

Veritas™ Cluster Server 6.0.1 Release Notes - HP-UX

6.0.1

Veritas™ Cluster Server Release Notes

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

Product version: 6.0.1

Document version: 6.0.1 Rev 4

Legal Notice

Copyright © 2014 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

Documentation

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

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

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Veritas Cluster Server Release Notes

This document includes the following topics:

- [About this document](#)
- [Component product release notes](#)
- [About Veritas Cluster Server](#)
- [About Symantec Operations Readiness Tools](#)
- [Important release information](#)
- [Changes introduced in VCS 6.0.1](#)
- [VCS system requirements](#)
- [No longer supported](#)
- [VCS: Issues fixed in 6.0.1](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation](#)

About this document

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

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

This is "Document version: 6.0.1 Rev 4" of the *Veritas Cluster Server Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

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

Component product release notes

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

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

/docs/product_name

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

This release includes the following component product release notes:

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

About Veritas Cluster Server

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

About VCS agents

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

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

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

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

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

About compiling custom agents

Custom agents must be developed using compilers from one of the products listed below:

- HP C/ANSI C Developer's Bundle (S800), part number B3901BA.
- HP aC++ Compiler (S800), part number B3913DB.

These products may be identified by various part numbers in addition to those listed.

Existing custom agents written to run on VCS versions earlier than 1.2 must be recompiled for use with VCS .

About Symantec Operations Readiness Tools

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

SORT can help you do the following:

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

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

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

To access SORT, go to:

<https://sort.symantec.com>

Important release information

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

Changes introduced in VCS 6.0.1

This section lists the changes in Veritas Cluster Server 6.0.1.

New versioning process for SFHA Solutions products

Symantec made some changes to simplify the versioning process to ensure that customers have a unified experience when it comes to deploying our different products across Storage, Availability, Backup, Archiving and Enterprise Security products. With this change, all the products will have a 3 digit version. In complying with this approach, the current SFHA Solutions release is available as version 6.0.1.

New directory location for the documentation on the software media

The PDF files of the product documentation are now located in the `/docs` directory on the software media. Within the `/docs` directory are subdirectories for each of the bundled products, which contain the documentation specific to that product. The `sfha_solutions` directory contains documentation that applies to all products.

Changes related to installation and upgrades

The product installer includes the following changes in 6.0.1.

Locally-installed installation and uninstallation scripts now include the release version

When you run local scripts (`/opt/VRTS/install`) to configure Veritas products, the names of the installed scripts now include the release version.

Note: If you install your Veritas product from the install media, continue to run the `installvcs` command without including the release version.

To run the script from the installed binaries, run the `installvcs<version>` command.

Where `<version>` is the current release version with no periods or spaces.

For example, to configure the 6.0.1 version of your product, run this command:

```
# /opt/VRTS/install/installvcs601 -configure
```

Additional installation postcheck options

The `postcheck` option has been enhanced to include additional checks.

You can use the installer's post-check option to perform the following checks:

- General checks for all products.
- Checks for Volume Manager (VM).
- Checks for File System (FS).
- Checks for Cluster File System (CFS).

Support for tunables file templates

You can use the installer to create a tunables file template. If you start the installer with the `-tunables` option, you see a list of all supported tunables, and the location of the tunables file template.

Installer support to configure Coordination Point servers

You can now use the `-configcps` option in the installer to configure CP servers. This functionality to configure CP servers is now integrated with the installer. The `configure_cps.pl` script used earlier to configure CP servers is now deprecated.

You can also configure CP servers by generating response files. You can use the `-responsefile '/tmp/sample1.res'` option in the installer to configure CP servers.

See the *Installation Guide* for more details.

Attributes introduced in VCS 6.0.1

The following section describe the attributes introduced in VCS 6.0.1.

Service group attribute

- `UserAssoc`: This attribute can be used for any purpose.

Cluster level attribute

- `FipsMode`: Indicates whether FIPS mode is enabled for the cluster. The value depends on the mode of the broker on the system.

Changes to VCS bundled agents

This section describes changes to the bundled agents for VCS.

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

Changes to Share agent

On HP-UX, VCS 6.0.1 requires at least one entry in `/etc/dfs/dfstab` file on a node in order to configure Share agent.

Use the following command to share a directory across reboots:

```
#share /xyz
```

Enhancement to the CoordPoint agent

The CoordPoint agent monitors changes to the Coordinator Disk Group constitution, such as when a disk is deleted from the Coordinator Disk Group due to accidental execution of a VxVM administrative command or if the VxVM private region of a disk is corrupted.

