# Symantec™ Cluster Server 6.1 Release Notes - AIX

January 2014

Symantec

# Symantec™ Cluster Server Release Notes

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

Product version: 6.1

Document version: 6.1 Rev 2

## Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

# Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

## Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

## Licensing and registration

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

www.symantec.com/business/support/

## Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

## Support agreement resources

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

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

## Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

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

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

## About Symantec Connect

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

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

# Symantec Cluster Server Release Notes

This document includes the following topics:



## About this document

This document provides important information about Symantec Cluster Server (VCS) version 6.1 for AIX. 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: 6.1 Rev 2" of the *Symantec Cluster Server Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec 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:

- *Symantec Storage Foundation Release Notes* (6.1)

# About Symantec Cluster Server

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

## About VCS agents

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

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

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

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

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

### About compiling custom agents

Custom agents developed in C++ must be compiled using the IBM XL C/C++ for AIX Compiler Version 8.0. Use the -brtl flag for runtime linking with the framework library.

# About Symantec Operations Readiness Tools

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

SORT can help you do the following:

| | |
|---|---|
| Prepare for your next installation or upgrade | ■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture.<br>■ Analyze systems to determine if they are ready to install or upgrade Symantec products and generate an Installation and Upgrade custom report.<br>■ List patches by product or platform, and in the order they need to be installed. Display and download the most recent patches or historical patches.<br>■ Display Array Support Library (ASL) details by vendor, platform, or Storage Foundation and High Availability (SFHA) version. ASLs make it easier to manage arrays that are connected to SFHA-based servers.<br>■ List VCS and ApplicationHA agents, documentation, and downloads based on the agent type, application, and platform. |
| Identify risks and get server-specific recommendations | ■ Analyze your servers for potential environmental risks. Generate a Risk Assessment custom report with specific recommendations about system availability, storage use, performance, and best practices.<br>■ Display descriptions and solutions for thousands of Symantec error codes. |

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

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

To access SORT, go to:

https://sort.symantec.com

# 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/TECH211540
- For the latest patches available for this release, go to: https://sort.symantec.com/
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware, visit the following URL: http://www.symantec.com/docs/TECH211575

- The software compatibility list summarizes each Storage Foundation and High Availability (SFHA) Solutions product stack and the product features, operating system versions, and third-party products it supports. For the latest information on supported software, visit the following URL:
  http://www.symantec.com/docs/TECH213121

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

# Changes introduced in 6.1

This section lists the changes in Symantec Cluster Server 6.1.

## Attributes introduced in VCS 6.1

The following section describes the attributes introduced in VCS 6.1.

### NFS agent attributes

| | |
|---|---|
| Protocol | Specifies the protocol to run the nfsd daemon. The agent uses this attribute to ensure that the NFS daemon is running using the specified protocol. |

### LPAR agent attributes on AIX

| | |
|---|---|
| ProfileFile | Specifies the path to the LPAR profile configuration file. |
| RemoveProfileOnOffline | Enables deleting the LPAR profile from the physical source server while bringing the LPAR resource offline or when the LPAR resource faults to facilitate the LPAR migration in future. |

### NotifierMngr agent attribute

| | |
|---|---|
| MessageExpiryInterval | Specifies time in seconds after which the messages expire. If the engine does not send a message to the notifier within MessageExpiryInterval, it deletes the message from the message queue of the engine. |

## DiskGroup agent attribute

| | |
|---|---|
| ClearClone | Clears the "clone" and "udid_mismatch" flags form the disks of the disk group while importing it and laso updates the UDID if required. |

## New attributes to enable AdaptiveHA

| | |
|---|---|
| Statistics | Indicates if statistics gathering is enabled and whether the FaultOverPolicy can be set to BiggestAvailable. Statistics are gathered for system resources like CPU, Memory, Swap, and so on. |
| MeterWeight | Represents the weight given for the cluster attribute's HostMeters key to determine a target system for a service group when more than one system meets the group attribute's Load requirements. |
| HostAvailableMeters | Lists the meters that are available for measuring system resources. You cannot configure this attribute in main.cf. |
| HostMeters | Indicates the parameters (CPU, Mem, or Swap) that are currently metered in the cluster. |
| MeterControl | Indicates the intervals at which metering and forecasting for the system attribute AvailableCapacity are done for the keys specified in HostMeters. |
| HostAvailableForecast | Indicates the forecasted available capacities of the systems in a cluster based on the past metered AvailableCapacity. |
| MeterRecord | Acts as an internal system attribute with predefined keys. This attribute is updated only when the Cluster attribute Statistics is set to Enabled. |

| | |
|---|---|
| ReservedCapacity | Indicates the reserved capacity on the systems for service groups which are coming online and with FailOverPolicy is set to BiggestAvailable. It has all of the keys specified in HostMeters, such as CPU, Mem, and Swap. The values for keys are set in corresponding units as specified in the Cluster attribute MeterUnit. |
| CapacityReserved | Indicates whether capacity is reserved to bring service groups online or to fail them over. Capacity is reserved only when the service group attribute FailOverPolicy is set to BiggestAvailable. |
| UnSteadyCount | Represents the total number of resources with pending online or offline operations. This is a localized attribute. |
| MemThresholdLevel | Determines the threshold values for memory utilization based on which various levels of logs are generated. |

Refer to *Symantec Cluster Server Administrator's Guide* for more information.

# Changes related to installation and upgrades

The product installer includes the following changes in 6.1.

## Support for SFHA 6.1 installations from any supported operating system to any other supported operating system

You can use the Deployment Server or the web-based installer to install your 6.1 Symantec products on a target system that runs any supported UNIX or Linux platform, even if the source system and target system are running on different UNIX or Linux platforms. Prior to 6.1, releases still require the same platform, architecture, distribution, and version of the operating system.

See the *Installation Guide* for more information.

## Automatic download of installer hot fixes

If you are running the 6.1 product installer, and your system has Internet access, the installer automatically imports any needed installer hot fix, and begins using it.

If your system does not have Internet access, you can still download installer hot fixes manually using the Symantec Operations Readiness Tools patch finder tool.

Automatic downloading of installer hot fixes requires the installer to make outbound networking calls. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network attempts by running the installer using the no Internet patch center (`-noipc`) option.

See the *Installation Guide* for more information.

## Support for centralized installations using the Deployment Server

The Deployment Server is a script that makes it easier to install or upgrade SFHA releases. The Deployment Server lets you store multiple release images in one central location and deploy them to systems of any supported UNIX or Linux operating system (6.1 or later). Prior to 6.1, releases still require the same platform, architecture, distribution, and version of the operating system. You can use the Deployment Server if you want to install or upgrade multiple releases and or multiple platforms.

The Deployment Server lets you do the following as described in Table 1-1.

**Table 1-1** Deployment Server functionality

| Feature | Description |
|---|---|
| Manage release images | ■ View available Storage Foundation releases.<br>■ Download maintenance and hot fix release images from the Symantec Operations Readiness Tools (SORT) website into a repository.<br>■ Load the downloaded release image files from FileConnect and SORT into the repository.<br>■ View and remove release image files stored in the repository. |
| Check versions | ■ Discovers filesets and patches installed on designated systems and informs you of the product and version installed, including installed hot fixes.<br>■ Identify base, maintenance, and hot fix level upgrades to your system and download maintenance and hot fix releases.<br>■ Query SORT for the most recent updates. |

**Table 1-1** Deployment Server functionality *(continued)*

| Feature | Description |
|---|---|
| Install or upgrade systems | ■ Install or upgrade a release stored in the repository on selected systems.<br>■ In release 6.1 and later:<br>  ■ Install hot fix level releases.<br>  ■ Install SFHA from any supported UNIX or Linux operating system to any other supported UNIX or Linux operating system.<br>  ■ Automatically load the script-based installer hot fixes that apply to that release. |

**Note:** The Deployment Server is available only for the script-based installer, not the web-based installer.

See the *Installation Guide* for more information.

## Improved patching and updating process

You can now download product maintenance releases and public hot fix releases directly from the Symantec Operations Readiness Tools (SORT) website using the installer. When you use the `installer` command with the `-version` option, the installer now lists the available GA releases, maintenance releases, and hot fix releases. If you have Internet access, you can follow the installer prompts to download available patches and hot fixes to your local system.

