart attribute for the resources is set to 0

This issue occurs for service groups linked with online local firm dependency, where the AutoStart attribute for all the resources of the parent service group is set to 0 (false) and at least one of these resources is up and running outside VCS control before VCS comes up. The AutoStart attribute of the parent service group itself does not matter.

If you take the resources of the parent service group offline individually and then switch or fail over the child service group to another node in the cluster, the child service group comes online on the node but the parent service group does not come online on that node. [1363506]

The following error is displayed in the VCS Engine logs for resources of the parent service group: "VCS WARNING V-16-1-10285 Cannot online: resource's group is frozen waiting for dependency to be satisfied"

Workaround: In such a scenario, while taking the parent service group resources offline, use the following command for the last resource:

```
hagrpr -offline service_group -sys system_name -clus cluster_name
```

Here, *service_group* is the name of the parent service group, *system_name* is the name of the system on which the service group is brought offline, and *cluster_name* is the name of the cluster to which the system belongs.

Engine may hang in LEAVING state

When the `hares -online` command is issued for a parent resource when a child resource faults, and the `hares -online` command is followed by the `hastop -local` command on the same node, then the engine transitions to the leaving state and hangs.

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

Timing issues with AutoStart policy

Consider a case where the service group is offline and engine is not running on node 1. If you restart the engine on node 1 after HAD is killed on node 2 and before the engine is restarted on node 2, then VCS does not initiate the autostart policy of the group.

Issues related to the VCS bundled agents

This section covers issues related to the VCS bundled agents.

Issue with the Share agent

The Share resource that was online may go into the OFFLINE state if you restart VCS after executing the `hastop -local -force` command. You must manually bring the Share resource online. [1845516]

RemoteGroup agent's monitor function may time out when remote system is down

If a RemoteGroup agent tries to connect to a system (specified as IpAddress) that is down, the monitor function of the RemoteGroup agent times out for the resource. [1397692]

Problem in failing over the IP resource

When a system panics, the IP address remains plumbed to the system for a while. In such a case, VCS may not succeed in failing over the IP resource to another system. This can be observed when a system panics during I/O Fencing.

Workaround: Increase the value of the OnlineRetryLimit attribute for the IP resource type.

DiskReservation might fault Mount resources

When the DiskReservation resource is brought online, the agent also does a `BLKRRPART ioctl` on the disks. This causes the block subsystem to see new block devices. Consequently, the OS launches the block hotplug agent to handle the events. The hotplug agent, as part of its work, unmounts any stale entries.[364315]

Because the hotplug agents are asynchronous, it is difficult to find whether all the hotplug agents have finished processing. So, the DiskReservation resource goes ONLINE even while the hotplug agents are running, which is fine with SCSI

reservation as the disks are reserved. However, when the DiskReservation resource goes ONLINE, a dependent Mount resource could also come up. And it is possible that the hotplug agent does its unmount after the Mount agent does its mount and a monitor cycle. If the Monitor entry point of the Mount resource is called after the unmount, VCS will never see the Mount resource as online. If the Monitor is called before the unmount, the resource goes ONLINE, and then in the next Monitor cycle, goes to FAULTED state.

Workaround: To avoid this, the DiskReservation agent is hard coded so that the Monitor entry point of the DiskReservation resource is called HOTPLUG_DELAY seconds after the Online entry point of the DiskReservation resource completes. HOTPLUG_DELAY is hard-coded to 5 seconds so that the first monitor happens 5 seconds after the DiskReservation resource is brought online.

If the hotplug agent cannot complete within the default HOTPLUG_DELAY time, set the OnlineRetryLimit and RestartLimit of the Mount type to 1.

NFS cannot handle minor number greater than 255

NFS cannot handle minor numbers greater than 255. [292216]

Workaround: Ensure that minor number of the VxVM diskgroup is not greater than 255.

NFS security feature

The NFS security feature does not work in a VCS environment. The NFSSecurity attribute is reserved for future use. [568498]

LVMVolumeGroup agent error on SE-Linux

If the LVMVolumeGroup agent is configured for SE-Linux, you may notice that the following error is logged in the audit log every monitor cycle: [1056433]

```
msg=audit(1189772065.053:232113): avc:  
denied { search } for pid=29652  
comm="vgdisplay" name="LVMVolumeGroup" ...
```

Workaround: Use the following command to stop getting these messages.

```
# setsebool -P vcs_lvmsupport true
```

Issues related to the I/O fencing for VCS

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

Fencing utilities require ssh or rsh configuration to the same node

To run the fencing utilities, `vxferswap` and `vxfersthdw`, ensure that you have `ssh/rsh` configured for password free logins even to the node on which these commands are executed. [1846607]

The `vxferclearpre` script displays error messages

The `vxferclearpre` script displays error messages if the `/etc/vxfentab` file is commented. [1512956]

The `/etc/vxfentab` file has the following comments:

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

When you run the `vxferclearpre` script, the following errors are displayed:

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

However, the `vxferclearpre` script operation is successful.

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

Preexisting split brain after rebooting nodes

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

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

Stopping vxfen when the fencing module is being configured

Trying to stop the `vxfen` driver when the fencing module is being configured results in the following error.

```
VXFEN vxfenconfig ERROR V-11-2-1013 Unable to unconfigure vxfen
VXFEN vxfenconfig ERROR V-11-2-1022 Active cluster is currently fencing.
```

Workaround: This message may be safely ignored.

The vxfen driver reports an error when stopping I/O fencing

If I/O fencing is started in disabled mode, and if you change the `/etc/vxfenmode` file later to `scsi3`, the `vxfen` driver throws the following spurious error when you stop `vxfen`:

```
ERROR: failed to stop vxfen
```

However, the actual I/O fencing stop operation is successful. [1301698]

Workaround: This error may be safely ignored.

I/O fencing configuration fails if fencing module is running on another node

The `vxfenconfig -c` command fails if any of the following commands are running on other nodes in the cluster:

```
/sbin/vxfenconfig -U
/sbin/vxfenconfig -c
```

```
vxfenconfig -U
vxfenconfig -c
```

Some vxfenadm options do not work with DMP paths

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

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

Issues related to global service groups

This section covers the issues related to global service groups.