The agent performs detailed monitoring on the CoordPoint resource and reports faults. You can tune the frequency of the detailed monitoring by setting the `LevelTwoMonitorFreq` attribute introduced in this release. For example, if you set this attribute to 5, the agent monitors the Coordinator Disk Group constitution in every fifth monitor cycle.

For more information on the CoordPoint agent, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

For information on configuring the CoordPoint agent using script-based installer and manually configuring the CoordPoint agent to monitor coordinator disks, see the *Veritas Cluster Server Installation Guide*.

For more information on replacing I/O fencing coordinator disks or coordinator diskgroup when the cluster is online, see the *Veritas Cluster Server Administrator's Guide*.

Application agent enhancements

Application agent has undergone the following enhancement:

- The ProPCV support for Application agent to prevent the processes configured in the `MonitorProcesses` attribute is supported since earlier release. If the ProPCV attribute is enabled for a service group, the processes configured under `MonitorProcesses` attribute for the Application resource in the group are prevented from starting on other cluster nodes once the resource is online on a cluster node.

With this release the ProPCV feature is enhanced for Application agent. With this enhancement, the Application agent now supports the prevention of the program configured under `StartProgram` attribute. When the ProPCV is enabled, the `StartProgram` will be prevented from starting on other cluster nodes once the Application resource is online on a node.

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

Changes related to IMF

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

Open IMF architecture

The Open IMF architecture builds further upon the IMF functionality by enabling you to get notifications about events that occur in user space. The architecture uses an IMF daemon (IMFD) that collects notifications from the user space notification providers (USNPs) and passes the notifications to the AMF driver, which in turn passes these on to the appropriate agent. IMFD starts on the first registration with AMF by an agent that requires Open IMF.

The Open IMF architecture provides the following benefits:

- IMF can group events of different types under the same VCS resource and is the central notification provider for kernel space events and user space events.
- More agents can become IMF-aware by leveraging the notifications that are available only from user space.
- Agents can get notifications from IMF without having to interact with USNPs.

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

New IMF-aware agent in VCS

The following agent is IMF-aware in VCS :

- DiskGroup agent

Changes to the VCS engine

Enhanced -propagate functionality to support more dependency types

The `-propagate` option can be used if the dependency tree contains global and/or remote dependency. The following dependency types are supported for both online propagate and offline propagate options:

- online global soft
- online global firm
- online remote soft
- online remote firm

Cluster security with FIPS mode

VCS provides an option to secure your cluster with FIPS. With this option, the communication with the cluster is encrypted using FIPS approved algorithms. The FIPS compliance is introduced with the following guiding factors:

- FIPS compliance is a configurable option available with VCS 6.0.1. When existing VCS deployments are upgraded from VCS 6.0 or earlier versions to 6.0.1, FIPS compliance is not automatically enabled.
- To enable FIPS mode, you must ensure that the cluster is new and configured without setting any security condition. To configure FIPS mode on a cluster which is already secured, refer to the steps under *Enabling and disabling secure mode for the cluster* in *Veritas Cluster Server Administrator Guide*.
- 6.0.1 does not support FIPS in GCO or CP server based cluster.

Postonline and postoffline triggers must be enabled after a manual upgrade

The preonline and postoffline triggers must be enabled if you perform a manual upgrade from VCS versions 5.x to 6.0 or later. You can enable the triggers if required by setting the TriggersEnabled attribute of the service group.

PreOnline, TriggersEnabled and ContainerInfo have a global (cluster-wide) value

The service group attributes PreOnline, TriggersEnabled and ContainerInfo have a global (cluster-wide) value. The value can be localized for every system.

Changes to LLT

This release includes the following change to LLT:

Setting the value of peerinact in the `/etc/llttab` file

Symantec recommends not to set the value of peerinact to 0. To achieve the infinite timeout functionality for peerinact, you must set peerinact to a large value. The supported range of value is between 1 through 2147483647.

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. All cluster nodes must be at the same patch level.

See [“Hardware compatibility list”](#) on page 16.

Hardware compatibility list

VCS supports vpar/npar from release VCS 5.0.1 and later.

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

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

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

Supported HP-UX operating systems

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

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

Table 1-1 Supported operating systems

Operating system	Operating system version	Architecture
HP-UX 11i Version 3 March 2011 Operating Environments Update Release or later	HP-UX B.11.31.1103	PA-RISC
	HP-UX B.11.31.1109	Itanium
	HP-UX B.11.31.1203	

Supported software for VCS

VCS supports the following volume managers and file systems:

- Logical Volume Manager (LVM)
- HP File System (HFS)

VCS supports the following versions of Veritas Storage Foundation:

Veritas Storage Foundation: Veritas Volume Manager (VxVM) with Veritas File System (VxFS)

- Storage Foundation 6.0.1
 - VxVM 6.0.1 with VxFS 6.0.1
- Storage Foundation 6.0
 - VxVM 6.0 with VxFS 6.0

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

Supported enterprise agents

[Table 1-2](#) lists the agents for enterprise applications and the software that the agents support.