Downloading patches and hot fixes requires the installer to make outbound networking calls. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network attempts by running the installer using the no Internet patch center (`-noipc`) option. When using the `-noipc` option, the installer does not try to connect to SORT website. For example:

```
# ./installer -version -noipc system1 system2
```

See the *Installation Guide* for more information.

## Support for simultaneously installing or upgrading base releases, maintenance patches, and hot fixes

Beginning with version 6.1, Symantec offers you a method to easily install or upgrade your systems directly to a base, maintenance, or hot fix level in one step using

Install Bundles. Install Bundles is the ability for installers to merge so customers can install or upgrade directly to maintenance or hot fix levels in one execution. Install Bundles consists of executing the installer from a GA release with a pointer to a higher maintenance or hot fix release. The installer installs them both as if they were combined in the same release image. The various scripts, filesets, and patch components are merged and multiple releases are installed together as if they are one install entity.

---

**Note:** This feature is not supported by the Deployment Server.

---

There are five possible methods of integration. All upgrades must be executed from the highest level script.

- Base + maintenance
- Base + hot fix
- Maintenance + hot fix
- Base + maintenance + hot fix
- Base or maintenance + multiple hot fixes

See the *Installation Guide* for more information.

# Changes related to virtualization support in VCS

## Live migration of service groups

VCS now supports live migration capabilities for service groups that have resources to monitor virtual machines. The process of migrating a service group involves concurrently moving the service group from the source system to the target system with minimum downtime. A new entry point titled “migrate” is introduced for agent developer for this process. This entry point is available with the Script60Agent. The behavior of migrate entry point can be controlled using new attributes - MigrateTimeout, MigrateWaitLimit and SupportedOperations.

For more information, see the *Symantec Cluster Server Administrator's Guide* and *Symantec Storage Foundation and High Availability Solutions Virtualization Guide*.

# Changes to VCS bundled agents

This section describes changes to the bundled agents for VCS.

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

## IMF support for Apache HTTP server agent

The Apache HTTP server agent is IMF-aware and uses the AMF kernel driver for IMF notification. The agent also performs detailed monitoring on the Apache resource. You can tune the frequency of detailed monitoring with the LevelTwoMonitorFreq attribute. The SecondLevelMonitor attribute is deprecated.

## Support for level two monitoring in Application agent when MonitorProgram attribute is configured

If the Application resource is configured with MonitorProcesses, PidFiles or both along with MonitorProgram, you can configure the Application resource to run MonitorProgram as a level two monitor. To enable level two monitoring, set the LevelTwoMonitorFreq attribute to a value greater than zero. The default value of LevelTwoMonitorFreq attribute for Application resource is 1 (one).

With this change, the Application agent can leverage AMF for instant notification even when MonitorProgram is configured along with MonitorProcess or PidFiles or both.

## Proxy agent logs improved to provide more detail

The Proxy agent log messages now provide more detail such as the reason for the agent going to unknown or faulted state. Debug messages are also logged when the Proxy resource goes online or offline.

## Apache agent takes a resource offline when process stops [2978005]

Apache agent is now modified to take the resource offline immediately when the Apache processes stop as part of offline entry point.

## NFS agent enhancement

NFS agent supports running nfsd daemon in a specified protocol.

## New agent function for the Mount agent

The Mount agent supports the attr_changed function. This function unlocks the mount when you change the value of the VxFSMountLock attribute from either 1 or 2 to 0.

## LPAR agent enhancements

LPAR agent for AIX has been enhanced to include the following capabilities:

### Support for migration of LPAR resource through VCS

A new migrate entry point is added to the LPAR agent to initiate live migration of LPAR through VCS.

### Added support for LPAR profile management with LPAR live migration

LPAR agent is enhanced to support LPAR failover along with live migration of LPAR. When an LPAR resource is live migrated from a physical source server to a physical target server, the migration process deletes the LPAR profile on the physical source server. If the migrated LPAR on target physical server faults or if the LPAR service group is required to switch back from the target server, the LPAR cannot be brought online on the physical source server due to unavailability of the LPAR profile configuration. In order to facilitate the failover of the LPAR back to the physical source server, you must first create the LPAR profile configuration. The LPAR agent is enhanced to read LPAR configuration file as configured in ProfileFile attribute and create the LPAR on failover physical server while bringing the LPAR resource online. Similarly, the LPAR agent is enhanced to delete the LPAR configuration while bringing LPAR offline depending on RemoveProfileOnOffline attribute value.

## Changes to the VCS engine

### OpenVCSCommunicationPort attribute to determine whether to allow external communication port

The OpenVCSCommunicationPort attribute determines whether or not the external communication port for VCS is open for communication.

If the external communication port for VCS is not open, the following restrictions apply:

- You cannot use Java Console to administer VCS.
- RemoteGroup resources and users set up with the `hawparsetup` command cannot access VCS.

### AdaptiveHA

AdaptiveHA enables VCS to make dynamic decisions about selecting the cluster node with maximum available resources to fail over an application. VCS dynamically monitors the available unused capacity of systems in terms of CPU, Memory, and Swap to select the most resourceful system. For more information on AdaptiveHA, refer to the *Symantec Cluster Server Administrator's Guide*.

### Attributes modified to implement AdaptiveHA

To implement AdaptiveHA in VCS, the following attributes have been modified:

- HostUtilization: Indicates the percentage usage of the resources on the host as computed by the HostMonitor agent.
- FailOverPolicy: Governs how VCS calculates the target system for failover. Added a new policy value BiggestAvailable to this service group attribute.
  BiggestAvailable: VCS selects a system based on the forecasted available capacity for all the systems in the SystemList. The system with the highest forecasted available capacity is selected. This policy can be set only if the cluster attribute Statistics is enabled and the service group attribute Load is defined. Load must be defined in terms of CPU, Memory, or Swap in absolute units as specified in MeterUnit attribute.
- Load: Indicates the multidimensional value expressing load exerted by a service group of the system.
- HostMonitor: Contains list of host resources that the HostMonitor agent monitors.
- AvailableCapacity: Indicates the system's available capacity.
- Capacity: Represents total capacity of a system.

---

**Note:** AvailableCapacity, Capacity, Load, and DynamicLoad attributes have multi-dimensional values

---

## Changes to the Oracle agent

This section mentions the changes made to the Symantec Cluster Server agent for Oracle.

### VCS agent for Oracle uses the Oracle health check APIs to determine intentional offline of an Oracle instance

The Symantec Cluster Server agent for Oracle uses the Oracle health check APIs to determine whether the Oracle instance on a node was shut down gracefully or aborted. When an oracle instance is shut down gracefully outside of VCS control the agent acknowledges the operation as intentional offline.

From the VCS 6.1 release onwards, the pre-built health check binaries will not be shipped. You need to run the `build_oraapi.sh` script to build the oracle health check binaries based on the Oracle Version.

For more information, refer to the *Symantec Cluster Server Agent for Oracle Installation and Configuration Guide*.

# Changes to LLT, GAB, and I/O fencing

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

## Disable LLT, GAB, and I/O fencing on a single node cluster

Disable LLT, GAB, and I/O fencing kernel modules on a single node Symantec Cluster Server (VCS) cluster if you only want to manage applications and use the application restart capabilities of VCS for the node.

Note that disabling the kernel modules means that you cannot provide high availability to applications across multiple nodes. However, in future, if you decide to extend the cluster to multiple nodes, you can enable the modules and make the applications highly available.

For more information, refer to the *Symantec Cluster Server Installation Guide*.

## Changes to LLT

Symantec Cluster Server includes the following changes to LLT in 6.1:

### LLT command changes

The following command changes are introduced in this release.

Updates in `lltconfig`:

- A new option `lltconfig -l` is introduced. When you add a new link, you can use the `-l` option to specify that the link is a low priority link.

## Changes to I/O fencing

Symantec Cluster Server (VCS) includes the following changes to I/O fencing in 6.1:

### Set the order of coordination points while configuring I/O fencing

You can use the `-fencing` option in the installer to set the order of coordination points.

Decide the order of coordination points (coordination disks or coordination point servers) in which they participate in a race during a network partition. The order of coordination points you set in the installer is updated to the /etc/vxfenmode file. I/O fencing approaches the coordination points based on the order listed in the vxfenmode file.

So, the order must be based on the possibility of I/O Fencing reaching a coordination point for membership arbitration.

For more information, refer to the *Symantec Cluster Server Installation Guide*.

### Refresh keys or registrations on the existing coordination points using the install program

You can use the `-fencing` option with the installer to refresh registrations on the existing coordination points.