Switch across clusters may cause concurrency violation

If you try to switch a global group across clusters while the group is in the process of switching across systems within the local cluster, then the group may go online on both the local and remote clusters. This issue affects only global groups. Local groups do not experience this behavior.

Workaround: Ensure that the group is not switching locally before attempting to switch the group remotely.

Using LVM in a GCO environment may cause concurrency violation

Logical Volume Manager (LVM) on Linux operating system is not supported with all replication technologies. Before using LVM with a VCS replication agent, read the documentation for the agent and the late breaking news for the Agent Pack.

<http://support.veritas.com/docs/282004>.

Global service group does not go online on AutoStart node

At cluster startup, if the last system where the global group is probed is not part of the group's AutoStartList, then the group does not AutoStart in the cluster. This issue affects only global groups. Local groups do not display this behavior.

Workaround: Ensure that the last system to join the cluster is a system in the group's AutoStartList.

Declare cluster dialog may not display highest priority cluster as failover target

When a global cluster fault occurs, the Declare Cluster dialog enables you to fail groups over to the local cluster. However, the local cluster may not be the cluster assigned highest priority in the cluster list.

Workaround: To bring a global group online on a remote cluster, do one of the following:

- From the Java Console, right-click the global group in the Cluster Explorer tree or Service Group View, and use the Remote Online operation to bring the group online on a remote cluster.
- From the Web Console, use the Operations links available on the Service Groups page to bring the global group online on a remote cluster.

Issues related to the VCS Agent for DB2

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

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

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

Log message when running DB2 agent

When the Veritas agent for DB2 is run on a VCS system with ja_JP.eucJP locale, VCS logs the following message:

```
Failed to open file /opt/VRTSvcs/messages/ja/HAD.bmcmmap.
```

Workaround: The agent is functioning correctly, and this message can be safely ignored.

Issues related to the VCS Agent for Oracle

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

NOFAILOVER action specified for certain Oracle errors

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

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

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

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

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

Health check may not work

If you set MonitorOption to 1, health check monitoring may not function when the following message is displayed [589934]:

```
Warning message - Output after executing Oracle Health  
Check is: GIM-00105: Shared memory region is corrupted.
```

Workaround: Set MonitorOption to 0 to continue monitoring the resource.

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

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

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

Issues related to the Cluster Manager (Java Console)

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

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

You can access the Template View in the Cluster Manager from the Tools > Templates menu. If you have Storage Foundation configured in a VCS cluster setup, the following error may occur while the Cluster Manager loads the templates. (1433844)

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

Workaround: Ignore the error.

VCS wizards fail on RHEL5

On RHEL5 VCS clusters, the VCS wizards fail to load because the JRE requirements are not met. To run the wizards, you must install the deprecated libxp packages. [974041]

Workaround: Install the following RPM.

```
# rpm -qlp xorg-x11-libXp-32bit-7.2-18.x86_64.rpm
```

```
/usr/lib/libXp.so.6  
/usr/lib/libXp.so.6.2.0
```

Cluster Manager (Java Console) hangs on Linux

The Cluster Monitor may hang while adding a new cluster panel and may fail to refresh the GUI. [527301]

Workaround: Kill the VCSGui process and restart the Cluster Manager (Java Console).

Exception when selecting preferences

On Windows systems, selecting the Java (Metal) look and feel of the Java Console may cause a Java exception. [585532]

Workaround: After customizing the look and feel, close restart the Java Console.

Java Console errors in a localized environment

When connected to cluster systems using locales other than English, the Java Console does not allow importing resource types or loading templates from localized directories. [585532]

Workaround: The workaround is to copy the types files or templates to directories with english names and then perform the operation.

Printing to file from the VCS Java Console throws exception

VCS Java Console and Help throw an exception while printing to a file from a system that does not have a printer configured. Also, the content is not written to the file.

Workaround: Before printing, make sure at least one printer is configured on the system where the VCS Java Console is launched.

Common system names in a global cluster setup

If both local and remote systems have a common system name in a global cluster setup, group operations cannot be performed on those systems using the Java console.

Workaround: Use command-line interface to perform group operations.

Some Cluster Manager features fail to work in a firewall setup

In certain environments with firewall configurations between the Cluster Manager and the VCS cluster, the Cluster Manager fails with the following error message [1392406]:

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

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

Issues related to the VCS Management Console

This section covers the issues that are related to the VCS Management Console.

Cluster Connector installation and configuration for fresh install of VCS 5.0 MP3 is not available

If you want to manage clusters that have VCS 5.0MP3, you must use the direct connection mode.

Upgrading from VCS 4.x (with CC5.0) leaves the CMC group frozen

If a VCS 4.x cluster running Cluster Connector 5.0 is upgraded to VCS 5.0 MP3, the CMC service group and other groups (except for ClusterService) are frozen and cannot come online after upgrading and rebooting the node. [1367670]

Workaround: The CMC group must be brought offline manually before upgrading from VCS 4.x (with CC 5.0) to VCS 5.0MP3.

Upgrade of the Cluster Connector can make the cluster state UNKNOWN on Management Server

If Cluster Connector is upgraded to 5.0 MP3, the cluster may show the state as "UNKNOWN" on the Management Server. This is due to missing attribute values. [1361602, 1202222]

Workaround: Execute the following commands to restore the CMC_ClusterConfig resource to the original state.

```
# haconf -makerw  
  
# hagrps -offline CMC -any  
  
# hares -modify CMC_ClusterConfig MSPort 14145
```

```
# hares -modify CMC_ClusterConfig ClusterType vcs  
# haconf -dump -makero  
# hagrps -online CMC -any
```

Restart the CMC group.

MSPort attribute of CMC_ClusterConfig is set to 0 after upgrading from VCS 4.x and 5.0 MP1 to 5.0 MP3

Cluster Connector cannot connect to Management Server after upgrading VCS from 4.x and 5.0 MP1 to 5.0 MP3. [1364500]

This is because the MSPort attribute of CMC_ClusterConfig has a value of 0, instead of 14145. This issue is not observed if VCS is upgraded from version 5.0 to 5.0 MP3 directly.