Table 1-2 Supported software for the VCS agents for enterprise applications

Agent	Application	Application version	HP-UX version
DB2	DB2 Enterprise Server Edition	9.1, 9.5, 9.7	HP-UX 11i v3
Oracle	Oracle	10gR2, 11gR1, 11gR2	HP-UX 11i v3
Sybase	Sybase Adaptive Server Enterprise	12.5.x, 15.x	HP-UX 11i v3

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

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

No longer supported

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

No longer supported agents and components

VCS no longer supports the following:

- The `configure_cps.pl` script used to configure CP server is now deprecated and is no longer supported.

VCS: Issues fixed in 6.0.1

This section covers the incidents that are fixed in VCS 6.0.1.

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

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

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

Incident	Description
2845244	<p><code>vxfen</code> startup script gives error <code>grep: can't open /etc/vxfen.d/data/cp_uid_db</code>.</p> <p>The error comes because <code>vxfen</code> startup script tries to read a file that might not be present. This error is typically seen when starting <code>vxfen</code> for the very first time after installation.</p>
2554167	Setting <code>peerinact</code> value to 0 in the <code>/etc/llttab</code> file floods the system log file with large number of log messages.

Bundled agents fixed issues in 6.0.1

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

Table 1-4 Bundled agents fixed issues

Incident	Description
2850904	Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and <code>PanicSystemOnDGLoss</code> is set to 0.
2728802	If the <code>'httpd'</code> binary or the <code>'ab'</code> binary is not present at the location that you specified in the <code>'httpdDir'</code> attribute, the Apache agent cannot perform detail monitoring or start the HTTP server.
2850905	IMF registration for Mount resource for file systems type other than VxFS and NFS should be blocked.
2850916	Mount resource does not get registered with IMF if the attributes <code>BlockDevice</code> and/or <code>MountPoint</code> have a trailing slash in their values.
2850917	LVMLogicalVolume agent does not support zero byte logical volume and the resource goes into a faulted state.
2822920	DNSAgent goes to UNKNOWN state if the Top Level Domain (TLD) is more than 4 characters in length.

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

Incident	Description
2846389	In releases prior to VCS 6.0.1, the upper bound value of FaultTolerance attribute of the CoordPoint agent was the one less than the number of coordination points. If the majority number of coordination points fault, the entire cluster panicked under network partition scenario. Therefore, the upper bound value of the FaultTolerance attribute of CoordPoint agent had to be set to less than the majority of the coordination points. Subsequent to VCS 6.0.1, the FaultTolerance attribute of CoordPoint agent is less than the majority of coordination points.

VCS engine fixed issues in 6.0.1

[Table 1-5](#) lists the fixed issues for VCS engine.

Table 1-5 VCS engine fixed issues

Incident	Description
2879413	You may see two instances of CmdServer running on a node. One of these using IPv4 and the other IPv6.
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility <code>hagrps</code> gives incorrect output with the "-clear", "-flush", "-state" options.
2741299	CmdSlave gets stuck in a tight loop when it gets an EBADF on a file descriptor(fd). The CmdSlave process keeps retrying on the FD and eventually dumps core.
2850906	If a group is auto-enabled, the engine clears the Start attribute even if the resource is online.
2692173	Engine does not check whether remote parent is online when <code>-nopre</code> option is selected.
2684818	If the following attributes are specified before SystemList attribute in <code>main.cf</code> , then the value got rejected when HAD started: <ul style="list-style-type: none"> ■ PreOnline ■ ContainerInfo ■ TriggersEnabled ■ SystemZones
2696056	Memory leak occurs in the engine when <code>haclus -status <cluster></code> command is run.

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

Incident	Description
2746802	When failover group is probed, VCS engine clears the MigrateQ and TargetCount.
2746816	The syslog call used in gab_heartbeat_alarm_handler and gabsim_heartbeat_alarm_handler functions is not async signal safe.

Enterprise agents fixed issues in 6.0.1

[Table 1-6](#) lists the fixed issues for enterprise agents.

Table 1-6 Enterprise agents fixed issues

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

Agent framework fixed issues in 6.0.1

[Table 1-7](#) lists the fixed issues for agent framework.

Table 1-7 Agent framework fixed issues

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

Veritas Cluster Server fixed issues in 6.0 RP1

[Table 1-8](#) lists the issues fixed in this release.

