Symantec™ Cluster Server 6.2 Release Notes - Solaris

August 2015

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: 6.2

Document version: 6.2 Rev 1

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_apj@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 6.2 for Solaris. 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: 6.2 Rev 1" 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 website 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* (6.2)

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 Oracle Solaris Studio. The following is the layout of libvcsagfw.so in usr/lib:

```
/usr/lib/libvcsagfw.so --> . /libvcsagfw.so.2
```

If you use custom agents compiled on older compilers, the agents may not work with VCS 6.2. If your custom agents use scripts, continue linking to ScriptAgent. Use Script50Agent for agents written for VCS 5.0 and above.

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. ■ Analyze systems to determine if they are ready to install or upgrade Symantec products and generate an Installation and Upgrade custom report. ■ 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. ■ 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. ■ 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. ■ 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. ■ 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. ■ List and download Symantec product documentation including product guides, manual pages, compatibility lists, and support articles. ■ 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. ■ 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/TECH225259
- 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/TECH225258

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 6.2

This section lists the changes in Symantec Cluster Server 6.2.

Release level terminology changes

With the 6.2 release, terms that are used to describe patch-based releases have changed as follows:

Table 1-1 Release level terminology changes

Pre 6.0.1	6.0.x, 6.1, 6.1.x	6.2 and forward	Status	Available from
P-Patch	Public hot fix	Patch	Official	SORT
Hot fix	Private hot fix	Hot fix	Unofficial	Customer support

Official patch releases are available from SORT. This release was previously referred to as a P-Patch or a Public hot fix and is now referred to as a Patch. Unofficial patch releases are available from customer support. Hot fix is the only unofficial patch release.

Cluster Manager Java GUI support consideration

The Cluster Manager Java GUI is End of Life but continues to be supported by Symantec to ensure that it works with the core clustering solutions of VCS and ApplicationHA in Linux and Windows environments. The Java GUI remains available for download with support for all VCS features available in pre-6.0 releases. Customers can manage service groups, generate new configurations, and perform other traditional cluster management operations. The Java GUI will be supported only on the Linux and Windows platforms.

Additional feature capabilities and platform support added in VCS 6.0 and later releases are available exclusively through Veritas Operations Manager (VOM). Symantec recommends the use of VOM to manage clusters and for all advanced capabilities.

Attributes introduced in VCS 6.2

The following section describes the attributes introduced in VCS 6.2.

Cluster level attributes

DefaultGuestAccess	Enables guest access for any authenticated user to the secure cluster.
GuestGroups	Contains a list of user groups that have guest access.

LDom agent attributes

Meter	Defines the meters based on which the failover decision is taken for the service group containing the LDom resource.
MeterControl	Defines the interval after which meter entry point should get called.
MeterTimeout	The maximum time for the meter entry point to complete.
AvailableMeters	Defines the meters that the agent supports.
MeterRetryLimit	Defines the number of times the meter operation can be retried before it succeeds.
MeterRegList	It is an ordered list of attributes. If MeterRegList attribute or any attribute that are defined in MeterRegList is changed then the meter entry point is called immediately.

SFCache agent attributes

CacheArea	Specifies the name of the cache area.
CacheMode	Specifies the caching mode.
CacheFaultPolicy	Specifies the action to be performed in case of a cache fault.
CacheObjectName	Specifies the cache object name; it can be a mount point or disk group/volume.
FaultOnMonitorTimeouts	Defines whether VCS interprets the Monitor timeout as a resource fault. By default, the FaultOnMonitorTimeouts attribute is set to 4, but the SFCache agent overrides this value and sets it to 0.
NumThreads	Number of threads that are used within the agent process for managing resources. This number does not include the number of threads that are used for other internal purposes.

IP agent attributes

IpadmIfProperties	Interface properties for the `ipadm set-ifprop` command. IpadmIfProperties attribute is applicable for Solaris 11 only. On Solaris 10, this attribute value is ignored.
IpadmAddrProperties	Address properties for the `ipadm set-addrprop` command. IpadmAddrProperties attribute is applicable for Solaris 11 only. On Solaris 10, this attribute value is ignored.

Resource level attributes

Utilization	The virtual machine agent meters the CPU and memory requirement of the virtual machine and populates this attribute value for the virtual machine resource.

System attributes

ServerAvailableCapacity	The HostMonitor agent meters the free CPU and memory on the physical server which are available to other virtual machines and populates this attribute.
ServerAvailableForecast	The HostMonitor agent forecasts the free CPU and memory of the physical server and populates this attribute value.
ServerCapacity	The HostMonitor agent meters the total CPU and memory of the physical server and populates this attribute value.
ServerReservedCapacity	This is an internal attribute populated by the VCS engine.