Registration loss on the existing coordination points may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs. You must refresh registrations on coordination points when the CoordPoint agent notifies VCS about the loss of registrations on any of the existing coordination points.

You can also perform a planned refresh of registrations on coordination points when the cluster is online without application downtime on the cluster.

For more information, refer to the *Symantec Cluster Server Installation Guide*.

### CPI automatically installs a CP server-specific license while configuring CP server on a single-node VCS cluster

The installer automatically installs a CP server-specific license if you are configuring CP server on a single-node VCS cluster. It also ensures that Veritas Operations Manager (VOM) identifies the license on a single-node coordination point server as a CP server-specific license and not as a VCS license.

For more information, see the *Symantec Cluster Server Installation Guide*.

### Site-based preferred fencing policy

The fencing driver gives preference to the node with higher site priority during the race for coordination points. VCS uses the site-level attribute Preference to determine the node weight.

For more information, see the *Symantec Cluster Server Administrator's Guide*.

### Support for HTTPS communication between CP server and application client cluster nodes

CP server and its application client cluster nodes can communicate securely over HTTPS, an industry standard protocol. Prior to release 6.1, communication between the CP server and its clients happened over the Inter Process Messaging (IPM) protocol, which is a Symantec proprietary protocol. Secure communication over IPM-based communication uses Symantec Product Authentication Services (AT) to establish secure communication between CP server and client nodes. With secure communication using HTTPS, CP server functionality is backward-compatible with previous releases. To support client nodes on releases before 6.1, CP server

supports IPM-based communication in addition to HTTP-based communication. However, client nodes from 6.1 onwards only support HTTPS-based communication.

For more information, refer to the Symantec Cluster Server Installation Guide and Symantec Cluster Server Administrator's Guide.

### The security attribute in `/etc/vxfenmode` file is obsolete

From VCS 6.1, the Coordination Point (CP) client will communicate with CP server using HTTPS protocol. The 'security' parameter in `/etc/vxfenmode` is therefore deprecated and setting it to 1 or 0 has no effect whatsoever.

### Rolling upgrade of an application cluster to release version 6.1 requires CP server running release version 6.1

The application clusters and CP servers running on release version 6.1 communicate over the HTTPS protocol. Hence, an application cluster which is using CP server as a fencing coordination point can no longer access the pre-6.1 CP server after the cluster is upgraded to 6.1. To ensure a smooth upgrade, either application cluster must use CP servers running release version 6.1 or the CP servers running an earlier release version must be upgraded to 6.1. Note that CP server running release version 6.1 can still work with pre-6.1 application clusters.

### Checks introduced in `vxfentsthdw` utility for disk size and option to override errors

The `vxfentsthdw` utility is enhanced to check the disks for size compatibility and new error messages are introduced for better error evaluation. The utility also provides the override option (`-o`) to override size-related errors and continue testing.

### New command for `hacli` in `vxfenswap` utility

A new option `-p` is introduced to specify a protocol value that `vxfenswap` utility can use to communicate with other nodes in the cluster.The supported values for the protocol can be `ssh`, `rsh`, or `hacli`.

# Changes to campus clusters

## Multi-site management

You can create sites to use in an initial failover decision in campus clusters by configuring the SiteAware cluster level attribute. You can define sites and add systems to the sites that you have defined. A system can belong to only one site. Site definitions are uniform across VCS, Veritas Operations Manager, and VxVM. You can define site dependencies to restrict connected applications to fail over within the same site.

If sites are configured for a cluster, a service group tries to stay within its site before choosing a host in another site. For example, in a campus cluster with two sites, site A and site B, you can define a site dependency among service groups in a three-tier application infrastructure consisting of Web, application, and database to restrict the failover within the same site.

You must have the Veritas Operations Manager 6.0 to define sites and dependencies and configure site for a cluster.

Refer to the *Symantec Cluster Server Administrator's Guide* for more information.

## Changes related to product name branding

Beginning with the 6.1 release, Storage Foundation and High Availability Solutions product names are rebranded.

Table 1-2 lists the rebranded Storage Foundation and High Availability Solutions products.

**Table 1-2** Rebranded Storage Foundation and High Availability Solutions products

| Old product name | New product names with Symantec branding |
|---|---|
| Veritas Storage Foundation | Symantec Storage Foundation (SF) |
| Veritas Dynamic Multi-Pathing | Symantec Dynamic Multi-Pathing (DMP) |
| Veritas Replicator Option | Symantec Replicator Option |
| Veritas Volume Replicator | Symantec Volume Replicator (VVR) |
| Veritas Storage Foundation Cluster File System HA | Symantec Storage Foundation Cluster File System HA (SFCFSHA) |
| Veritas Storage Foundation for Oracle RAC | Symantec Storage Foundation for Oracle RAC (SFRAC) |
| Veritas Storage Foundation HA | Symantec Storage Foundation HA (SFHA) |
| Veritas Cluster Server | Symantec Cluster Server (VCS) |
| Veritas Disaster Recovery Advisor | Symantec Disaster Recovery Advisor (DRA) |
| Veritas Storage Foundation and High Availability Solutions | Symantec Storage Foundation and High Availability Solutions (SFHAS) |
| Veritas High Availability Agent Pack | Symantec High Availability Agent Pack |

**Table 1-2** Rebranded Storage Foundation and High Availability Solutions products *(continued)*

| Old product name | New product names with Symantec branding |
|---|---|
| Veritas File System Software Development Kit | Symantec File System Software Development Kit |

Symantec rebranding does not apply to the following:

- Product acronyms
- Command names
- Error messages
- Alert messages
- Modules and components
- Feature names
- License key description
- Veritas Operations Manager product branding

# 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 supports an environment where a few nodes in the cluster are hosted on LPARs with storage and network connectivity presented to the OS using VIOS. The remaining nodes in the cluster are hosted on physical systems with storage and network connectivity presented to the OS directly. However SCSI3 I/O fencing will be supported in this environment only if storage is available through NPIV in the LPARs. If NPIV is not available in LPARs, non-SCSI3 fencing is supported.

VCS requires that all nodes in the cluster use the same processor architecture and all nodes in the cluster must run the same VCS version. Each node in the cluster may run a different version of the operating system, as long as the operating system is supported by the VCS version in the cluster.

See "Hardware compatibility list" on page 25.

See "Supported AIX operating systems " on page 25.

## Hardware compatibility list

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

http://www.symantec.com/docs/TECH211575

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

## Supported AIX operating systems

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

Table 1-3 shows the supported operating systems for this release.

**Table 1-3** Supported operating systems

| Operating systems | Levels | Chipsets |
|---|---|---|
| AIX 7.1 | TL0, TL1, or TL2 | Power 5, Power 6, or Power 7 |
| AIX 6.1 | TL6, TL7, or TL8 | Power 5, Power 6, or Power 7 |

## Supported software for VCS

VCS supports the following volume managers and file systems:

- Logical Volume Manager (LVM)
- Journaled File System (JFS) and Enhanced Journaled File System (JFS2) on LVM

VCS supports the following versions of Symantec Storage Foundation:

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

- Storage Foundation 6.1
  - VxVM 6.1 with VxFS 6.1
- Storage Foundation 6.0.1
  - VxVM 6.0.1 with VxFS 6.0.1

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

For supported database versions of enterprise agents, refer the support matrix at http://www.symantec.com/business/support/index?page=content&id=DOC4039.

## Supported enterprise agents

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

| | |
|---|---|
| Oracle | Support matrix for Oracle |
| DB2 | Support matrix for DB2 |
| Sybase | Support matrix for Sybase |

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

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

# No longer supported

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

## No longer supported agents and components

VCS no longer supports the following:

- The configure_cps.pl script used to configure CP server is now deprecated and is no longer supported.
- The 'security' parameter has been deprecated since communication with CP Server will always happen over HTTPS. Hence enabling or disabling this parameter in `/etc/vxfenmode` will not have any effect.

## Deprecated attributes

The following table lists the attributes deprecated in this release.

**Table 1-4** Attributes deprecated in this release

| Attribute name | Agent type |
|---|---|
| SecondLevelMonitor | Apache<br>**Note:** The SecondLevelMonitor attribute is deprecated in VCS 6.1. Instead, LevelTwoMonitorFreq attribute at the Apache resource type level may be used |
| DetailMonitor | Oracle, Sybase<br>**Note:** If you manually upgrade VCS to 6.1 with detail monitoring enabled in the previous version, set the value of LevelTwoMonitorFreq attribute to that of DetailMonitor. |

# Fixed issues

This section covers the incidents that are fixed in this release.