Table 1-8 Veritas Cluster Server fixed issues

Incident	Description
2646789	Fault propagation does not work if the parent is in faulted state on one or more nodes.

Table 1-8 Veritas Cluster Server fixed issues (*continued*)

Incident	Description
2636874	AMF calls VXFS API while holding a spin lock and causes panic.
2646793	The following message appears as an error message: VCS ERROR V-16-25-50036 The child service group came online (recovered) before the parent was offlined.
2653701	The HAD process fails to maintain heartbeat with GAB due to a deadlock issue. Once the heartbeat timeout is elapsed, GAB terminates the HAD process by sending a SIGABRT signal.
2684818	If a pure local attribute is specified before SystemList in main.cf then it gets rejected when HAD is started.

Known issues

This section covers the known issues in this release.

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

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

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

Recovery and rollback to original configuration may not succeed if the system reboots while the online migration setup is in partial state (2611423)

During online migration from LVM to VxVM volumes, if there is a system reboot when the migration setup is in partial state, that is, the start operation has not completed successfully, then the recover and abort operations might not be able to recover and rollback the configuration.

Workaround: This needs manual intervention for cleanup, depending on the state, to restore the original configuration.

During online migration from LVM to VxVM volumes, LVM sometimes incorrectly reports the remapped LVM device paths as valid LVM volumes

Problem: In a migrated or committed configuration, only the renamed LVM names of the form `<lvolname>_yxlx` are valid LVM volumes. The original LVM names, in turn, point to target VxVM volumes. However, LVM sometimes incorrectly reports these original LVM device paths pointing to VxVM volumes, as valid LVM volumes.

Do not assume these as LVM volumes or do any operations on them, as it would disrupt the application's access to the target VxVM volumes.

Issues related to installing and upgrading VCS

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

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

Workaround:

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

To unfreeze the service groups manually

- 1 List all the frozen service groups

```
# hagrpl -list Frozen=1
```

- 2 Unfreeze all the frozen service groups:

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

Manual upgrade of VRTSvlic depot loses keyless product levels [2737124]

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

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

- ```
vxkeyless display
```
2. Set the product level to NONE.
 

```
vxkeyless set NONE
```
  3. Upgrade the VRTSvlic depot
 

```
swremove VRTSvlic
```

This step may report a dependency, which can be safely overridden.

```
swinstall -s 'pwd' VRTSvlic
```
  4. Restore the list of products that you noted in step 1.
 

```
vxkeyless set product[,product]
```

## Errors recorded in the swremove logs of VRTSgab during VCS upgrade from 4.1 to 5.0.1 (1719136)

When VCS is upgraded from 4.1 to 5.0.1 on HP-UX 11i v3 using the Veritas product installer, the installer reports errors for GAB and errors are recorded in the swremove logs related to VRTSgab.

**Workaround:** You can safely ignore these error messages.

## VCS agents dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the update-ux command (1630968)

On PA-RISC architecture, the VCS agents (Oracle, Netlsnr, Sybase, SybaseBk, MultiNICB, and so on) may dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the `update-ux` command.

This is because on HP-UX PA-RISC systems, the default thread stack size is limited to 64k. When the agent requires more than 64k stack memory, it may dump core due to SIGBUS error.

**Workaround:** Before running the `update-ux` command, edit the `/opt/VRTSvcs/bin/vcsenv` file to append following lines to it:

```
PLATFORM=`uname -s`
ARCHITECTURE=`uname -m`
if ["${PLATFORM}" = "HP-UX"] && ["${ARCHITECTURE}" = "9000/800"]; then
 PTHREAD_DEFAULT_STACK_SIZE=524288
 export PTHREAD_DEFAULT_STACK_SIZE
fi
```

## Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

After upgrading from 5.0.1 to higher versions of VCS, some keyless licenses may be left in the system. As a result, you may see periodic reminders being logged if the VOM server is not configured.

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

To resolve this issue, perform the following steps:

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

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

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

---

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

---

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

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

## Installer does not detect the duplicate cluster ID in an already configured SF Oracle RAC cluster [2368898]

When you run the installer using `installsfrac -configure` command and if you choose to check the cluster ID, the installer correctly checks if the cluster ID is in use by any other setup. However, if you perform the same check on an already configured SF Oracle RAC cluster, it is unable to detect it.

**Workaround:** No workaround.

## VxSS may go to a faulted state After stack and OS upgrade [2564568]

During the upgrade using the installer, if cluster is in secure mode and the upgrade is from 1123 to 1131, the installer may send a the following warning message.

```
Warning: /opt/VRTSat/bin/vxatd is not running
on <system name>. Will be unable to setup trust with shared
broker, however secure upgrade can still proceed.
```

Workaround: You can ignore this warning and proceed with the upgrade.