Oracle agent attributes:

PDBName	This attribute must be configured for pluggable database (PDB) and the value should be set for a PDB database name. Do not set this attribute for traditional and container (CDB) database.

Mount agent attributes:

CacheRestoreAccess	Determines whether to perform restore access operation or not. This attribute is applicable only if: ■ File system type is VxFS. ■ Writeback caching is enabled for the SmartIO feature.

Changes related to installation and upgrades

The product installer includes the following changes in 6.2.

Connecting to the SORT website through a proxy server

The product installer connects to the Symantec Operations Readiness Tools (SORT) website for several purposes, such as downloading the latest installer patches, and uploading installer logs. Deployment Server can connect to SORT to automatically download Maintenance or Patch release images. In this release, before running the product installer or Deployment Server, you can use the following proxy settings to connect to SORT through proxy servers:

```
# https_proxy=http://proxy_server:port
# export https_proxy
# ftp_proxy=http://proxy_server:port
# export ftp_proxy
```

Symantec Cluster Server gets installed in secure mode by default

Symantec Cluster Server gets installed in secure mode by default. You are advised to install VCS in secure mode to be able to control guest user access to secure clusters and encrypt communication between VCS components. You can choose

the non-secure mode during installation; however, the product installer warns you during the installation with the following message:

```
Symantec recommends that you install the cluster
in secure mode. This ensures that communication between
cluster components is encrypted and cluster information
is visible to specified users only.
```

The upgrade from non-secure mode continues to happen in non-secure mode. The upgrade from secure mode advises you to control user access to secure clusters.

Upgrade Symantec Cluster Server online, while keeping your applications online

You can perform an online upgrade of Symantec Cluster Server using the installer, while keeping your applications online. Your applications can run seamlessly when the upgrade is in progress. The upgrade behavior otherwise remains same as the one in the previous release. Note that the application monitoring does not happen as long as the upgrade is in progress.

Support for upgrading VCS using the web-based installer for Solaris 10 Live Upgrade

You can use the Symantec web-based installer to upgrade VCS as part of the Live Upgrade.

On a node in the cluster, run the web-based installer on the DVD to upgrade VCS on all the nodes in the cluster.

The program uninstalls the existing version of VCS on the alternate boot disk during the process. At the end of the process, VCS 6.2 is installed on the alternate boot disk.

Support for setting up ssh and rsh connection using the pwdutil.pl utility

The password utility, `pwdutil.pl`, is bundled in the 6.2 release under the scripts directory. The users can run the `pwdutil.pl` utility to set up the `ssh` and `rsh` connection automatically.

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 platform (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 1-2.

Table 1-2 Deployment Server functionality

Feature	Description
Install or Upgrade systems with Install Bundle and Install Template	■ Install or upgrade systems with an Install Bundle. ■ Install packages on systems based on the information stored in the Install Template.
Define or modify Install Bundles	Define or modify Install Bundles and save them using the Deployment Server.
Create Install Templates	Discover installed components on a running system that you want to replicate on new systems.
Connecting the Deployment Server to SORT using a proxy server	Use a proxy server, a server that acts as an intermediary for requests from clients, for connecting the Deployment Server to the Symantec Operations Readiness Tools (SORT) website.
Platform Filtering	On the Set Preference menu, choose Selected Platforms to filter the platforms that are currently being used in the deployment environment.

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

See the *Installation Guide* for more information.

Changes to VCS bundled agents

This section describes changes to the bundled agents for VCS.

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

SFCache agent

SFCache agent is a new agent introduced in this release. The SFCache agent enables, disables, and monitors cache. In case of a cache fault, the application still

runs without any issues on the very same system, but with degraded I/O performance. Considering this, the SFCache agent provides an attribute to control the agent behavior. You can either choose to receive “IGNORE” or initiate “FAILOVER” in case of cache fault.

For more information, see *Bundled Agents Reference Guide*.

New agent function for the LDom agent

In this release, the meter entry point has been introduced. This entry point measures the VCPU and Memory requirement of the LDom based on the keys in the Meters attribute.

Changes in Zpool agent

Zpool agent does not monitor the ZFS (Zettabyte File Systems) which have mount point property set to none or canmount property set to off. The agent does not monitor the file systems even when ChkZFSMounts attribute is set to 1. The agent considers these values as intentionally set and does not display a warning in such scenarios.