The HostMonitor group is displayed in the GUI when connected to CMC (or) Java GUI

The HostMonitor group (VCSHmg) is displayed in the GUI when connected to CMC (or) Java GUI on clusters running 5.0MP3. This group is not displayed to the user using the command line interface on cluster nodes. [1315640]

VRTScmc directory still exists in /opt/ after uninstalling VCS

After uninstalling VCS, the VRTScmc directory still exists in /opt/ even though the rpm is removed. This directory should be removed after the rpm is removed. [1315832]

The following directories are undeleted after uninstalling VCS.

```
/etc/VRTScmc  
/opt/VRTScmc/bin  
/opt/VRTScmc/conf  
/opt/VRTSweb
```

Login issue for CMC on a secure cluster

Entering any string in the Domain Name allows you to log into the Cluster Management Console for a cluster running in secure mode. [1208237]

Blank page is displayed on the cluster details page of the DR alert tab

A blank page is displayed on DR alert tab of the cluster details page after multiple DR alerts are generated. [1194391]

When the cluster has only one DR alert, the alerts page is displayed properly.

Warning messages in the log file when running the installmp command

The installmp command logs the following messages. [782798]

```
warning: user vcsbuild does not exist - using root
warning: group fcf does not exist - using root
warning: user vcsbuild does not exist - using root
warning: group fcf does not exist - using root
```

Workaround: None. You may ignore these messages.

Known issue for the Migrate Site task

The Migrate Site task starts the Migrate Site wizard that enables you to migrate one or more global service groups to one or more remote target clusters. The Cluster Management Console does not migrate a faulted service group. If you attempt to migrate a faulted service group, you may see an entry similar to the following in the management server log:

```
2006-11-20 10:38:33 INFO    Unable to use the -force option when
the cluster that has Authority for the group is not completely
down {vrts.vxcs.mcm.gui.web.actions.wizard.MigrateSiteLastPage
lookupFinish() }
```

Workaround: In the Global Application Group Selection panel, select only service groups that are in the online or partial state. Do not select service groups that are in the faulted state.

Erroneous output from gares command

The gares command returns a value for the Start attribute that is different from what the hares command returns. The local values are inverted (exchanged). For example, if gares returns 1, hares returns 0. [853969]

Workaround: This situation can result if the attribute values with local scope are missing for a newly-added system in the system list of a service group. Use the switch command for the CMC_CC service group (for configurations that use the

cluster connector) or reconnect to the cluster (for configurations that use direct connection).

Cluster Management Console displays fatal errors

CMC displays fatal errors when it encounters an invalid XML file for an agent. [595973]

Workaround: None. Make sure the XML files for custom agents are valid.

The database fails to back up or restore to a path with Japanese characters

The database fails to back up or restore to the specified path with Japanese characters in it, when the command `gadp -backup` is run. [767796]

Workaround: Use English folder names when backing up, then copy the database file to the Japanese folder manually, if required.

Cannot uninstall updates on Windows management server

On Windows, uninstalling the VCS 5.0 MP1 management server using Add or Remove Programs removes only the entry from the Add or Remove Programs list. No files are removed. You must perform a management server uninstallation using the original VCS 5.0 uninstallation program. You cannot revert a VCS 5.0 MP1 management server back to a VCS 5.0 management server. [841149]

View displays incorrect version

After upgrading to the Cluster Management Console for VCS 5.0 MP1, the Admin:Management Server view (Admin -> Management Server) shows an incorrect version of 5.0.1136.0 and an incorrect installation history. The correct information is in the About box. [856103]

Default SMTP and SNMP addresses in notification policies for Cluster Management Console

When you configure notification settings, the Edit SMTP Settings task asks you to provide default email or default SNMP console addresses. The policy configuration wizard uses these addresses only to populate the recipient lists during policy configuration. The wizard does not automatically configure policies with these addresses.

When you launch the Notification Policy Configuration wizard, the default email address you specified appears in the Notification Recipients dialog box.

If you add email addresses to this list, the wizard adds them to the policy along with the default address. However, if you delete all the addresses from the Email Recipients list, including the default email address, the wizard configures no email addresses in the policy.

Leave default email addresses in the recipients list to configure them into the policy.

The same behavior applies to specifying default SNMP addresses.

Console displays logs in English and Japanese

If your management server is configured to run in the Japanese locale, but the managed cluster does not have the Japanese language pack installed, the management server displays a mix of logs in English and Japanese. [778176]

Workaround: Make sure the managed cluster has the Japanese language pack installed.

Some Cluster Management Console controls not immediately active

In some versions of Internet Explorer, you may need to click Flash-based screens, popups, and wizards once before the controls become active. Controls that require this activating click show the following message when you roll over them with your mouse pointer [603415]:

```
Press SpaceBar or Click to activate this Control
```

Login screen may not display after inactivity timeout

If your Cluster Management Console is inactive and the session times out, your next action in the console should return you to the login screen. However, if your next action is to request a sort or a new page, the console will not sort the data or load the page.

Workaround: Use the browser refresh feature and the login screen will display.

Very large clusters may not load into Cluster Management Console

Very large clusters may not load into Cluster Management Console. [493844]

Workaround: To accommodate very large clusters, increase the value of the `loadClusterQueryTimeout` property in the management server configuration file, `/opt/VRTScmc/conf/ManagementServer.conf`. The management server generates this file upon startup.

To load large clusters into Cluster Management Console

- 1 Stop the Cluster Management Server web console:

```
/opt/VRTSweb/bin/stopApp cmc
```

- 2 Add the following line to the file `/opt/VRTScmc/conf/ManagementServer.conf`:

```
loadClusterQueryTimeout=60000
```

Adjust the value as needed to allow complete initial load of your cluster information.

- 3 Start the Cluster Management Server web console:

```
/opt/VRTSweb/bin/startApp cmc ../VERITAS
```

Log entries in the Management Server:Logs view

The Management Server:Logs view might contain log entries for the management server and for the cluster. [610333]

Management server log entries have the value site in the Object Type column. Cluster log entries have the value cluster in the Object Type column.

Cannot install if VxAT 4.3 is installed

If you have installed Symantec Product Authentication Services on a system using the 4.3 client/server installer, install of Cluster Management Console will not succeed because the path to the AT binaries is not in the path. Since this path is not present, the custom action DLL in our MSI will not be able to run certain AT-related commands. [617861]

Workaround: Add the path for the AT binaries before attempting a Cluster Management Console install.