## Manual install of VRTSvc depot using /usr/sbin/swinstall -s 'pwd' VRTSvc may fail [2399744]

Manual installation of VRTSvc package using `/usr/sbin/swinstall -s `pwd`` VRTSvc might fail on freshly installed HP-UX machine.

AVXFS is an HP-owned integration product and has dependency on VRTSvlic 3.02.24.0. This causes `/usr/sbin/swinstall` to not select VRTSvc dependent products from the depot.

Workaround: The workaround for this issue is to remove the AONLINEJFS, OnlineJFS01, and AVXFS depots manually before stack installation. Moreover, before installing VCS on freshly installed HP-UX machine, uninstall older VRTSvlic 3.02.24.0 depot (if installed).

## Warning messages may be seen during script-based installation (2615500)

When you install VCS using the script-based installer, you may see the following warning message:

```
interpreter "/opt/VRTSperl/bin/perl" not found
```

**Workaround:** You must install perl to resolve the issue.

### To install perl

- 1 Exit the installer.
- 2 Install the VRTSperl depot from the product media manually:

```
cd /dvd_path/depot
/usr/sbin/swinstall -x enforce_dependencies=false
-x autoreboot=false -s `pwd` VRTSperl
```

- 3 Start the installer again.

## **During rolling upgrade swverify command displays errors [2439492]**

While performing Rolling upgrade from VCS 5.1SP1 to VCS 6.0 or later, after phase 1 of the Rolling Upgrade process, `swverify` command throws the following errors for all VCS related man pages:

```
"/opt/VRTS/man/man1m/hastatus.1m" missing.
```

This does not cause any product functionality issues. The man pages reappear after phase 2 of the Rolling Upgrade is complete.

**Workaround:** Not required, as the man pages reappear after Rolling Upgrade phase 2. In you wish to retain all the man pages even after phase 1, copy the `/opt/VRTS/man` directory in a safe location before starting the Rolling Upgrade procedure.

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

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

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

## **Errors seen during verification of VRTSamf depot (2599242)**

If VRTSvcsea depot is manually upgraded using `swinstall` command from VCS 5.1SP1RP1 to VCS 6.0, permissions for `/opt/VRTSamf/imf/imf_register` file are affected. Due to this, errors are seen during verification of VRTSamf.

**Workaround:** First uninstall the VRTSvcsea depot from VCS5.1SP1RP1 using `swremove` command and then proceed to install VRTSvcsea depot from VCS 6.0 release.

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

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

**Workaround:** Do one of the following:

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

## Operational issues for VCS

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

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

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

Workaround:

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

### **Saving large configuration results in very large file size for main.cf (616818)**

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.

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

## Issues related to the VCS engine

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

the agent process is running, HAD does not recognize it and hence does not perform any resource operations.

Workaround: Terminate the agent process.

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

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

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

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

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

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

### **If secure and non-secure WAC are connected the engine\_A.log receives logs every 5 seconds [2653695]**

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

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

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

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

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

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

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

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

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

When you initiate an offline of a service group and before the offline is complete, if you initiate a forced flush, the offline of the service group which was initiated earlier is treated as a fault. As start bits of the resources are already cleared, service group goes to OFFLINE|FAULTED state but TargetCount remains elevated.

Workaround: No workaround.

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

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

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

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

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

## GCO clusters remain in INIT state [2848006]

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

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

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

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

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

#### **Workaround**

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

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

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

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

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

or

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

## **Issues related to the bundled agents**

### **Process and ProcessOnOnly agent rejects attribute values with white spaces [2303513]**

Process and ProcessOnOnly agent does not accept Arguments attribute values that are separated by multiple whitespaces. The Arguments attribute specifies the set of arguments for a process. If a script controls the process, the script is passed as an argument. You must separate multiple arguments by using a single whitespace. A string cannot accommodate more than one space between arguments, or allow leading or trailing whitespace characters. This attribute must not exceed 80 characters.

Workaround: You should use only single whitespace to separate the argument attribute values. Make sure you avoid multiple whitespaces between the argument attribute values or trailing whitespace characters.

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

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

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

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

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

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

### **NIC agent may report incorrect interface state due to less traffic [2556355]**

When PingOptimize is set to 1 and no NetworkHosts is specified, NIC agent depends on packet count to report the health of the interface. If the traffic on the interface is not sufficient enough, NIC agent may report incorrect state of the interface.