Solaris 11: `ipadm` command support for IP and NIC agents

On Solaris 11, IP and NIC agents now support the `ipadm` command. IP and NIC agents have been enhanced for performing online, offline, and monitor operations using the `ipadm` command.

For more information, refer to the *Bundled Agents Reference Guide*.

Coordpoint agent

The Coordpoint agent is always in the online state when the I/O fencing is configured in the majority or the disabled mode. For both these modes the I/O fencing does not have any coordination points to monitor. Thereby, the Coordpoint agent is always in the online state.

Changes to the VCS engine

AdaptiveHA enhancement (virtual machine service group)

In a virtualized environment, VCS monitors and forecasts the available capacity of the physical server in terms of SCPU and SMem. For a virtual machine service group (VMSG) if you set FailOverPolicy (service group attribute) to BiggestAvailable, AdaptiveHA enables VCS to dynamically select the biggest available target physical server to online, switch, and failover the VMSG.

The following new attributes have been introduced:

- ServerAvailableCapacity
- ServerAvailableForecast
- ServerCapacity
- ServerReservedCapacity
- Utilization

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

New environment variables

In this release, the following VCS environment variables have been introduced:

- VCS_CONN_INIT_QUOTA
- VCS_CONN_HANDSHAKE_TIMEOUT

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

Atleast resource dependency

A new type of resource dependency has been introduced in this release wherein a parent resource can depend on a set of child resources. The parent resource is brought online or remains online only if a minimum number of child resources in this resource set are online. The system creates a set of child IP resources and the application resource depends on this set.

For example, if an application depends on five IPs and if this application has to be brought online or has to remain online, at least two IPs must be online. If two or more IP resources come online, the application attempts to come online. If the number of online resources falls below the minimum requirement, resource fault is propagated up the resource dependency tree.

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

Changes to VCS agent framework

The following changes are introduced to the VCS agent framework.

Enhanced failure data capturing process to better troubleshoot unexpected resource behavior

The first failure data capture (FFDC) of unexpected events is extended to resource level to record VCS events. The event logs generated during such events will be logged in their respective agent log files for easier viewing of the logs.

The examples of such events are:

- When resource faults and times out of its entry point.
- When resource moves to UNKNOWN state.
- When any resource operation, such as ONLINE or OFFLINE, fails.
- When resource is detected as ONLINE or OFFLINE for the first time.

Log unification of VCS agent entry points

The logs of VCS agent entry points implemented using script and C/C++ language will be logged under the respective agent log file. Prior to VCS 6.2, these logs were shared between the engine and the agent log files. For example, logs of script based agents used to go into engine logs and logs of C/C++ based agent into agent log file. This change will reduce the load on VCS engine process and remove the clutter from the engine log.

To restore the pre-VCS 6.2 logging behavior, you can change the LogViaHalog attribute value to 1 (default 0).

Note: Log messages of one or more script entry points which are executed inside the container will continue to be logged in the engine log file.

FireDrill attribute can be overridden

The FireDrill attribute is modified such that it can be overridden in VCS 6.2 or later releases.

AdaptiveHA is extended for Solaris LDom

The CPU and Memory requirement of LDom is metered and used in the BiggestAvailable failover policy for service groups containing LDom resource.

Changes to LLT, GAB, and I/O fencing

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

Changes to I/O fencing

Symantec Cluster Server (VCS) includes the following changes to I/O fencing in 6.2:

I/O fencing supports majority-based fencing mechanism, a new fencing mechanism that does not need coordination points

I/O fencing supports a new fencing mode called majority-based I/O fencing. Majority-based I/O fencing mode does not need coordination points to provide protection against data corruption and data consistency in a clustered environment. Use majority-based I/O fencing when there are no additional servers and or shared SCSI-3 disks to be used as coordination points. It provides a reliable arbitration method and does not require any additional hardware setup, such as CP Servers or shared SCSI3 disks.

In the event of a network failure, the majority sub-cluster wins the fencing race and survives the race. Note that even if the majority sub-cluster is hung or unresponsive, the minority sub-cluster loses the fencing race and the cluster panics. The cluster remains unavailable till the issue is resolved.

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

Clear coordination point server registrations using the vxfenclearpre utility

The vxfenclearpre utility is enhanced to clear registrations from coordination point servers for the current cluster in addition to the existing functionality to remove SCSI3 registrations and reservation keys from the set of coordinator disks and shared data disks. The local node from where you run the utility must have the UUID of the current cluster at `/etc/vx/.uuids` directory in the `clusuuid` file.