## LLT, GAB, and I/O fencing fixed issues

Table 1-5 lists the fixed issues for LLT, GAB, and I/O fencing.

**Table 1-5** LLT, GAB, and I/O fencing fixed issues

| Incident | Description |
|---|---|
| 2619600 | After you execute Live Partition Mobility (LPM) on an SFHA or SFCFSHA node with SCSI-3 fencing enabled for data disks, I/O fails on devices or disks with reservation conflict. |
| 2869763 | When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/lltab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files. |
| 2991093 | The preferred fencing node weight does not get reset to the default value when HAD is terminated. In spite of lack of high availability on that node, fencing may give preference to that node in a network partition scenario. |

**Table 1-5** LLT, GAB, and I/O fencing fixed issues *(continued)*

| Incident | Description |
|---|---|
| 2995937 | The default value of preferred fencing node weight that vxfen uses is 1 (one). However, when HAD starts without any service group or if HAD is stopped or terminated, the node weight is reset to 0 (zero). Since vxfen resets the preferred fencing weight to its default value when HAD gets terminated, stopping HAD and killing HAD shows different preferred fencing weight. |
| 2110148 | Installer is unable to split a cluster that is registered with one or more CP servers. |
| 2802682 | Server-based fencing may fail to start if you use the existing configuration files after reinstalling the stack. |
| 2858190 | If VRTSvxfen fileset is not installed on the system, then certain script files that are needed for the vxfentsthdw utility to function are not available. So, without the VRTSvxfen fileset installed on the system you cannot run the utility from the install media. |
| 3101262 | GAB queue is overloaded causing memory pressure during I/O shipping. |
| 3218714 | GAB does not log messages about changing tunable values. |
| 2858076 | Changing the module parameter `gab_conn_wait` had no effect. |

# Installation related fixed issues

**Table 1-6** Installation related fixed issues

| Incident | Description |
|---|---|
| 2873102 | When you install, configure, or uninstall VCS, the installer prompts you to optionally upload installation logs to the Symantec Web site. If the installer encounters connectivity problems, you may see an error . |
| 2737124 | If you upgrade the `VRTSvlic` fileset manually, the product levels that were set using `vxkeyless` may be lost. The output of the `vxkeyless display` command does not display correctly. |
| 2141446 | After upgrading from VCS 5.1 to higher versions of VCS, some keyless licenses may be left in the system. As a result periodic reminders get logged if Veritas Operations Manager Server is not configured. |

# VCS engine fixed issues

Table 1-7 lists the fixed issues for VCS engine.

**Table 1-7** VCS engine fixed issues

| Incident | Description |
|---|---|
| 2858188 | If you attempt to reconfigure an already configured Global Cluster Option (GCO) using `gcoconfig`, the command does not change the existing GCO IP while reconfiguring the global cluster options. |
| 2941155 | Symantec Cluster Server (VCS) does not mark a group as offline on a failed cluster when a cluster failure is declared in a GCO environment . |
| 2954319 | On a heavily loaded system, the logger thread frequently picks the SIGABRT from GAB. The logger thread runs at a low priority and may not get scheduled. Hence, the SIGABRT is not processed and GAB panics the machine. |
| 2736627 | Remote cluster state remains in INIT state and Icmp heartbeat status remains UNKNOWN if IPv6 is disabled on the systems. |
| 3042450 | Parent service group which if frozen and configured with online local hard dependency is brought offline when its child service group faults. |
| 3079893 | Symantec Cluster Server does not retry to online a service group when a resource in the service group faults while the group is being brought online and when OnlineRetryLimit and OnlineRetryInterval for the service group is set to non -zero values. |
| 3090710 | High Availability Daemon (HAD) starts and stops before VxFEN driver configuration completes. |
| 3207663 | When user fires 'hauser -addpriv' command to set user privileges for a group and provides any string without dash (-) instead of the '-group' option syntax error is not seen and incorrect privileges are set. |
| 3112608 | Resource is unable to come online after switch fails for a service group. |

## Bundled agents fixed issues

Table 1-8 lists the fixed issues for bundled agents.

**Table 1-8** Bundled agents fixed issues

| Incident | Description |
|---|---|
| 2989861 | Incorrect command usage is displayed for `havmconfigsync` command. |
| 2952387 | LPAR agent needs to be modified for VCS to use live migration during switch over. |

**Table 1-8** Bundled agents fixed issues *(continued)*

| Incident | Description |
|---|---|
| 2962270 | Apache agent requires online monitoring IMF support. |
| 2954312 | If hostname of the DLPAR and name of DLPAR as seen from HMC are different, the MemCPUAllocator agent is unable to provide CPU or memory to the DLPAR. |
| 2964772 | NFSRestart Agent may unexpectedly stop NFS processes in a local container, if an NFSRestart resource is taken offline. |
| 2523171 | If ContainerInfo attribute is already set for a service group and the key "Enabled" is set to some value other than 1 (one), running `hawparsetup.pl` overwrites the value for key "Enabled" to 1. Thus, `hawparsetup.pl` does not check whether the key "Enabled" in "ContainerInfo" attribute has been set or not. |
| 3315273 | LVMVG resource fails to come online if volume group configured in the resource has more than 128 disks. |
| 3028760 | NFSRestart resource does not start NFS processes, such as `statd` and `lockd`, during online or offline operations. |

# Fixed issues related to AMF

**Table 1-9** AMF fixed issues

| Incident | Description |
|---|---|
| 2937673 | A race condition arises in the context of amfstat, group unregistration, and event notification, which causes the AMF driver to panic. |
| 2848009 | If an agent exits while AMF is notifying it about some events, sometimes AMF causes the node to panic. |
| 2703641 | VRTSamf patch gets installed or uninstalled when some events monitored by amf remains registered even after the patch is installed or uninstalled. |
| 3030087 | The `amfconfig -Uo` command must stop IMFD and other functions internally started or setup by AMF. |
| 2954309 | Unconfigure AMF forcefully from the AMF stop script to remove any dependencies that agents might have on the AMF. |

**Table 1-9** AMF fixed issues *(continued)*

| Incident | Description |
|---|---|
| 3191098 | AMF driver panics the machine when a VXFS filesystem is unmounted. Sometimes a function calls itself infinitely causing stack corruption and panic due to the AMF driver. |
| 3090229 | The libusnp_vxnotify.so library used for disk group notifications, goes into an infinite loop when vxconfigd daemon is unresponsive. This causes AMF to enter an inconsistent state as a result of which AMF driver panics the node. |
| 3145047 | Due to the way AMF interacts with VXFS, AMF has access into VXFS driver even if no mounts are online, without actually holding a module reference on it. Therefore VXFS can get unloaded despite AMF having access into it. |
| 3133181 | Due to an operational error in AMF driver, in some cases an ioctl made by IMFD into AMF gets stuck inside AMF. The IMFD process cannot exit until this ioctl returns back to userspace. |
| 3018778 | Perl errors seen while using `haimfconfig` command. |
| 3259682 | If `vxconfigd` hangs, then registration thread of imfd trying to get disk group status from `vxconfigd` also hangs. Therefore, the `amfregister` command waiting for IMFD gets stuck. |
| 3279336 | If AMF is unconfigured while a disk group resource registration with AMF is going on, then both the contexts may enter hung state. |
| 3177476 | If a process online registration with AMF is unregistered after it has already been triggered, the machine panics. |

# Enterprise agents fixed issues

Table 1-10 lists the fixed issues for enterprise agents.

**Table 1-10** Enterprise agents fixed issues

| Incident | Description |
|---|---|
| 1938138 | The health check monitoring in Oracle agent for VCS does not work due to incompatibility of the health check APIs provided by Oracle. |
| 3088915 | VCS reports the status of Oracle resources configured inside the container as OFFLINE even when Oracle processes are running inside the container. |

**Table 1-10** Enterprise agents fixed issues *(continued)*

| Incident | Description |
|---|---|
| 2847994 | The ASMDG agent delays the exit of offline entry point when it finds the device (any one of the volume) busy as indicated by the user command. For each of the disk group mentioned in ASMDG agent's DiskGroups attribute, agent runs an SQL command and gets the list of volumes used by it. |
| 3240209 | During the Oracle online operation, the Oracle agent unnecessarily tries to back up the database due to an incorrect pattern match. |
| 1805719 | Due to issues with health check monitoring, Intentional Offline does not work for VCS agent for Oracle. |

## Fixed operational issues

Table 1-11 lists the fixed issues for enterprise agents.