Windows management server uninstall using Add or Remove Programs does not remove folder

After using Add or Remove Programs to remove (uninstall) the Windows management server, an empty Cluster Management Console folder remains:

The default path is `C:\Program Files\VERITAS`.

Workaround: Delete the empty folder after the uninstall.

Windows cluster monitor uninstall does not remove folder

After a Windows cluster monitor uninstall, an empty folder remains:

The default path is C:\Program Files\VERITAS.

Workaround: Delete the empty folder after the uninstall.

Uninstalling Cluster Connector does not remove entry from Add\Remove Programs on Windows

After you uninstall cluster connector on Windows cluster nodes, the Add or Remove Programs control panel continues to show an entry for cluster connector. This persistent entry prevents any reinstallation of cluster connector. [599424]

Workaround: Remove the Veritas Cluster Management Console entry from the list using Windows Installer Cleanup Utility. Run the utility to remove the entry on each node. If you do not have the utility, you may download it from the Microsoft support site.

Windows install over Terminal Services needs Service Pack 4

Per Microsoft, Windows 2000 without at least Service Pack 4 has problems installing multiple MSI files that alter the same registry key over Terminal Services.

Workaround: If you want to install to a Windows 2000 host using Terminal Services, first ensure that the system has Windows 2000 Service Pack 4 installed.

Removing the CMC_SERVICES domain

Uninstalling the management server in multi-cluster environments does not remove the *CMC_SERVICES* domain. [612176]

You can verify the existence of this domain using the following command:

```
vssat showpd --pdrtype ab --domain CMC_SERVICES
```

You must manually remove the *CMC_SERVICES* domain using the command line. To manually remove all the peripherals in the *CMC_SERVICES* domain, enter the following command:

```
vssat deletoprpl --pdrtype ab --domain  
CMC_SERVICES --prplname principalname
```

Enter the following command to remove the domain:

```
vssat deletepd --pdrtype ab --domain CMC_SERVICES@hostname
```

You can determine the host name using the following command:

```
vssat showpd
```

Issues related to VCS Simulator

This section covers the issues related to VCS Simulator.

Simulator clusters with Windows configurations fail to start on UNIX host platforms

The following clusters are affected: Win_Exch_2K3_primary, Win_Exch_2K3_secondary, Win_Exch_2K7_primary, Win_Exch_2K7_secondary, WIN_NTAP_EXCH_CL1, WIN_NTAP_EXCH_CL2, Win_SQL_Exch_SiteA, Win_SQL_Exch_SiteB, WIN_SQL_VVR_C1, WIN_SQL_VVR_C2. [1363167]

Workaround: For each of these clusters, there is a separate directory named after the cluster under the VCS Simulator installation directory

C:\Program Files\VERITAS\VCS Simulator on Windows

/opt/VRTScssim on Unix

Perform the following steps:

- Navigate to the conf/config directory under this cluster specific directory.
- Open the types.cf file in an editor and change all instances of the string "i18nstr" to "str".
- Open the SFWTypes.cf file in an editor if it exists in this directory and change all instances of the string "i18nstr" to "str".
- Repeat these steps for the following files if they exist: MSSearchTypes.cf, SQLServer2000Types.cf, ExchTypes.cf, SRDFTypes.cf.

Error in LVMVolumeNFSGroup template for AIX

In the VCS Simulator, the AIX_NFS cluster gives error while loading the LVMVolumeGroupNFS template. [1363967]

This problem can also affect real AIX clusters if they try to load this template.

Workaround: For the Simulator, navigate to the Templates/aix directory under the VCS Simulator installation directory (C:\Program Files\VERITAS\VCS Simulator on Windows, /opt/VRTScssim on Unix). Open the LVMVolumeNFSGroup.tf file and look for all instances of the MajorNumber = "". Remove the empty double-quotes and set the correct integer value for MajorNumber.

For real clusters, make identical changes to
`/etc/VRTSvcs/Templates/LVMVolumeNFSGroup.tf`.

Issues related to VCS in Japanese locales

This section covers the issues that apply to VCS 5.0 in a Japanese locale.

Installer does not create user account and password

The product installer does not ask for a VCS user account and password in a Japanese locale. Only the English installer provides this function.

Workaround: Use the `hauser` command to create VCS user accounts after installation is complete.

Symantec Web Server (VRTSWeb) requires restart after installing language packs

Cluster Management Console does not list Japanese as a language option after installing the language pack. [588560]

Workaround: Restart Symantec Web Server.

Error running CmdServer in Japanese eucjp locale

The command servers displays an unsupported encoding error when you run the Java Console in the Japanese `eucjp` locale. The error does not appear when you run the console in the Japanese UTF-8 locale. [533291]

Remote system log displays in English in Japanese locale

Log messages in the Java Console from remote systems display in English in the Japanese locale. [859457]

The perform check option in the virtual fire drill wizard does not work in Japanese locale

Running the `perform check` command in the virtual fire drill wizard in the Japanese locale results in the following error message [865446]:

```
No fire drill found for type <typename> of resource.
```

Workaround: Change the locale to English, when you run fire drill wizard with the `perform check` option.

Other known issues

This section covers other known issues.

Issue with LLT over bonded NICs

LLT does not work with Linux NIC bonded interfaces upon initial startup in VCS 5.0 RU3.

Workaround: Refer to the following technote for setting up LLT over bonded NICs.

<http://library.veritas.com/docs/308109>

Rebooting may fail due to disk reservation conflict

If a shared disk is reserved by one node, the other node may not boot up [315015]. During reboot, when the VxVM is coming up, the VM tries to read the private region data for all disks it can see, including the one reserved by the DiskReservation on some other node. The read goes through Dynamic Multipathing, which propagates only the EIO. Consequently, the rest of VM does not realize that a RESERVATION_CONFLICT has occurred, and retries the private region read. The rebooted node comes up only after the VM has done all the retries.

Workaround: Use vxdiskadm utility to remove all disks from VM that the DiskReservation uses. Reboot each node on which you remove the disks.

Software limitations

The following limitations apply to this release.