Workaround: Any of the following workaround must resolve the issue:

- Setting PingOptimize = 0. This makes NIC agent ping the broadcast address whenever there is no traffic on the interface.
- Setting valid NetworkHosts value. This makes NIC agent to ping NetworkHosts to check health of status.

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

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

Workaround:

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

## CoordPoint agent remains in faulted state [2852872]

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

Workaround: After HAD has stopped, reconfigure fencing.

## Process resource fails to come online if call to exec are present in the profile of the root user [2611530]

Process agent fails to bring the resource online if there are calls to exec in the shell profile file of the root user.

Workaround: Make sure there no calls to exec in the shell profile of the root user.

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

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

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

## CoordPoint agent reports faulted if agile disk naming scheme is used [2871893]

If LevelTwo monitoring is enabled, i.e., if LevelTwoMonitorFreq attribute value is greater than zero, CoordPoint agent reports faulted if agile disk naming scheme is used on HP-UX platform.

Workaround: Use a disk naming scheme other than agile naming scheme.

### **Probing LVMLogicalVolume resource makes it OFFLINE when LVMLogicalVolume attribute is " " and volume group is not activated [2858185]**

If volume group is not activated and the LogicalVolume attribute is set to " " (empty string), a probe of LVMLogicalVolume resource shows the resource state as OFFLINE instead of UNKNOWN.

Workaround: Activate the volume group before configuring LVMLogicalVolume resource under VCS control.

### **LVMVolumeGroup resource remains ONLINE if VolumeGroup is deactivated outside VCS [2858165]**

If LVMVolumeGroup resource is deactivated outside of VCS, clean entry point cannot clean the resource. As a result, the resource state remains ONLINE and clean is called repeatedly.

Workaround: Do not deactivate LVMVolumeGroup from outside VCS.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

### **Service group with LVMLogicalVolume resource does not failover if the underlying storage is not available (2916108)**

If service group with LVMLogicalVolume resource is ONLINE and the underlying storage is disconnected or is unavailable, then it does not failover. This is because the clean entry point is not able to clean the resource, causing the service group to remain ONLINE.

Workaround: Make sure the underlying storage is always available.

## Issues related to the VCS database agents

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

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

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

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

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

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

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

Workaround:

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

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

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

### **NOFAILOVER action specified for certain Oracle errors**

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

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

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

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

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

## Issues related to the agent framework

### Agent may fail to heartbeat under heavy load [2073018]

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

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

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

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

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

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

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

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

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

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

Workaround: No workaround.

## Issues related to global clusters

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

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

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

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

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

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

## LLT known issues

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

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

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

## On reboot of cluster nodes that are connected via a single switch, a race condition may cause one of the llT links to not come up (2848001)

If cluster nodes are connected via a single switch and nodes are rebooted multiple times then sometimes a race condition may cause one of the links to be down. Run the `lltstat -nvv` command to know the link that is down.

**Workaround:** Restart LLT on the rebooted node.

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

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

**Workaround:** None

## GAB known issues

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

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

After you unregister the port with `-nodeinit` option, the `gabconfig -C` command shows `refcount` as 1. But when forceful `deinit` option (`gabdebug -R GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

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

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

**Workaround:** There is no workaround for this issue.

### Cluster panics during reconfiguration (2590413)

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

**Workaround:** There is no workaround for this issue.

## I/O fencing known issues

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

### Startup or shutdown failure messages reported for LLT, GAB, and VXFEN (1666327)

If you need to reboot the system when you install VCS, the init scripts for LLT, GAB, and VXFEN report start or stop failure messages. This is because VCS is not yet configured and the required configuration files are not yet generated for these components. These messages may be ignored.

### CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**Workaround:** Start VxFEN again after some time.

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

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSat depot is not removed from the system, the `cpsadm` command loads the old security libraries present on the system. As the installer runs the `cpsadm` command on the CP server to add or upgrade the VCS cluster (application cluster), the installer also fails.

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

To resolve this issue

- 1 Rename `cpsadm` to `cpsadmbin`:

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

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

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

- 3 Change the permissions of the new file to 775:

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

## ping\_cps and server\_security fail to communicate with the secure CPS Server if the CPS variables are not exported manually (2791763)

The `cpsadm -a ping_cps` and `cpsadm -a server_security` commands fail to communicate with the secure CPS Server from a client if the CPS variables are not exported manually.