**Table 1-11** Fixed operational issues

| Incident | Description |
|---|---|
| 3210553 | If the system tags are modified without selecting the fencing option in a Replicated Data Cluster (RDC) setup, the Stretch site wizard fails to modify tags. |

# Known issues

This section covers the known issues in this release.

## 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:

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

2 Unfreeze all the frozen service groups:

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

## Erroneous resstatechange trigger warning

You may encounter the following warning when you restart resources:

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

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

## The VRTSsfcpi fileset is retained after you upgrade to 6.1 on an alternate disk (2811749)

On AIX, if you run the command `alt_disk_scenario` to perform a disk clone and upgrade from 6.0 or later to 6.1, the older version of the VRTSsfcpi fileset is retained.

**Workaround**: Optionally uninstall the older VRTSsfcpi60 fileset after upgrading. Retaining the older version will not cause any harm.

## VRTSvcsea package cannot be uninstalled from alternate disk in manual live upgrade

Description: In manual live upgrade procedure from 5.1x to 5.1SP1 , all packages are copied to an alternate root disk. However, VRTSvcsea package cannot be uninstalled from alternate disk to upgrade it to 5.1SP1.

Workaround : Instead of removing the VRTSvcsea package, you must apply a patch to upgrade this package to 5.1SP1 version.

## 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.

## Perl messages seen in engine log during rolling upgrade [2627360]

While performing a rolling upgrade from VCS 5.1SP1 to 6.0 with MultiNICA resource configured, if VRTSperl fileset is upgraded but VRTSvcsag fileset is not yet upgraded on the system, Perl code related messages may be seen. The messages seen are similar to the following:

```
Using a hash as a reference is deprecated at MultiNICA.pm line 108.
```

Workaround: Complete the rolling upgrade to VCS 6.0.

## 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.

## If you have a non-shared (detached) WPAR configured, when you install, upgrade, or install any Symantec product, the filesets in the WPAR cannot be installed, upgraded, or uninstalled correspondingly (3313690)

On AIX, if you have a non-shared (detached) workload partition (WPAR) configured, when you perform an install, upgrade, or uninstall task on any Symantec product by the Symantec product installer, the filesets cannot be installed, upgraded, or uninstalled inside the WPAR correspondingly.

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

### If you have a shared (system) WPAR configured, when you install, upgrade, or uninstall any Symantec product, the filesets in the WPAR are not synchronized correspondingly (3313690)

On AIX, if you have a shared (system) workload partition (WPAR) configured, when you perform an install, upgrade, or uninstall task on any Symantec product by the Symantec product installer, the filesets may not be installed, upgraded, or uninstalled correspondingly.

**Workaround:** After an install, upgrade, or uninstall task, execute the following command to synchronize your WPAR with global systems:

```
# /usr/sbin/syncwpar -A
```

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

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

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

## Operational issues for VCS

### Connecting to the database outside VCS control using sqlplus takes too long to respond

Connecting to start the database outside VCS control, using sqlplus takes more than 10 minutes to respond after pulling the public network cable. [704069]

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

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

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

- Setting up trust relationships between servers fails.

Workaround:

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

## 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.

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

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

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

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

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

Workaround: No workaround.

### CP server service group fails to come online with the default database path after the CP server is upgraded from 6.0 to 6.1 on a multi-node cluster [3326639]

If the CP server is configured on a multi-node cluster before the upgrade with security enabled, you must reconfigure the CP server after the CP server upgrade. If you reuse the old credentials with the old database path, the CP server service group does not come online. Since the default database paths of CP server in 6.0 and 6.1 are different, reusing the old credentials and default database path prevents the CP server service group from coming online.

**Workaround:**

If the CP server multi-node cluster is configured with security enabled and if the old credentials such as database path are expected to be reused in reconfiguration of the CP server after the upgrade of the CP server, use the same database path before and after the upgrade.

## 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.

## VCS fails to validate processor ID while performing CPU Binding [2441022]

If you specify an invalid processor number when you try to bind HAD to a processor on a remote system, HAD does not bind to any CPU. However, the command displays no error to indicate that the specified CPU does not exist. The error is logged on the node where the binding has failed and the values are reverted to default.

Workaround: Symantec recommends that you modify CPUBinding from the local system.

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

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

Workaround: Remove the extra leading or trailing '/' characters from the path.

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

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

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

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

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

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

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

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

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

Workaround: Terminate the agent process.

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

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

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

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

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

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

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

Two WACs in a global service group 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 a global service group.

## 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.

## System sometimes displays error message with `vcsencrypt` or `vcsdecrypt` [2850899]

If random number generator is not configured on your system and you run vcsencrypt or vcsdecrypt, the system sometimes displays the following error message:

```
VCS ERROR V-16-1-10351 Could not set FIPS mode
```

Workaround: Ensure that the random number generator is defined on your system for encryption to work correctly. Typically, the files required for random number generator are /dev/random and /dev/urandom.

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

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

1. If it is able to confirm that clus2 is down, it will mark clus2 as FAULTED.

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

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

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

## GCO clusters remain in INIT state [2848006]

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

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

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

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

The `ha` commands fail to work 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).

## Every `ha` command takes longer time to execute on secure FIPS mode clusters [2847997]

In secure FIPS mode cluster, `ha` commands take 2-3 seconds more time than in secure cluster without FIPS mode for non-root users. This additional time is required to perform the FIPS self-tests before the encryption module can be used in FIPS mode.

Workaround: No workaround.

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

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

Workaround: No workaround.

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

VCS enters into admin_wait state when started locally if:

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

Workaround:

1. Stop VCS on all nodes in the cluster.
2. Perform any one of the following steps:
   - Edit the `main.cf` on one of the nodes in the cluster and set the Statistics attribute to Disabled or MeterHostOnly.
   - Remove the Group Load and System Capacity values from the `main.cf`.
3. Run `hacf -verify` on the node to verify that the configuration is valid.
4. Start VCS on the node and then on the rest of the nodes in the cluster.

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

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

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

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

### Site preference fencing policy value fails to set on restart of a site-aware cluster [3380584]

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

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

## Issues related to the bundled agents

### VCS resources may time out if NFS server is down [2129617]

The VCS resources may time out if the server NFS mounted file system and the NFS server is down or inaccessible. This behavior is exclusive to AIX platform.

Workaround: You must unmount the NFS mounted file system to restore the cluster to sane condition.

### MultiNICB resource may show unexpected behavior with IPv6 protocol [2535952]

When using IPv6 protocol, set the LinkTestRatio attribute to 0. If you set the attribute to another value, the MultiNICB resource may show unexpected behavior.

Workaround: Set the LinkTestRatio attribute to 0.

### 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.

### Bringing the LPAR resource offline may fail [2418615]

Bringing the LPAR resource offline may fail with the following message in the engine_A.log file.

```
<Date Time> VCS WARNING V-16-10011-22003 <system_name>
LPAR:<system_name>:offline:Command failed to run on MC
<hmc_name> with error HSCL0DB4 An Operating System
Shutdown can not be performed because the operating system image
running does not support remote execution of this task from the HMC.
This may be due to problem in communication with
MC <hmc_name>
```

This is due to RMC failure between HMC and management LPAR. Since the LPAR could not be shutdown gracefully in offline, the LPAR is shutdown forcefully in the clean call, hence it shows as Faulted.

Workaround: In order to recycle the RSCT deamon for LPAR and HMC, refer the *Symantec Storage Foundation™ and High Availability Solutions Virtualization Guide*.

## LPAR agent may not show the correct state of LPARs [2425990]

When the Virtual I/O server (VIOS) gets restarted, LPAR agent may not get the correct state of the resource. In this case, the LPAR agent may not show the correct state of the LPAR.

Workaround: Restart the management LPAR and all the managed LPARs that depend on the VIOS.

## 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.

## Prevention of Concurrency Violation (PCV) is not supported for applications running in a container [2536037]

For an application running in a container, VCS uses a similar functionality as if that resource is not registered to IMF. Hence, there is no IMF control to take a resource offline. When the same resource goes online on multiple nodes, agent detects and reports to engine. Engine uses the offline monitor to take the resource offline. Hence, even though there is a time lag before the detection of the same resource coming online on multiple nodes at the same time, VCS takes the resource offline.