Limitation with the ASM disk group

During the time the ASMDG resource is online, if the permissions of the ASM disk groups are changed, the resource does not go into the FAULTED state.

DB2 RestartLimit value

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

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

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.

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]

Systems in a cluster must have same system locale setting

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

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

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

- Any group that you defined as VCSshm 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 VCSshm.

No support for NIC names larger than 8 characters

VCS does not support NIC names that are longer than eight characters. [605163]

GAB panics the systems while VCS gets diagnostic data

On receiving a SIGABRT signal from GAB, VCS engine forks off vcs_diag script. When VCS engine fails to heartbeat with GAB, often due to heavy load on the system, the vcs_diag script does a sys req to dump the stack trace of all processes in the system to collect diagnostic information. The dump of stack trace is intended to give useful information for finding out which processes puts heavy load. However, the dumping puts extra load on the system that causes GAB to panic the system in such heavy loads. See *Veritas Cluster Server User's Guide* for more information. [383970]

Workaround: Disable the `vcs_diag` script. To disable, rename the file `/opt/VRTSvcs/bin/vcs_diag` to `/opt/VRTSvcs/bin/vcs_diag.backup`.

Using agents in NIS

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can hang if the host is disconnected from the network. If such a program is used as an agent entry point, a network disconnect can cause the entry point to hang and possibly time out. For example, if the host is configured to use NIS maps as a client, basic commands such as `ps -ef` can hang if there is network disconnect. Symantec recommends creating users locally. To reflect local users, configure:

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

Fire drill does not support volume sets

The fire drill feature for testing fault readiness of a VCS configuration supports only regular Volume Manager volumes. Volume sets are not supported in this release.

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.

VxVM site for the diskgroup 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 diskgroup defined. The DiskGroupSnap agent invokes VCS action entry points to run VxVM commands to detach the site. These

commands must be run on the node where the diskgroup is imported, which is at the primary site.

If you attempt to shut down the node where the fire drill service group or the diskgroup is online, the node goes to a LEAVING state. The VCS engine attempts to take all the service groups offline on that node and rejects all action entry point requests. Therefore, the DiskGroupSnap agent cannot invoke the action to reattach the fire drill site to the target diskgroup. The agent logs a message that the node is in a leaving state and then removes the lock file. The agent's monitor function declares that the resource is offline. After the node restarts, the diskgroup site still remains detached. [1272012]

Workaround:

You must take the fire drill service group offline using the `hagrps -offline` command before you shut down the node or before you stop VCS locally.

If the node has restarted, you must manually reattach the fire drill site to the diskgroup that is imported at the primary site.

Manually removing VRTSat package erases user credentials

Symantec recommends saving user credentials before manually removing the VRTSat package. If you need the credentials again, you can restore them to their original locations.

To save user credentials

- 1 Run the `vssat showbackuplist` command. The command displays the data files and backs them up into the Snapshot directory `/var/VRTSatSnapshot`. Output resembles the following:

```
vssat showbackuplist
B| /var/VRTSat/.VRTSat/profile/VRTSatlocal.conf
B| /var/VRTSat/.VRTSat/profile/certstore
B| /var/VRTSat/RBAAuthSource
B| /var/VRTSat/ABAAuthSource
B| /etc/vx/vss/VRTSat.conf
Quiescing ...
Snapshot Directory :/var/VRTSatSnapshot
```

- 2 Move the credentials to a safe location. Preserving the directory structure makes restoring the files easier.

To restore user credentials

- 1 Navigate to the SnapShot directory or the safe location where you previously saved credentials:

```
cd /var/VRTSatSnapShot/
```

- 2 Restore the files:

```
cp ABAuthSource /var/VRTSat/  
cp RBAuthSource /var/VRTSat/  
cp VRTSat.conf /etc/vx/vss/  
cd /var/VRTSatSnapShot/  
cp -rp profile /var/VRTSat/.VRTSat/
```

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.

NFS locking

Due to RHEL 4 Update 2, update 3, and SLES 9 SP3 issues, lock recovery is not yet supported. Refer to issue 73985 for RHEL issues and bugzilla id 64901 for SLES 9 issues.