**Workaround:** Set and export the following variables manually on the client cluster:

```
CPS_DOMAINTYPE="vx"
export CPS_DOMAINTYPE
EAT_HOME_DIR="/opt/VRTScps"
export EAT_HOME_DIR
CPS_HOME="/opt/VRTScps"
export CPS_HOME
CPS_USERNAME="CPSADM@VCS_SERVICES"
export CPS_USERNAME
```

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

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

### Workaround:

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

## Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later (2824472)

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

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

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

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

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

## vxfen module does not come up after phased upgrade from release version 4.1MP1 to 6.0.1 (2846209)

With HP-UX 11iv3, after upgrade, vxfen module does not allow raw disks to be specified as coordinator disks. So, even if you set the `vxfen_disk_policy` attribute to **raw** in the `/etc/vxfenmode` file fencing does not come up.

Workaround: Set the `vxfen_disk_policy` to `dmp` in the `/etc/vxfenmode` file.

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

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

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

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

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

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

## Unable to customize the 30-second duration (2551621)

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

**Workaround:** There is no workaround for this issue.

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

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

**Workaround:** There is no workaround for this issue.

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

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

**Workaround:**

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

### **Cannot run the `vxfersthdw` utility directly from the install media if `VRTSvxfer` package is not installed on the system (2858190)**

If `VRTSvxfer` package is not installed on the system, then certain script files that are needed for the `vxfersthdw` utility to function are not available. So, without the `VRTSvxfer` package installed on the system you cannot run the utility from the install media.

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

### **During a race scenario, the larger subcluster of a cluster can lose to a smaller subcluster, which may cause the large subcluster to panic (2858189)**

It may happen that during a split-brain scenario, GAB and `vxfer` modules may take more time to confirm memberships of nodes on a larger subcluster than the time taken to for the same action on a smaller subcluster. So, GAB and `vxfer` modules on the larger subcluster may lose the race to confirm new node memberships. Hence, the larger subcluster may panic.

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

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

#### **Workaround:**

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

### **Fencing command, `vxferadm`, does not print the registration keys in character format (2760308)**

The `vxferadm` command does print character format of keys with leading NULL bytes. This behavior happens because the `vxferadm` command prints entire registration key as a string and if there is a leading NULL byte in the string key the character format of the key is not printed.

Workaround: None

## Fencing may cause nodes panic in a network partitioning scenario (2858189)

After a network partition is done, on the racer node of the larger subcluster, GAB waits for 1 minute in order to get and confirm new memberships. After that, there is a delay of around 20 seconds due to GAB-vxfen integration wherein GAB waits for vxfen to fence out any nodes before GAB itself sends iofence to any of them. Due to these delays on the racer of the larger subcluster, the racer entered the race late and hence loses the race to the racer of the smaller subcluster, in spite of the fact that the racer of the smaller subcluster waited for 60 seconds before it starts racing. The 1 minute wait time by GAB is dictated by the numerical product of the value of `gab_conn_wait` tunable and the GAB stable timeout value.

### Workaround:

No workaround is available.

## The vxfentsthaw utility fails to launch before you install the VRTSvxfen package (2858190)

Before you install the VRTSvxfen package, the file of `/etc/vxfen.d/script/vxfen_scriptlib.sh` where stores the vxfentsthaw utility does not exist. In this case, the utility bails out.

### Workaround:

Besides installing the VRTSvxfen package, run the vxfentsthaw utility directly from the installation DVD.

## After you run the vxfenswap utility the CoordPoint agent may fault (3462738)

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

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

## Issues related to Intelligent Monitoring Framework (IMF)

### Registration error while creating a Firedrill setup [2564350]

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

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

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

**Workaround:** No workaround.

## Perl errors seen while using haimfconfig command

Perl errors seen while using `haimfconfig` command:

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

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

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

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

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

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

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

## IMF does not provide notification for a registered disk group if it is imported using a different name (2730774)

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

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

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

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

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

## Error messages displayed during reboot cycles [2847950]

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

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

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

Workaround: No workaround.

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

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

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

Workaround: No Workaround.

## Error message seen during system shutdown [2954309]

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

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

The system continues to proceed with the shutdown.

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

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

For example:

```
kill -9 -27824
```

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

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

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

Workaround: Use the correct commands:

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

## VCS 5.0.1 Rolling Patch 1 known issues

The VCS issues in this release are as follows:

- The Oracle agent with 11g Release 2 does not support Health check monitoring using the `MonitorOption` attribute. If the database is 11g Release 2, the `MonitorOption` attribute for the Oracle agent should be set to 0.