PCV does not function for an application running inside a WPAR on AIX.

Workaround: No workaround.

## VCS does not monitor applications inside an already existing WPAR [2494532]

If a WPAR is already present on the system at the time of VCS installation, and this WPAR or an application running inside this WPAR needs to be monitored using VCS, then VCS does not monitor the application running in that WPAR. This is because the VCS packages/files are not visible inside that WPAR.

Workaround: Run `syncwpar` command for that WPAR. This makes the VCS packages/files visible inside the WPAR and VCS can then monitor the applications running inside the WPAR.

## 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 and DNS 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.

## Error messages for wrong HMC user and HMC name do not communicate the correct problem

The wrong HMC user and wrong HMC name errorsd are not reflective of the correct problem. If you see the following errors in engine_A.log for LPAR resource, it means wrong HMC user:

```
Permission denied, please try again
Permission denied, please try again
```

If you see the following errors in engine_A.log for LPAR resource, it means wrong HMC name:

```
ssh: abc: Hostname and service name
not provided or found.
```

You must see the applicationha_utils.log file to confirm the same.

## LPAR agent may dump core when all configured VIOS are down [2850898]

When using Virtual Input Output Servers (VIOS), the LPARs need a restart after VIOS restart/reboot/crash. If management LPAR is not restarted after VIOS is rebooted, then LPAR agent may dump core.

Workaround: Restart the management LPAR which was depended on the rebooted VIOS.

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

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

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

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

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

2 Enable the preonline trigger for the service group.

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

## WPAR-aware agents cannot run in a non-shared WPAR [3313698]

Non-shared WPAR has writable `/usr` file system and `/opt` file system local to WPAR. The common product installer installs VCS packages in `/opt/VRTSvcs`and

libraries in `/usr/lib` in a global environment. As VCS packages cannot be synchronized with a local copy of `/usr` and `/opt` on a non-shared WPAR, they are not available to the non-shared WPAR. Therefore in absence of the VCS packages, agents which are configured to monitor applications inside non-shared WPAR cannot run.

Workaround: No workaround.

## WPAR-aware agent fails to log messages in a secure cluster [3343222]

WPAR-aware agent fails to log messages in a secure cluster.

**Workaround: Perform the following steps on every cluster node:**

1 If the WPAR resource does not exist, then run `hawparsetup` utility:

```
# /opt/VRTSvcs/bin/hawparsetup.pl <service_group> <resource>
  <WPAR> <password> <systems>
```

2 To find domain name in FQDN format, run:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat
  listpd -t ab|grep -w vcs_lzs | awk {'print $NF'}
```

The above command will give the fully qualified domain name.

For example: `vcs_lzs@81a4f374-1dd2-11b2- 80a4-163d778711bd`

3 If the output of the command mentioned in step 2 provides a domain name, go to step 5, or else generate the domain name in an FQDN format using the following command:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat
  createpd -t ab -d vcs_lzs
```

4 Find domain name in FQDN format

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat
  listpd -t ab|grep -w vcs_lzs | awk {'print $NF'}
```

5 Use this above domain name in the below following commands where FQDN is mentioned.

```
# /opt/VRTS/bin/hauser -add
  w_<wpar_resource_name>_<clustername>@FQDN
#  /opt/VRTS/bin/hauser  -addpriv
  w_<wpar_resource_name>_<clustername>@FQDN
  Administrator -group <service_group>
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat
  showprpl --pdrtype ab --domain vcs_lzs
--prplname w_<wpar_resource_name>_<clustername>
#  echo $?
```

If the return code is non-zero then run command,

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat
  addprpl -t ab -d vcs_lzs -p w_<wpar_resource_name>_<clustername>
  -s <password> -q service
# clogin <wpar_name> "VCS_HOST=<host_name>; export VCS_HOST;
  VCS_DOMAIN=<FQDN>; export VCS_DOMAIN; VCS_DOMAINTYPE=vx; export
  VCS_DOMAINTYPE; /opt/VRTSvcs/bin/halogin
  w_<wpar_resource_name>_<clustername>  <password>""
```

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

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

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

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

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

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

Workaround: No workaround.

## Issues related to the VCS database agents

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

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

Workaround:

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

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

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

### NOFAILOVER action specified for certain Oracle errors

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

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

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

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

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

### IMF registration fails if sybase server name is given at the end of the configuration file [2365173]

AMF driver supports a maximum of 80 characters of arguments. In order for AMF to detect the start of the Sybase process, the Sybase server name must occur in the first 80 characters of the arguments.

Workaround: Must have the server name, -sSYBASE_SERVER, as the first line in the configuration file: `ASE-15_0/install/RUN_SYBASE_SERVER`.

## Issues related to the agent framework

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

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

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

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

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

### Agent framework cannot handle leading and trailing spaces for the dependent attribute (2027896)

Agent framework does not allow spaces in the target resource attribute name of the dependent resource.

Workaround: Do not provide leading and trailing spaces in the target resource attribute name of the dependent resource.

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

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

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

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

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

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

Workaround: No workaround.

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

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

Some of the commands are mentioned below:

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

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

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

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

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

   Where <TypeName> is the name of the agent with which you are facing delays and <value> is any numerical value appropriate for your environment.
3. If the agent supports IMF, increase the value of MonitorFreq attribute of IMF.

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

   Where <value> is any numerical value appropriate for your environment.
4. Wait for several minutes to ensure that VCS has executed all pending commands, and then execute any new VCS command.
5. If the delay persists, repeat step 2 or 3 as appropriate.
6. If the CPU usage returns to normal limits, revert the attribute changes to the backed up values to avoid the delay in detecting the resource fault.

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

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

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

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

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

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

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

# Issues related to global clusters

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

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

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

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

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

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

# LLT known issues

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

### LLT may fail to make connections with LLT on peer nodes in virtual environment (2343451/2376822)

After you upgrade from 5.0 MP3 or earlier releases to version 6.0, LLT may fail to make connections with LLT on the peer nodes in AIX virtual environment.

This is a known IBM VIOS issue. Install APAR IV00776 on your VIOS server. Without this fix, VIOS fails to handle new LLT packet header and drops packets.

**Workaround:** Disable the largesend attribute of the SEA adapter. Check the properties of the SEA adapter (on which the virtual links are configured under LLT maps) on each VIOS using the following command:

```
# lsattr -El SEA
```

If the largesend is set to 1, then set it to 0 using the following command:

```
# chdev -l SEA -a largesend=0
```

### 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.

### The node may panic after you stop the LLT service and LLT unload is in progress [3333290]

LLT uses the `xmfree()` function of the AIX operating system to free network messages. This function takes a heap as an argument. The heap is created before allocating memory in AIX. In rare cases, during LLT unload, it may happen that LLT destroys this heap while LLT frees messages using the `xmfree()` function. This issue causes LLT to panic the node.

No workaround. You can restart the node and resume normal operations.

## GAB known issues

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

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

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

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

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

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

### Cluster panics during reconfiguration (2590413)

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

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

## I/O fencing known issues

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**To resolve this issue**

1 Rename `cpsadm` to `cpsadmbin`:

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

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

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

3 Change the permissions of the new file to 775:

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

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

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

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

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

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

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

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

When you configure fencing in customized mode and do no provide default port, fencing comes up. However, the `vxfenconfig -l` command output does not list the port numbers.

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

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

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

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

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

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

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

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

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

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

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

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

**Workaround:**

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

## The CP server process, vxcpserv, communicates with client nodes only on those VIPs that are available when CP server process starts (3156922)

When you configure a CP server, the CPSSG service group is configured to manage the vxcpserv process (CP server process) and its dependency (quorum resource). The CP server is managed by a process agent and its dependent virtual IP addresses (VIPs) are managed by a quorum resource. Quorum resource comes online only if it achieves a quorum among the VIPs.

When VCS brings the CPSSG group online, the vxcpserv process listens only on VIPs that come up before the quorum resource is online. The vxcpserv does not listen on these VIPs even if they come up after the quorum resource is online. So, the CP server process serves client nodes only on those VIPs that are available when the CP server process starts.

Note that you can get the list of VIPs on which the vxcpserv process is listening by issuing the netstat command with platform-specific flags.

**Workaround**: Restart CP server configured under the vxcpserv resource using the following commands

```
# hares -offline vxcpserv -sys <system >
```

```
# hares -online vxcpserv -sys <system >
```

where <system> refers to the node where the CPSSG group is online.