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
```

Security-Enhanced Linux is not supported on RHEL 4, SLES 9, and SLES 10

VCS does not support Security-Enhanced Linux (SELinux) on RHEL 4, SLES 9, and SLES 10. [1056433]

To disable SELinux at boot time on RHEL 4, SLES 9, or SLES 10 distributions, set the kernel boot parameter `selinux` to 0 (`selinux=0`) and reboot the machine.

For example, if the system is configured to boot from the grub, then do the following:

- Edit the file `/boot/grub/menu.lst` to include `selinux=0` on the kernel line.
- Reboot the machine to ensure the setting takes effect.

Bundled agent limitations

This section covers the software limitations for VCS 5.0 bundled agents.

NFS wizard limitation

The NFS wizard allows only one NFS service group to be created. You need to create additional groups manually.

Volume agent clean may forcibly stop volume resources

When the attribute `FaultOnMonitorTimeouts` calls the Volume agent clean entry point after a monitor time-out, the `vxvol -f stop` command is also issued. This command forcibly stops all volumes, even if they are still mounted.

NFS failover

This issue occurs on SLES 9 systems.

If the NFS share is exported to the world (*) and the NFS server fails over, NFS client displays the following error, "Permission denied".

Workaround: Upgrade `nfs-utils` to the package version "nfs-utils-1.0.6-103.28".

False concurrency violation when using PidFiles to monitor application resources

The PID files created by an application contain the PIDs for the processes that are monitored by Application agent. These files continue to exist even after a node running the application crashes. On restarting the node, the operating system may assign the PIDs listed in the PID files to other processes running on the node.

Thus, if the Application agent monitors the resource using the PidFiles attribute only, the agent may discover the processes running and report a false concurrency violation. This could result in some processes being killed that are not under VCS control.

Networking agents do not support IPv6 protocol

The bundled IP, NIC, IPMultiNIC, and MultiNICA agents for VCS 5.0 do not support the IPv6 enhanced IP protocol.

VCS does not provide a bundled agent for volume sets

5.0 RU3 does not provide a bundled agent to detect Volume Manager volume sets. Problems with volumes and volume sets can only be detected at the DiskGroup and Mount resource levels.

Workaround: Set StartVolumes and StopVolumes attributes of the DiskGroup resource that contains volume set to 1. If a file system is created on the volume set, use a Mount resource to mount the volume set.

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.

Cluster Management Console limitations

This section covers the software limitations for Cluster Management Console.

Cluster connector not supported on some OS versions

Cluster Management Console does not support cluster connector on AIX 5.1, Solaris 7, and RHEL 3.0. If your cluster runs on any of these platforms, you must use direct connection to manage the cluster from a management server.

Limited peer management server support

Peer management server support is limited to a configuration of two management servers in an enterprise. An enterprise of three or more management servers is not supported in this release.

Management server cannot coexist with GCM 3.5 Master

The Cluster Management Console management server should not be installed on the same system with a GCM 3.5 Master. These two products will conflict with each other and are not supported running on the same system.

Agent info files needed for Agent Inventory report

By design, the Agent Inventory report requires agent info files that supply the information reported on individual agents. These files are shipped with agents in VCS.

Global clusters must be CMC-managed clusters

All clusters forming a global cluster (using the VCS 4.0 Global Cluster Option) must be managed clusters in order for Veritas Cluster Management Console views to display correct and consistent information. Managed clusters are running the cluster connector or have a direct connection with the management server.

Windows Active Directory installation requires NetBIOS

If you install Cluster Management Console management server in a Windows Active Directory domain, NetBIOS must be turned on. A native (non-NetBIOS) Active Directory environment is not supported in this release.

Remote root broker not supported on Windows

If you set up a management server on a Windows system, you must configure a root broker on the management server system. This release does not support specifying a remote root broker during management server install [841739].

The root broker can be changed after install using the `configureRemoteRoot.exe` installed in `C:\Program Files\VERITAS\Cluster Management Console\bin` (default install directory).

Cluster Manager (Java console) limitations

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

Use the VCS 5.0 Java Console to manage clusters

Cluster Manager (Java Console) from previous VCS versions cannot be used to manage VCS 5.0 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 system 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 and wizards do not work if the hosts file contains IPv6 entries

VCS Cluster Manager and Wizards fail to connect to the VCS engine if the `/etc/hosts` file contains IPv6 entries.

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

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the `UseFence` attribute is set to the default, "None."

Undocumented commands, command options, and libraries

VCS contains undocumented commands and command options intended for development use only. Undocumented commands are not supported.

Documentation errata

Veritas Cluster Server User's Guide

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

Logging in to VCS non-interactively using the LDAP credentials

User's guide does not include the procedure to log in to VCS non-interactively using the LDAP credentials. [1382394]

If you do not have root privileges, VCS prompts for your VCS user name and password when you execute haxxx commands. You can use the halogin command to save the authentication information so that you do not have to enter your credentials every time you run a VCS command.

To store the LDAP credentials

- 1 Set the following environment variables:
 - VCS_HOST—Node where VCS is running to which the user must connect (can be a remote host)
 - VCS_DOMAIN—Name of the LDAP domain to which the user belongs
 - VCS_DOMAINTYPE—Type of VxSS domain: ldap

- 2 Log on to VCS using the halogin command:

```
# halogin vcsusername password
```

To run haxxx commands non-interactively using the LDAP credentials you stored

- 1 Set the VCS_HOST environment variable.
- 2 Unset the VCS_DOMAIN and VCS_DOMAINTYPE environment variables if these are already set. After unsetting if you run the haxxx commands, VCS does not prompt for the password.

Corrections for I/O fencing utilities and procedures

Replace the following topics in the *Veritas Cluster Server User's Guide* for 5.0 RU3 with the information in this Release Notes.

About I/O fencing utilities

The I/O fencing feature provides the following utilities that are available through the VRTSvxfen package:

vxfentsthdw	Tests hardware for I/O fencing
vxfenconfig	Configures and unconfigures I/O fencing Checks the list of coordinator disks used by the vxfen driver.
vxfenadm	Displays information on I/O fencing operations and manages SCSI-3 disk registrations and reservations for I/O fencing

vxfcntlpre	Removes SCSI-3 registrations and reservations from disks
vxfcntlswap	Replaces coordinator disks without stopping I/O fencing
vxfcntl	Generates the list of paths of disks in the diskgroup. This utility requires that Veritas Volume Manager is installed and configured.

Refer to the corresponding manual page for more information on the commands.

Removing and replacing a failed disk

If a disk in the coordinator disk group fails verification, remove the failed disk or LUN from the vxfcntlcoordg disk group, replace it with another, and retest the disk group.

To remove and replace a failed disk

- 1 Use the vxdiskadm utility to remove the failed disk from the disk group.
Refer to the *Veritas Volume Manager Administrator's Guide*.
- 2 Add a new disk to the node, initialize it, and add it to the coordinator disk group.
See the *Veritas Cluster Server Installation Guide* for instructions to initialize disks for I/O fencing and to set up coordinator disk groups.
If necessary, start the disk group.
See the *Veritas Volume Manager Administrator's Guide* for instructions to start the disk group.
- 3 Retest the disk group.

Removing preexisting keys

If you encountered a split brain condition, use the vxfcntlpre utility to remove SCSI-3 registrations and reservations on the coordinator disks as well as on the data disks in all shared disk groups.

You can also use this procedure to remove the registration and reservation keys created by another node from a disk.

To clear keys after split brain

- 1 Stop VCS on all nodes.

```
# hastop -all
```

- 2 Make sure that the port `h` is closed on all the nodes. Run the following command on each node to verify that the port `h` is closed:

```
# gabconfig -a
```

Port `h` must not appear in the output.

- 3 Stop I/O fencing on all nodes. Enter the following command on each node:
- 4 If you have any applications that run outside of VCS control that have access to the shared storage, then shut down all other nodes in the cluster that have access to the shared storage. This prevents data corruption.
- 5 Start the `vxfclearpre` script:

```
# cd /opt/VRTSvcs/vxfen/bin  
# ./vxfclearpre
```

- 6 Read the script's introduction and warning. Then, you can choose to let the script run.

```
Do you still want to continue: [y/n] (default : n) y
```

The script cleans up the disks and displays the following status messages.