Note that you may experience delays while clearing registrations on the coordination point servers because the utility tries to establish a network connection with IP addresses used by the coordination point servers. The delay may occur because of a network issue or if the IP address is not reachable or is incorrect.

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

Raw disk I/O fencing policy is not supported

Symantec does not support raw disk policy for I/O fencing. Use DMP as the I/O fencing policy for coordinator disks that have either a single hardware path or multiple hardware paths to nodes.

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

Changes to wizard support

New VCS configuration wizards introduced on Linux and UNIX

VCS Cluster Configuration Wizard and Symantec High Availability Configuration Wizard are introduced on all supported Linux and UNIX distributions in this release.

The two new wizards replace the Symantec High Availability Configuration Wizard that earlier provided a combined workflow for cluster configuration and application (high availability) configuration, and was supported only on Linux.

You can launch the wizards from the Symantec High Availability view. You can continue to access the view as required from Veritas Operations Manager, VMware vSphere Client, or a browser.

For more information, see the *Symantec Cluster Server Installation Guide* and *Symantec Cluster Server Administrator's Guide.* For steps to use the Symantec High Availability wizard, see the application-specific VCS agent installation and configuration guides. For VMware specific information on this feature, see the *Symantec High Availability Solutions Guide for VMware*.

Changes to database agents

Changes to the Oracle agent

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

VCS agent for Oracle supports management of container and pluggable databases

VCS supports the multitenant architecture introduced in Oracle 12c Release 1 (12.1). The multitenant architecture enables Oracle database to function as a multitenant container database (CDB) and one or many customized pluggabble databases (PDBs).

Note: IMF monitoring is not supported in a PDB resource.

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

Changes introduced in Symantec Cluster Server 6.1.1

The following changes were introduced in VCS 6.1.1.

Changes in Zpool agent

The Zpool agent doesn't monitor the ZFS (Zettabyte File System) file systems with the mountpoint property set to none or canmount set to off. The agent doesn't monitor those file systems even though the ChkZFSMounts attribute is 1. Because

the agent considers users choose these file systems intentionally, it doesn't throw warning in such scenarios

Changes to hazonesetup and hazoneverify utilities

The `hazonesetup` command is able to create multiple zone resources in a service group. For more details, refer to the `hazonesetup` man page.

The `hazoneverify` command is able to verify multiple zone resources in a service group.

VCS system requirements

This section describes system requirements for VCS.

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

VCS requires that all nodes in the cluster use the same processor architecture and run the same operating system.

For example, in a cluster with nodes running Solaris, all nodes must run Solaris SPARC.

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 23.

See “Supported Solaris operating systems ” on page 23.

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 Solaris 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
Solaris 10	Update 9, 10, and 11	SPARC
Solaris 11	Support for Oracle Solaris 11.2 and Support Repository Updates (SRUs) up to 11.2.6.5. Solaris 11.1 and up to Support Repository Update (SRU) 11.1.21.0.4.1 Solaris 11.2 and up to Support Repository Update (SRU) 11.2.2.0.8.0	SPARC

This release is not supported on the x86-64 architecture.

This release supports native brand zones on Solaris 10 operating system and solaris brand and solaris10 brand zones on the Solaris 11 operating system. This release does not support the Kernel Zones feature of Solaris 11 Update 2.

Supported software for VCS

VCS supports the following versions of Symantec Storage Foundation:

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

Oracle Solaris 11

- Storage Foundation 6.2
  - VxVM 6.2 with VxFS 6.2
- Storage Foundation 6.1
  - VxVM 6.1 with VxFS 6.1

Oracle Solaris 10

- Storage Foundation 6.2
  - VxVM 6.2 with VxFS 6.2
- Storage Foundation 6.1
  - VxVM 6.1 with VxFS 6.1

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 Oracle VM Server for SPARC

Supported Oracle VM Server for SPARC versions are OVM 2.0, OVM 2.1, OVM 2.2, OVM 3.0, and OVM 3.1.

For supported OS version for Oracle VM Server for SPARC, refer to *Oracle VM server for SPARC Release Notes*.

The version of the Oracle Solaris operating system (OS) that runs on a guest domain is independent of the Oracle Solaris OS version that runs on the primary domain. Therefore, if you run the Oracle Solaris 10 OS in the primary domain, you can still run the Oracle Solaris 11 OS in a guest domain. Likewise if you run the Oracle Solaris 11 OS in the primary domain, you can still run the Oracle Solaris 10 OS in a guest domain.