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

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

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

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

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

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

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

**Workaround**: Ignore the message.

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

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

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

## Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups [3335137]

If SysDownPolicy of one or more online service groups is configured to AutoDisableNoOffline, fencing configurations such as server-based, disk-based and disable mode fail. Since the service groups is configured with `SysDownPolicy`

= { AutoDisableNoOffline }, stopping VCS fails which leads to the failure of fencing configuration.

**Workaround:** When configuring fencing and before stopping VCS, you must offline the service groups configured with SysDownPolicy = { AutoDisableNoOffline } manually.

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

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

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

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

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

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

**Workaround:** Fix the CP server.

## Symantec Cluster Server agents for Volume Replicator known issues in 6.1

The following are new additional Symantec Cluster Server agents for Volume Replicator known issues in 6.1 release.

### Stale entries observed in the sample main.cf file for RVGLogowner agent [2872047]

Stale entries are found in sample main.cf file for RVGLogowner agent. The stale entries are present in main.cf.seattle file on the RVGLogowner agent which includes CFSQlogckd resource. However, CFSQlogckd is not supported since VCS 5.0.

Workaround: In the cvm group remove the following two lines:

```
CFSQlogckd qlogckd (
             Critical = 0
            )
```

## Issues related to Intelligent Monitoring Framework (IMF)

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

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

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

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

Workaround: No workaround.

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

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

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

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

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

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

## Error messages displayed during reboot cycles [2847950]

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

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

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

Workaround: No workaround.

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

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

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

Workaround: No Workaround.

## 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.

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

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

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

This is an expected behavior and not an issue.

Workaround: No workaround.

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

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

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

For example:

```
# kill -9 -27824
```

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

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

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

Workaround: No workaround.

### AMF may panic the system if it receives a request to unregister an already unregistered resource [3333913]

If AMF encounters any internal error, it unregisters all the resources which it cannot support. During such an event, if any agent calls unregister for one of such resources, AMF may panic the machine.

Workaround: No Workaround.

## 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 77.

## Limitations related to installing and upgrading VCS

### Upgrade of secure clusters not supported using native operating system tools

This release does not support the upgrade of secure clusters using native operating system tools such as Alternate Disk Installation (ADI) and Network Install Manager Alternate Disk Migration (NIMADM).

### Limitation on upgrading to 6.1 on a Symantec Storage Foundation and High Availability cluster

Symantec Storage Foundation (SF) 6.1 requires the AIX operating system to be at 6.1 TL6 or above. To upgrade SF to 6.1 from a release prior to 5.0 MP3 RP5, you must first upgrade SF to the 5.0 MP3 RP5 release. If upgrading to 5.0 MP3 RP5 requires an intermediate operating system upgrade, the operating system level cannot exceed 6.1 TL1. After upgrading to 5.0 MP3 RP5, you must upgrade the operating system to AIX 6.1 TL6, which is the minimum requirement for the 6.1 release. You must upgrade SF to 5.0 MP3 RP5 to avoid a system panic or crash that can occur when a node running AIX 6.1 TL2 or above with a release prior to 5.0 MP3 RP5 is removed from the Symantec Storage Foundation and High Availability cluster. Removing the node causes file system threads to exit. The panic is caused by a check introduced from AIX 6.1 TL2 that validates the lockcount values when a kernel-thread-call exits.

For more information, see the following TechNote:

http://www.symantec.com/docs/TECH67985

## 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/netsvc.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 [2162929]

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

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

## WPAR agent registered to IMF for Directory Online event

The Directory Online event monitors the WPAR root directory. If the parent directory of the WPAR root directory is deleted or moved to another location, AMF does not provide notification to the WPAR agent. In the next cycle of the WPAR monitor, it detects the change and reports the state of the resource as offline.

## Application agent limitations

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

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

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

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

### Live Partition Mobility (LPM) of management LPAR is not supported

Live Partition Mobility (LPM) of management LPAR is not supported.

## Limitations related to VCS engine

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

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

Workaround: No workaround.

### Preferred fencing ignores the forecasted available capacity [3077242]

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

## Limitations related to IMF

- If a process is registered with IMF for offline monitoring, IMF may not detect the process being executed if the length of the process and related arguments exceed 70 characters. In case of ProPCV, IMF may not be able to prevent the process from coming online if the length of the process and related arguments

exceeds 70 characters. This limitation affects Application agent and Process agent. Refer to the *Symantec Cluster Server Bundled Agents Reference Guide* for more information. [2768558]

## 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.

## 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.

## Virtualizing shared storage using VIO servers and client partitions

In an Advanced POWER™ Virtualization (APV) environment, AIX uses the VIO Server to monitor and manage the I/O paths for the virtualized client partitions. At a very high level, the VIO server provides a partition's access to storage that is external to the physical computer. The VIO server encapsulates the physical hardware into virtual adapters called virtual SCSI adapters (server adapter). On

the client side, you can create virtual adapters (client adapters) that map to the server adapter and enable a partition to connect to external storage.

The VIO server provides similar mechanisms to share limited networking resources across partitions. Refer to the manual that came with your system to help set up partitions, and to configure and use the various components such as VIO server and HMC, which are integral parts of IBM's APV environment.

The minimum patch level for using VIO servers with VCS is: version 2.1.3.10-FP-23 and later.

## Supported storage

Refer to the IBM data sheet:

http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html

## Disk Restrictions

When using VCS in combination with VIO servers and their client partitions, you need to ensure that no reservations are placed on the shared storage. This enables client partitions on different systems to access and use the same shared storage.

- If the shared storage is under MPIO control, set the reserve_policy attribute of the disk to no_reserve.
- If the shared storage is not under MPIO control, look up the array documentation to locate a similar attribute to set on the disk.

Internal testing on EMC disks shows that this field maps as the reserve_lock attribute for EMC disks. In this case, setting it to no achieves the same result.

## Accessing the same LUNs from Client Partitions on different Central Electronics Complex (CEC) modules

This section briefly outlines how to set shared storage so that it is visible from client partitions on different CEC modules.

With the VIO server and client partitions set up and ready, make sure that you have installed the right level of operating system on the client partitions, and that you have mapped the physical adapters to the client partitions to provide access to the external shared storage.

To create a shareable diskgroup, you need to ensure that the different partitions use the same set of disks. A good way to make sure that the disks (that are seen from multiple partitions) are the same is to use the disks serial numbers, which are unique.

Run the following commands on the VIO server (in non-root mode), unless otherwise noted.

Get the serial number of the disk of interest:

```
$ lsdev -dev hdisk20 -vpd
 hdisk20
 U787A.001.DNZ06TT-P1-C6-T1-W500507630308037C-
 L401 0401A00000000  IBM FC 2107

 Manufacturer.............IBM
 Machine Type and Model......2107900
 Serial Number..............7548111101A
 EC Level.....................131
 Device Specific.(Z0).........10
 Device Specific.(Z1).........0100
 …
```

Make sure the other VIO server returns the same serial number. This ensures that you are viewing the same actual physical disk.

List the virtual SCSI adapters.

```
$ lsdev -virtual | grep vhost
 vhost0  Available  Virtual SCSI Server Adapter
 vhost1  Available  Virtual SCSI Server Adapter
```

**Note:** Usually vhost0 is the adapter for the internal disks. vhost1 in the example above maps the SCSI adapter to the external shared storage.

Prior to mapping hdisk20 (in the example) to a SCSI adapter, change the reservation policy on the disk.

```
$ chdev -dev hdisk20 -attr reserve_policy=no_reserve
hdisk20 changed
```

For hdisk20 (in the example) to be available to client partitions, map it to a suitable virtual SCSI adapter.

If you now print the reserve policy on hdisk20 the output resembles:

```
$ lsdev -dev hdisk20 attr reserve_policy
 value
 no_reserve
```

Next create a virtual device to map hdisk20 to vhost1.

```
$ mkvdev -vdev hdisk20 -vadapter vhost1 -dev mp1_hdisk5
 mp1_hdisk5 Available
```

Finally on the client partition run the cfgmgr command to make this disk visible via the client SCSI adapter.

You can use this disk (hdisk20 physical, and known as mp1_hdisk5 on the client partitions) to create a diskgroup, a shared volume, and eventually a shared file system.

Perform regular VCS operations on the clients vis-a-vis service groups, resources, resource attributes, etc.