The Oracle agent with 11g Release 2 database does not support the Intentional Offline feature. [1975007]

- The `ASMinst` agent does not support `pfile` or `spfile` for the ASM Instance on the ASM diskgroups in 11g Release 2. Symantec recommends that you store the file on the local file system. [1975010]

- If you try to enable debug logs for the DB2 agent, the logs are not written to the `engine_A.log` file. [1954752]

Workaround: Download and install the GNU Awk software from the GNU Web site. Then, create a soft link to the default `awk` binary on the cluster nodes as follows:

```
ln -s /usr/local/bin/gawk /bin/awk
```

- The `VRTSperl` patch takes more than 10 minutes to install on an HP Integrity system node:

On an HP Integrity system node, installing the `VRTSperl` patch takes more than 10 minutes and requires that VCS is offline during this period. The installation time may vary based on the configuration of the machine on which the `VRTSperl` patch is being installed.

## Issues related to the Cluster Manager (Java Console)

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

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

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

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

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

## Software limitations

This section covers the software limitations of this release.

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

See [“Documentation”](#) on page 56.

### Limitations related to bundled agents

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

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can stop responding if the host is disconnected from the network. If such a program is used as an agent entry point, a network disconnect can cause the entry point to stop responding and possibly time out.

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

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

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

#### Volume agent clean may forcibly stop volume resources

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

## False concurrency violation when using PidFiles to monitor application resources

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

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

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

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

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

## HPVirtualMachine agent does not support migration of Integrity Virtual Machines by Stop and Resume technique

HPVirtualMachine agent does not support migration of Integrity Virtual Machines (IVM) by `stop` and Resume technique with IVM 4.3. See the *Veritas Storage Foundation and High Availability Solutions Virtualizations Guide* for more information.

## Application agent limitations

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

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

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

**Workaround:**

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

```
cd /opt/VRTSagents/ha/bin/<ResourceType>
ln -s /opt/VRTSvcs/bin/Script51Agent <ResourceType>Agent
```
- 2 Run the following command to update the AgentFile attribute based on value of VCS\_HOME.
  - If VCS\_HOME is `/opt/VRTSvcs`:  

```
hatype -modify <ResourceType> AgentFile
/opt/VRTSvcs/bin/<ResourceType>/<ResourceType>Agent
```
  - If VCS\_HOME is `/opt/VRTSagents/ha`:  

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

## Limitations related to the VCS database agents

### DB2 RestartLimit value [1234959]

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

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

### Limitation with intentional offline functionality of VCS agent for Oracle

The Oracle resource never faults after an intentional offline.

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

## 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.

## Limitations with DiskGroupSnap agent [1919329]

The DiskGroupSnap agent has the following limitations:

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

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

## Cluster Manager (Java console) limitations

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

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

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

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

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

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

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

## VCS Simulator does not support I/O fencing

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

## Limited support from Cluster Manager (Java console)

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

## Port change required to connect to secure cluster [2615068]

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

## Limitations related to I/O fencing

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

### Preferred fencing limitation when VxFEN activates RACER node re-election

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

### Stopping systems in clusters with I/O fencing configured

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

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

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

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

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

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

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

## Limitations related to global clusters

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

## Documentation

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

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

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

## Documentation set

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

**Table 1-9** Veritas Cluster Server documentation

| Title                                                                                               | File name                      |
|-----------------------------------------------------------------------------------------------------|--------------------------------|
| <i>Veritas Cluster Server Installation Guide</i>                                                    | vcs_install_604_lin.pdf        |
| <i>Veritas Cluster Server Release Notes</i>                                                         | vcs_notes_604_lin.pdf          |
| <i>Veritas Cluster Server Administrator's Guide</i>                                                 | vcs_admin_601_hpx.pdf          |
| <i>Veritas Cluster Server Bundled Agents Reference Guide</i>                                        | vcs_bundled_agents_601_hpx.pdf |
| <i>Veritas Cluster Server Agent Developer's Guide</i><br>(This document is available online, only.) | vcs_agent_dev_601_unix.pdf     |
| <i>Veritas Cluster Server Agent for DB2 Installation and Configuration Guide</i>                    | vcs_db2_agent_601_hpx.pdf      |
| <i>Veritas Cluster Server Agent for Oracle Installation and Configuration Guide</i>                 | vcs_oracle_agent_601_hpx.pdf   |
| <i>Veritas Cluster Server Agent for Sybase Installation and Configuration Guide</i>                 | vcs_sybase_agent_601_hpx.pdf   |

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

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

| Document title                                                                         | File name                        |
|----------------------------------------------------------------------------------------|----------------------------------|
| <i>Veritas Storage Foundation and High Availability Solutions Solutions Guide</i>      | sphas_solutions_601_hpx.pdf      |
| <i>Veritas Storage Foundation and High Availability Solutions Virtualization Guide</i> | sphas_virtualization_601_hpx.pdf |

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

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

## Manual pages

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

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

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

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

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

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

See the `man(1)` manual page.

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

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