The only difference between running the Oracle Solaris 10 OS or the Oracle Solaris 11 OS on the primary domain is the feature difference in each OS.

Supported Solaris operating systems for CP server

Table 1-4 Supported Solaris OS versions for CP server

Operating systems	Levels	Chipsets
Solaris 10	Update 9, 10, and 11	SPARC
Solaris 11	Solaris 11.1 and up to Supported Repository Update (SRU) 11.1.21.0.4.1 Solaris 11.2 and up to SRU 11.2.2.0.8.0	SPARC

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:

- Raw disk I/O fencing policy is no longer supported.

Cluster Manager Java GUI support consideration

The Cluster Manager Java GUI is End of Life but continues to be supported by Symantec. The Java GUI remains available for download and use within clusters with support for all VCS features available in pre-6.0 releases. Customers can manage service groups, generate new configurations, and perform other traditional cluster management operations. Symantec supports the Java GUI only on the Linux and Windows platforms.

Additional feature capabilities and platform support added in VCS 6.0 and later releases are available exclusively through Veritas Operations Manager (VOM). Symantec recommends the use of VOM to manage clusters and for all advanced capabilities.

Deprecated attributes

The following table lists the attributes deprecated in this release.

Table 1-5 Attributes deprecated in this release

Attribute name	Agent type
SecondLevelMonitor	Apache **Note:** The SecondLevelMonitor attribute is deprecated in VCS 6.2. Instead, LevelTwoMonitorFreq attribute at the Apache resource type level may be used

Table 1-5 Attributes deprecated in this release *(continued)*

Attribute name	Agent type
ResLogLevel	Apache **Note:** Use type level attribute LogDbg to enable debug logs. Set LogDbg attribute to DBG_5 to enable debug logs for Apache agent. By default setting the LogDbg attribute to DBG_5, enables the debug logs for all Apache resources in the cluster. If specific Apache resource needs to be enabled for debug logs override LogDbg attribute.
DetailMonitor	Oracle, Sybase **Note:** If you manually upgrade VCS to 6.2 with detail monitoring enabled in the previous version, set the value of LevelTwoMonitorFreq attribute to that of DetailMonitor.
AgentDebug	DB2udb

Fixed issues

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

LLT, GAB, and I/O fencing fixed issues

Table 1-6 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-6 LLT, GAB, and I/O fencing fixed issues

Incident	Description
3156922	The CP server process, vxcpserv, communicates with client nodes only on those VIPs that are available when CP server process starts.
3335137	Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups.
3473104	When virtual NICs are configured under LLT without specifying the MTU size 1500 in lltab, cluster does not function properly. For example, VCS engine commands may hang and print below message in the engine logs: `VCS CRITICAL V-16-1-51135 GlobalCounter not updated`

Table 1-6 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
3548629	On Solaris 11, LLT, GAB and I/O fencing modules fails to configure when installed on an alternate boot environment.
3331801	SMF services for VCS kernel components may go into maintenance state when installed in a new boot environment.
3031216	The dash (-) in a disk group name causes vxfentsthdw(1M) and Vxfenswap(1M) utilities to fail.
3471571	Cluster nodes may panic if you stop the HAD process by force on a node and reboot that node.
3532859	The Coordpoint agent monitor fails if the cluster has a large number of coordination points.

Installation related fixed issues

Table 1-7 Installation related fixed issues

Incident	Description
3325954	On Solaris 10, xprtld will not be started if you use jumpstart to install product.
3326639	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.

VCS engine fixed issues

Table 1-8 lists the fixed issues for VCS engine.

Table 1-8 VCS engine fixed issues

Incident	Description
3381042	The checkboot utility core dump or time difference between a system and Network Time Protocol (NTP) time leads to unexpected deletion of the temporary files. The deletion causes the VCS agents to report an incorrect state.
3448510	The `hastatus` command fails and dumps core when it is run from a local zone.

Table 1-8 VCS engine fixed issues *(continued)*

Incident	Description
3468891	Inconsistencies in `/etc/VRTSvcs/conf/attributes/cluster_attrs.xml` file and hide resource-level ContainerInfo attribute from `hares -display` command.
3211834	CurrentLimits attribute value is not updated correctly when a service group faults.
3385820	Sometimes the high availability daemon (HAD) crashes if it runs for a long duration.
3436617	When invoking triggers if some arguments are left unassigned, the `hatrigger` command fails due to compilation errors.
3471819	The service group fails to go online if the CurrentCount attribute value is incorrect.
3498072	The hazonesetup (1M) utility reports a Perl warning message when the locale is set to non-English.
3464981	The file size of Engine_A.log file could not be increased beyond 32 MB.
3580940	VCS configuration becomes unclean when incorrect filename is provided with the `ha` command for SourceFile attribute.
3603275	HAD and other nodes abort while shutting down two nodes simultaneously.