# Cluster Manager (Java console) limitations

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

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

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

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

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

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

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

## VCS Simulator does not support I/O fencing

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

## 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.

## The operating system does not distinguish between IPv4 and IPv6 packet counts

In a dual-stack configuration, when you use packet counts and the IPv6 network is disabled, the NIC agent might not detect a faulted NIC. It might not detect a fault because while the IPv6 network is down its packet count still increases. The packet count increases because the operating system does not distinguish between the packet counts for IPv4 and IPv6 networks. The agent then concludes that the NIC is up. If you are using the same NIC device for IPv4 as well as IPv6 resources, set PingOptimize to 0 and specify a value for the NetworkHosts attribute for either the IPv6 or the IPv4 NIC resource. [1061253]

## A service group that runs inside of a WPAR may not fail over when its network connection is lost

For a WPAR configuration when the WPAR root is on NFS, the WPAR service group may not fail over if the NFS connection is lost. This issue is due to an AIX operating system limitation. [1637430]

## Limitations related to LLT

This section covers LLT-related software limitations.

### LLT over IPv6 UDP cannot detect other nodes while VCS tries to form a cluster (1907223)

LLT over IPv6 requires link-local scope multicast to discover other nodes when VCS tries to form a cluster. If multicast networking is undesirable, or unavailable in your environment, use the address of the peer nodes to eliminate the need for the multicast traffic.

Workaround: Add the set-addr entry for each local link into the /etc/llttab file. You add the entry to specify the address of the peer nodes that are available on the corresponding peer links. For example, you add the following lines into the llttab file to specify the set-addr entry for a node. In this example, the node's IPv6 address is fe80::21a:64ff:fe92:1d70.

```
set-addr 1 link1 fe80::21a:64ff:fe92:1d70
set-arp 0
```

### LLT does not start automatically after system reboot (2058752)

After you reboot the systems, if you had not completed the terminal setting procedure, LLT does not start automatically and does not log any error messages. You can manually start LLT using the `/etc/init.d/llt.rc` command.

If you reinstall a system, when the system reboots a message appears on the system console to set the terminal setting if you have not already done so. LLT does not start until you complete the terminal setting procedure.

**Workaround: To resolve the LLT startup issue**

1. After you reboot the systems, open the system console using any available method, for example, from HMC.
2. On the console, go to the terminal setting menu, and set the terminal of your choice.
3. Select the **Task Completed** menu option.

## 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.

### Limitation with RDAC driver and FAStT array for coordinator disks that use raw disks

For multi-pathing to connected storage, AIX uses the RDAC driver for FAStT arrays. Since it is an active/passive array, only the current active path is exposed to clients. The I/O fencing driver, vxfen, can use only a single active path and has no

foreknowledge of the passive paths to the coordinator disks on an array. If the single active path fails, all nodes in the cluster lose access to the coordinator disks.

The loss of the path to the coordinator disks can potentially go unnoticed until a reboot, split brain, or any other reason that leads to a cluster membership change occurs. In any of these conditions, the cluster cannot form, and all nodes panic to prevent data corruption. No data loss occurs.

Workaround: Use DMP and specify paths to coordinator disks as DMP paths rather than raw disks to avoid this limitation.

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

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

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

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

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

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

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

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

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

### Symantec Cluster Server documentation

Table 1-12 lists the documents for Symantec Cluster Server.

**Table 1-12** Symantec Cluster Server documentation

| Title | File name | Description |
|---|---|---|
| *Symantec Cluster Server Release Notes* | vcs_notes_61_aix.pdf | Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product. |
| *Symantec Cluster Server Installation Guide* | vcs_install_61_aix.pdf | Provides information required to install the product. |
| *Symantec Cluster Server Administrator's Guide* | vcs_admin_61_aix.pdf | Provides information required for administering the product. |
| *Symantec Cluster Server Bundled Agents Reference Guide* | vcs_bundled_agents_61_aix.pdf | Provides information about bundled agents, their resources and attributes, and more related information. |
| *Symantec Cluster Server Agent Developer's Guide*<br>(This document is available online only.) | vcs_agent_dev_61_unix.pdf | Provides information about the various Symantec agents and procedures for developing custom agents. |
| *Symantec Cluster Server Agent for DB2 Installation and Configuration Guide* | vcs_db2_agent_61_aix.pdf | Provides notes for installing and configuring the DB2 agent. |
| *Symantec Cluster Server Agent for Oracle Installation and Configuration Guide* | vcs_oracle_agent_61_aix.pdf | Provides notes for installing and configuring the Oracle agent. |
| *Symantec Cluster Server Agent for Sybase Installation and Configuration Guide* | vcs_sybase_agent_61_aix.pdf | Provides notes for installing and configuring the Sybase agent. |

## Symantec Storage Foundation and High Availability Solutions products documentation

Table 1-13 lists the documentation for Symantec Storage Foundation and High Availability Solutions products.

**Table 1-13** Symantec Storage Foundation and High Availability Solutions products documentation

| Document title | File name | Description |
|---|---|---|
| *Symantec Storage Foundation and High Availability Solutions—What's new in this release*<br>(This document is available online.) | sfhas_whats_new_61_unix.pdf | Provides information about the new features and enhancements in the release. |
| *Symantec Storage Foundation and High Availability Solutions Getting Started Guide* | getting_started.pdf | Provides a high-level overview of installing Symantec products using the Veritas script-based installer. The guide is useful for new users and returning users that want a quick refresher. |
| *Symantec Storage Foundation and High Availability Solutions Solutions Guide* | sfhas_solutions_61_aix.pdf | Provides information about how SFHA Solutions product components and features can be used individually and in concert to improve performance, resilience and ease of management for storage and applications. |
| *Symantec Storage Foundation and High Availability Solutions Virtualization Guide*<br>(This document is available online.) | sfhas_virtualization_61_aix.pdf | Provides information about Symantec Storage Foundation and High Availability support for virtualization technologies. Review this entire document before you install virtualization software on systems running SFHA products. |
| *Symantec Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide*<br>(This document is available online.) | sfhas_dr_impl_61_aix.pdf | Provides information on configuring campus clusters, global clusters, and replicated data clusters (RDC) for disaster recovery failover using Storage Foundation and High Availability Solutions products. |
| *Symantec Storage Foundation and High Availability Solutions Troubleshooting Guide* | sfhas_tshoot_61_aix.pdf | Provides information on common issues that might be encountered when using Symantec Storage Foundation and High Availability Solutions and possible solutions for those issues. |

## Symantec ApplicationHA documentation

Table 1-14 lists the documentation for Symantec ApplicationHA.

**Table 1-14** Symantec ApplicationHA documentation

| Document title | File name | Description |
|---|---|---|
| *Symantec ApplicationHA Release Notes* | applicationha_notes_61_lpar_aix.pdf | Describes the new features and software and system requirements. This document also contains a list of limitations and issues known at the time of the release. |
| *Symantec ApplicationHA Installation Guide* | applicationha_install_61_lpar_aix.pdf | Describes the steps for installing and configuring Symantec ApplicationHA. Some of the most common troubleshooting steps are also documented in this guide. |
| *Symantec ApplicationHA User's Guide* | applicationha_users_61_lpar_aix. pdf | Provides information about configuring and managing Symantec ApplicationHA in Logical Partition (LPAR) virtualization environments. Some of the most common troubleshooting steps are also documented in the guide. |
| *Symantec ApplicationHA Agent for Oracle Configuration Guide* | applicationha_oracle_agent_ 61_lpar_aix.pdf | Describes how to configure application monitoring for Oracle. |
| *Symantec ApplicationHA Generic Agent Configuration Guide* | applicationha_gen_agent_61_ lpar_aix.pdf | Describes how to configure application monitoring for a generic application. |
| *Symantec ApplicationHA Agent for DB2 Configuration Guide* | applicationha_db2_agent_61_ lpar_aix.pdf | Describes how to configure application monitoring for DB2. |
| *Symantec ApplicationHA Agent for Apache HTTP Server Configuration Guide* | applicationha_apache_agent_61_ lpar_aix.pdf | Describes how to configure application monitoring for Apache HTTP Server. |

Veritas Operations Manager (VOM) is a management tool that you can use to manage Symantec Storage Foundation and High Availability Solutions products. If you use VOM, refer to the VOM product documentation at:

https://sort.symantec.com/documents

## Manual pages

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

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

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

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

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

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

See the `man`(1) manual page.

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

https://sort.symantec.com/documents