```
Cleaning up the coordinator disks...
```

```
Cleaning up the data disks for all shared disk groups...
```

```
Successfully removed SCSI-3 persistent registration and  
reservations from the coordinator disks as well as the  
shared data disks.
```

```
Reboot the server to proceed with normal cluster startup...
```

```
#
```

- 7 Restart all nodes in the cluster.

Replacing I/O fencing coordinator disks when the cluster is online

Review the procedures to add, remove, or replace one or more coordinator disks in a cluster that is operational.

Warning: The cluster might panic if any node leaves the cluster membership before the `vxfsnwap` script replaces the set of coordinator disks.

To replace a disk in a coordinator diskgroup when the cluster is online

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# /sbin/vxfsnwap -d
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Import the coordinator disk group.

The file `/etc/vxfsnwap` includes the name of the disk group (typically, `vxfsnwapddg`) that contains the coordinator disks, so use the command:

```
# vxfsnwap -tfc import `cat /etc/vxfsnwap`
```

where:

-t specifies that the disk group is imported only until the node restarts.

-f specifies that the import is to be done forcibly, which is necessary if one or more disks is not accessible.

-C specifies that any import locks are removed.

- 4 Turn off the coordinator attribute value for the coordinator disk group.

```
# vxfsnwap -g vxfsnwapddg set coordinator=off
```

- 5 To remove disks from the coordinator disk group, use the VxVM disk administrator utility `vxfsnwapadm`.

- 6 Perform the following steps to add new disks to the coordinator disk group:

- Add new disks to the node.

- Initialize the new disks as VxVM disks.
- Check the disks for I/O fencing compliance.
- Add the new disks to the coordinator disk group and set the coordinator attribute value as "on" for the coordinator disk group.

See the *Veritas Cluster Server Installation Guide* for detailed instructions.

Note that though the disk group content changes, the I/O fencing remains in the same state.

- 7 Make sure that the `/etc/vxfenmode` file is updated to specify the correct disk policy.

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

- 8 From one node, start the `vxfenswap` utility. You must specify the diskgroup to the utility.

Do one of the following:

- If you use `ssh` for communication:

```
# /opt/VRTSvcs/vxfen/bin/vxfenswap -g diskgroup
```

- If you use `rsh` for communication:

```
# /opt/VRTSvcs/vxfen/bin/vxfenswap -g diskgroup -n
```

The utility performs the following tasks:

- Backs up the existing `/etc/vxfentab` file.
 - Creates a test file `/etc/vxfentab.test` for the diskgroup that is modified on each node.
 - Reads the diskgroup you specified in the `vxfenswap` command and adds the diskgroup to the `/etc/vxfentab.test` file on each node.
 - Verifies that the serial number of the new disks are identical on all the nodes. The script terminates if the check fails.
 - Verifies that the new disks can support I/O fencing on each node.
- 9 If the disk verification passes, the utility reports success and asks if you want to commit the new set of coordinator disks.

- 10 Review the message that the utility displays and confirm that you want to commit the new set of coordinator disks. Else skip to step 11.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfentab.test` file to the `/etc/vxfentab` file.

- 11 If you do not want to commit the new set of coordinator disks, answer `n`.
The `vxfenswap` utility rolls back the disk replacement operation.

Corrections for troubleshooting I/O fencing procedures

Replace the following topics in the *Veritas Cluster Server User's Guide* for 5.0 RU3 with the information in this Release Notes.

How `vxfen` driver checks for pre-existing split-brain condition

The `vxfen` driver functions to prevent an ejected node from rejoining the cluster after the failure of the private network links and before the private network links are repaired.

For example, suppose the cluster of system 1 and system 2 is functioning normally when the private network links are broken. Also suppose system 1 is the ejected system. When system 1 restarts before the private network links are restored, its membership configuration does not show system 2; however, when it attempts to register with the coordinator disks, it discovers system 2 is registered with them. Given this conflicting information about system 2, system 1 does not join the cluster and returns an error from `vxfenconfig` that resembles:

```
vxfenconfig: ERROR: There exists the potential for a preexisting  
split-brain. The coordinator disks list no nodes which are in  
the current membership. However, they also list nodes which are  
not in the current membership.
```

```
I/O Fencing Disabled!
```

Also, the following information is displayed on the console:

```
<date> <system name> vxfen: WARNING: Potentially a preexisting  
<date> <system name> split-brain.  
<date> <system name> Dropping out of cluster.  
<date> <system name> Refer to user documentation for steps  
<date> <system name> required to clear preexisting split-brain.  
<date> <system name>  
<date> <system name> I/O Fencing DISABLED!
```

```
<date> <system name>
<date> <system name> gab: GAB:20032: Port b closed
```

However, the same error can occur when the private network links are working and both systems go down, system 1 restarts, and system 2 fails to come back up. From the view of the cluster from system 1, system 2 may still have the registrations on the coordinator disks.

To resolve actual and apparent potential split brain conditions

- ◆ Depending on the split brain condition that you encountered, do the following:

- | | |
|---|--|
| Actual potential split brain condition—system 2 is up and system 1 is ejected | <ol style="list-style-type: none"> 1 Determine if system1 is up or not. 2 If system 1 is up and running, shut it down and repair the private network links to remove the split brain condition. 3 Restart system 1. |
|---|--|

- | | |
|---|---|
| Apparent potential split brain condition—system 2 is down and system 1 is ejected | <ol style="list-style-type: none"> 1 Physically verify that system 2 is down. Verify the systems currently registered with the coordinator disks. Use the following command: |
|---|---|

```
# vxfenadm -g all -f /etc/vxfentab
```

The output of this command identifies the keys registered with the coordinator disks.

- | |
|---|
| <ol style="list-style-type: none"> 2 Clear the keys on the coordinator disks as well as the data disks using the <code>vxfenclearpre</code> command. 3 Make any necessary repairs to system 2. 4 Restart system 2. |
|---|

Replacing defective disks when the cluster is offline

If the disk becomes defective or inoperable and you want to switch to a new diskgroup in a cluster that is offline, then perform the following procedure.

In a cluster that is online, you can replace the disks using the `vxfenswap` utility.

Review the following information to replace coordinator disk in the coordinator disk group, or to destroy a coordinator disk group.

Note the following about the procedure:

- When you add a disk, add the disk to the disk group `vxfencoorddg` and retest the group for support of SCSI-3 persistent reservations.

- You can destroy the coordinator disk group such that no registration keys remain on the disks. The disks can then be used elsewhere.

To replace a disk in the coordinator disk group when the cluster is offline

- 1 Log in as superuser on one of the cluster nodes.
- 2 If VCS is running, shut it down:

```
# hstop -all
```

Make sure that the port h is closed on all the nodes. Run the following command to verify that the port h is closed:

```
# gabconfig -a
```

- 3 Stop I/O fencing on each node:

This removes any registration keys on the disks.

- 4 Import the coordinator disk group. The file /etc/vxfendg includes the name of the disk group (typically, vxencoordg) that contains the coordinator disks, so use the command:

```
# vxdg -tfc import `cat /etc/vxfendg`
```

where:

-t specifies that the disk group is imported only until the node restarts.

-f specifies that the import is to be done forcibly, which is necessary if one or more disks is not accessible.

-C specifies that any import locks are removed.

- 5 To remove disks from the disk group, use the VxVM disk administrator utility, vxdiskadm.

You may also destroy the existing coordinator disk group. For example:

- Verify whether the coordinator attribute is set to on.

```
# vxdg list vxencoordg | grep flags: | grep coordinator
```

- If the coordinator attribute value is set to on, you must turn off this attribute for the coordinator disk group.

```
# vxdg -g vxencoordg set coordinator=off
```

- Destroy the disk group.

```
# vxdg destroy vxencoordg
```

- 6 Add the new disk to the node, initialize it as a VxVM disk, and add it to the vxencoorddg disk group.
- 7 Test the recreated disk group for SCSI-3 persistent reservations compliance.
- 8 After replacing disks in a coordinator disk group, deport the disk group:

```
# vxvg deport `cat /etc/vxfendg`
```
- 9 On each node, start the I/O fencing driver:
- 10 If necessary, restart VCS on each node:

```
# hastart
```

Incorrect I/O fencing tunable parameter name mentioned

The Veritas Cluster Server User's Guide mentions `vxfen_debug_sz` as one of the tunable parameters for I/O fencing. Replace `vxfen_debug_sz` with:

`dbg_log_size`

User's Guide does not mention about the OfflineWaitLimit attribute

The Appendix titled 'VCS Attributes' of the *Veritas Cluster Server User's Guide* does not mention about the OfflineWaitLimit attribute.

OfflineWaitLimit
(user-defined)

Note: This attribute can be overridden.

Description: Number of monitor intervals to wait for the resource to go offline after completing the offline procedure. Increase the value of this attribute if the resource is likely to take a longer time to go offline.

Type and dimension: integer-scalar

Default: 0

Veritas Cluster Server Bundled Agents Reference Guide

This section covers the additions or corrections to the *Veritas Cluster Server Bundled Agents Reference Guide* for 5.0 RU3.

ProcessOnOnly agent

Replace the agent's description with the following text:

The ProcessOnOnly agent starts and monitors a process that you specify. You can use the agent to make a process highly available or to monitor it. This resource's Operation value is OnOnly. VCS uses this agent internally to mount security processes in a secure cluster.

Replace the text under the Dependency heading with the following text:

No child dependencies exist for this resource.

Remove or ignore figure 5-4 under the Dependency header, it is incorrect.

Manual pages errata

The man page for `haremajor(1m)` includes AIX related information. Refer to the man page in the release notes for correct information on Linux.

haremajor

`haremajor` - Change the major numbers for disk partitions or volumes.

SYNOPSIS

```
haremajor -sd major_number  
haremajor -vx major_number_vxio major_number_vxspec  
haremajor -atf major-number
```

AVAILABILITY

VRTSvcs

DESCRIPTION

The `haremajor` command can be used to reassign major numbers of block devices such as disk partitions or Veritas Volume Manager volumes. NFS clients know the major and minor numbers of the block device containing the file system exported on the NFS server. Therefore, when making the NFS server highly available, it is important to make sure that all nodes in the cluster that can act as NFS servers have the same major and minor numbers for the block device. The `haremajor` command can be used to change major numbers used by a system when necessary. Use the `-sd` option to reassign a major number for a disk partition managed by the SD driver. Minor numbers will automatically be the same on both systems. Use the `-vx` option to reassign the major number for the Volume Manager volume device driver (`vxio`) as well as the `vxspec` device used by the Volume Manager process `vxconfigd`. Currently assigned major numbers can be determined by entering the command:

```
grep '^vx' /etc/name_to_major
```

Note that minor numbers for volumes can be changed using the `remisor` option of the `vxdg` command; see the manual page for `vxdg(1M)`. Use the `-atf` option to reassign the major number for the ATF (Application Transparent Failover) driver on a system.

OPTIONS

- `-sd majornumber`
Reassign the major number used by the SD driver on the system to *major_number*.
- `-vx major_number_vxio`
Reassign the major numbers used by the Volume Manager device drivers.
- `-atf major_number`
Reassign a major number used by ATF driver to *major_number*.

SEE ALSO

`vxdg(1M)`

COPYRIGHTS

Copyright © 2008 Symantec.

All rights reserved.

VCS documentation

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

You can access the VCS 5.0 MP3 documentation online at the following URL:

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

VCS documentation set

[Table 1-7](#) lists the documents that VCS includes.

Table 1-7 VCS documentation

Title	File name
<i>Veritas Cluster Server Installation Guide</i>	<code>vcs_install.pdf</code>
<i>Veritas Cluster Server Release Notes</i>	<code>vcs_notes.pdf</code>
<i>Veritas Cluster Server User's Guide</i>	<code>vcs_users.pdf</code>

Table 1-7 VCS documentation (*continued*)

Title	File name
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf
<i>Veritas Cluster Server Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_install.pdf
<i>Veritas Cluster Server Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_install.pdf
<i>Veritas Cluster Server Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_install.pdf

Table 1-8 lists the documentation for the VCS component - Symantec Product Authentication Service.

Table 1-8 Documentation for VCS components

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

VCS manual pages

The manual pages for VCS packages are installed in /opt/VRTS/man. 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
```

Documentation feedback

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