Bundled agents fixed issues

Table 1-9 lists the fixed issues for bundled agents.

Table 1-9 Bundled agents fixed issues

Incident	Description
2490296	Application agent cannot handle a case with user as root, envfile set and shell as `csh`.
2618482	Some agents may fail to come online after full upgrade to VCS 6.0 if they were online before the upgrade.

Table 1-9 Bundled agents fixed issues *(continued)*

Incident	Description
3536195	The Online operation of LDom resource fails when Disaster Recovery options are configured. This is because the online entry point fails to get the definition of a function.
3326591	The IPMultiNICB agent delays bringing IPv4 VIP online on Solaris 10 by 5 seconds.
3621042	NIC agent might incorrectly report NIC resource state as offline.
3590419	On Solaris, the MultiNICB agent does not add the default route when the resource state is ONLINE in the first probe.
3593137	Mount agent unmounts the incorrect file system due to incorrect pattern matching.
3422904	Zone agent fails to bring the Zone resource to online state when the state of the zone is unavailable.
3576701	The Share resource fails to unshare the path name when file system is unmounted.
3505202	VCS does not support non-default value for VCS_LOG environment variable.

Fixed issues related to AMF

Table 1-10 AMF fixed issues

Incident	Description
2848007	The libvxamf library encounters an error condition while doing a process table scan.
3333913	AMF may panic the system if it receives a request to unregister an already unregistered resource.
3407338	If one of the events of a group is in triggered state, then the group fails to unregister because of the triggered event.
3606494	During a branded zone boot operation, under certain circumstances the system panics because of memory corruption in Asynchronous Monitoring Framework.

Table 1-10 AMF fixed issues *(continued)*

Incident	Description
3338946	Sometimes when a process offline registration is requested and the system is under heavy load, AMF library fails to verify whether the resource is actually offline. As a result, registration fails.

Cluster configuration wizard fixed issues

Table 1-11 lists the cluster configuration wizard fixed issues.

Table 1-11 Wizard fixed issues

Incident	Description
3593390	In a Solaris 11 setup, if you execute the deprecated 'pkginfo' command for the VRTSwiz package, an incorrect error message appears, stating that VRTSwiz information cannot be found.

Known issues

This section covers the known issues in this release.

Issues related to installing and upgrading VCS

On Solaris 11, when you install the operating system together with SFHA products using Automated Installer, the local installer scripts do not get generated. (3640805)

On Solaris 11, when you use Automated Installer (AI) to install the Solaris 11 operating system together with SFHA products, the local installer scripts fail to get generated.

Workaround:

On the target system(s), execute the following script:

```
/opt/VRTSsfcpi62/bin/run-once
```

`installer -requirements` does not list RHEL 6 Update 6 and Oracle Linux 7 as supported platforms (3657260)

The `installer -requirements` command does not list RHEL 6 Update 6 and Oracle Linux 7 as supported platforms though they are qualified with version 6.2.

Workaround: The correct supported list is mentioned in the latest version of the product Release Notes. See the latest Release Notes on the Symantec website for the updated list.

https://sort.symantec.com/documents

Node panics after upgrade from Solaris 11 to Solaris 11.1 on systems running version 6.0.1 or earlier (3560268)

Nodes running version 6.0.1 or earlier panic after you upgrade the operating system from Solaris 11 to Solaris 11.1. This is due to changes introduced in the Solaris operating system.

Workaround: Perform the following steps during the operating system upgrade from Solaris 11 to Solaris 11.1 before you boot to the Solaris 11.1 boot environment. This will prevent the product from starting on the Solaris 11.1 boot environment.

Open the file `/etc/default/llt` on the new boot environment and set `LLT_START` to 0.

Open the file `/etc/default/gab` on the new boot environment and set `GAB_START` to 0

Open the file `/etc/default/amf` on the new boot environment and set `AMF_START` to 0

Open the file `/etc/default/vxfen` on the new boot environment and set `VXFEN_START` to 0

After the operating system is upgraded to Solaris 11.1, upgrade the product to a version that support Solaris 11.1.

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
```

Upgrade or uninstallation of VCS may encounter module unload failures

When you upgrade or uninstall VCS, some modules may fail to unload with error messages similar to the following messages:

```
llt failed to stop on node_name
gab failed to stop on node_name
```

The issue may be observed on any one or all the nodes in the sub-cluster.

Workaround: After the upgrade or uninstallation completes, follow the instructions provided by the installer to resolve the issue.

Erroneous resstatechange trigger warning [2277819]

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.

Installing VRTSvlic package on Solaris system with local zones displays error messages [2555312]

If you try to install VRTSvlic package on a Solaris system with local zones in installed state, the system displays the following error messages:

```
cp: cannot create /a/sbin/vxlicinst: Read-only file system
cp: cannot create /a/sbin/vxlicrep: Read-only file system
cp: cannot create /a/sbin/vxlictest: Read-only file system
```

Workaround: On the Solaris system, make sure that all non-global zones are started and in the running state before you install the VRTSvlic package.

Installing VRTSvlic package during live upgrade on Solaris system non-global zones displays error messages [3623525]

While installing VRTSvlic package during live upgrade on Solaris system with non-global zones following error messages are displayed:

```
cp: cannot create /a/sbin/vxlicinst: Read-only file system
cp: cannot create /a/sbin/vxlicrep: Read-only file system
cp: cannot create /a/sbin/vxlictest: Read-only file system
```

Workaround: This message can be ignored. The vxlicinst, vxlicrep, vxlictest utilities are present in `/opt/VRTSvlic/sbin/` inside a non-global zone.

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.

On Solaris 10, a flash archive installed through JumpStart may cause a new system to go into maintenance mode on reboot (2379123)

If a Flash archive is created on a golden host with encapsulated root disks, when this Flash archive is installed onto another host through JumpStart, the new system may go to maintenance mode when you initially reboot it.

This problem is caused by the predefined root disk mirror in the Flash archive. When the archive is applied to a clone system, which may have different hard drives, the newly cloned system may get stuck at root disk mirroring during reboot.

Workaround: Create the Flash archive on a golden host with no encapsulated root disks. Run `vxunroot` to clean up the mirrored root disks before you create the Flash archive.

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.

VCS Zone users must be added after upgrade to VCS 6.0 or later

If you upgrade your configuration containing Zone resources to VCS 6.0 or later from:

- VCS 5.1SP1RP1 or later VCS releases with DeleteVCSZoneUser attribute of Zone agent set to 1
- VCS 5.1SP1 or ealier VCS releases

You may see the following issue.

Zone agent offline/clean entry points delete VCS Zone users from configuration. After upgrade to VCS 6.0, VCS Zone users need to be added to the configuration. VCS Zone users can be added by running `hazonesetup` utility with new syntax after upgrade. See the *Symantec Storage Foundation and High Availability Solutions Virtualization Guide* for Solaris for more information on `hazonesetup` utility and see the *Symantec Storage Foundation and High Availability Solutions Virtualization Guide* for Solaris.

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.

VCS installation with CPI fails when a non-global zone is in installed state and zone root is not mounted on the node (2731178)

On Solaris 10, CPI tries to boot a zone in installed state during installation/ or uninstallation. The boot fails if the underlying storage for zone root is not imported and mounted onto the node, causing the installation or uninstallation to fail.

Workaround: Make sure that the non-global zones are in running or configured state when CPI is invoked for installation or uninstallation.

Log messages are displayed when VRTSvcs is uninstalled on Solaris 11 [2919986]

The following message is displayed when you uninstall VRTSvcs package on Solaris 11 OS.

```
The following unexpected or editable files and directories were
salvaged while executing the requested package operation; they
have been moved to the displayed location in the image:

  var/VRTSvcs/log -> /var/pkg/lost+found/var/VRTSvcs/log-20111216T122049Z
  var/VRTSvcs/lock -> /var/pkg/lost+found/var/VRTSvcs/lock-20111216T122049Z
  var/VRTSvcs -> /var/pkg/lost+found/var/VRTSvcs-20111216T122049Z
  etc/VRTSvcs/conf/config
->/var/pkg/lost+found/etc/VRTSvcs/conf/config-20111216T122049Z
```

You can safely ignore this message as this is an expected behavior of IPS packaging. The files mentioned in the above message are not part of the package. As a result, uninstallation moves them to `/var/pkg/lost+found` directory.

Cluster goes into `STALE_ADMIN_WAIT` state during upgrade from VCS 5.1 to 6.1 [2850921]

While performing a manual upgrade from VCS 5.1 to VCS 6.1, cluster goes in `STALE_ADMIN_WAIT` state if there is an entry of DB2udbTypes.cf in main.cf.

Installation of VRTSvcsea package in VCS 5.1 creates a symbolic link for `Db2udbTypes.cf` file inside `/etc/VRTSvcs/conf/config` directory which points to `/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf`. During manual upgrade, the VRTSvcsea package for VCS 5.1 gets removed, which in turn removes the symbolic link for file `Db2udbTypes.cf` inside `/etc/VRTSvcs/conf/config` directory. After the complete installation of VRTSvcsea for VCS 6.1, because of absence of

file `Db2udbTypes.cf` inside `/etc/VRTSvcs/conf/config`, cluster goes into STALE ADMIN WAIT state.

Workaround: Manually copy `DB2udbTypes.cf` from `/etc/VRTSagents/ha/conf/Db2udb` directory to the `/etc/VRTSvcs/conf/config` directory after the manual upgrade before starting HAD.

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.

After Live Upgrade to Solaris 10 Update 10/Update 11, boot from an alternate boot environment fails [2370250]

If your setup involves volumes in a shared disk group that are mounted as CFS in a cluster, then during Live Upgrade using the `vxlustart` command from any supported Solaris version to Solaris 10 Update 10/11, boot from an alternate boot environment may fail.

Workaround:

1 Run the `vxlufinish` command. Enter:

```
# vxlufinish
```

2 Manually delete the entries of all the volumes of shared disks that are mounted as CFS in the `/altroot.5.10/etc/vfstab` directory. Enter:

```
rm -rf /altroot.5.10/etc/vfstab
```

3 Restart the system.

On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail (2424410)

On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail with the following error:

```
Generating file list.
Copying data from PBE <source.24429> to ABE <dest.24429>.
```

```
99% of filenames transferredERROR: Data duplication process terminated
unexpectedly.
ERROR: The output is </tmp/lucreate.13165.29314/lucopy.errors.29314>.

29794 Killed
Fixing zonepaths in ABE.
Unmounting ABE <dest.24429>.
100% of filenames transferredReverting state of zones in PBE
<source.24429>.
ERROR: Unable to copy file systems from boot environment <source.24429>
to BE <dest.24429>.
ERROR: Unable to populate file systems on boot environment <dest.24429>.
Removing incomplete BE <dest.24429>.
ERROR: Cannot make file systems for boot environment <dest.24429>.
```

This is a known issue with the Solaris `lucreate` command.

Workaround: Check with Oracle for possible workarounds for this issue.

On Solaris 11, if a reboot is performed during upgrade from 6.0PR1 to 6.2, the pkg verify VRTSsfmh command results in an error (3624856)

On Solaris 11, if a reboot is performed during upgrade from 6.0PR1 to 6.2, the `pkg verify VRTSsfmh` command results in the following error:

```
pkg verify VRTSsfmh
    PACKAGE
STATUS
    pkg://Symantec/VRTSsfmh
ERROR
            dir: var/opt/VRTSsfmh
                    Group: 'root (0)' should be 'other (1)'
            dir: var/opt/VRTSsfmh/etc
                    Missing: directory does not exist
            dir: var/opt/VRTSsfmh/logs
                    Group: 'root (0)' should be 'other (1)'
            dir: var/opt/VRTSsfmh/tmp
                    Group: 'root (0)' should be 'other (1)'
            file: opt/VRTSsfmh/web/operator/cgi-bin/firedrill.pl
                    Missing: regular file does not exist
```

Workaround:

- Set the "Symantec" publisher repository pointing to `VRTSpkgs.p5p`.

```
# pkg set-publisher -P -g /mnt/release_train/sol/6.2/
SxRT-6.2-2014-10-01a/dvd1-sol_sparc/sol11_sparc/pkgs/VRTSpkgs.p5p
Symantec
```

- Run the `pkg fix VRTSsfmh` command.

```
# pkg fix VRTSsfmh
```

New agent types are not detected after SFHA stack upgrade to 6.2 [3654406]

On upgrading SFHA stack to 6.2 on Solaris 11, the new types may not be imported. As a result, newly introduced types like SFCache are not seen.

Perform the following steps to verify that you are not impacted by this issue and also for remediation steps.

1. Verify whether SFCache type is present on all running nodes with the command `hatype -list|grep SFCache` command.
2. If SFCache is present in the `hatpye -list` output, then you are not impacted by this issue.
3. If SFCache is not present in the `hatype -list` output, continue with the following steps.
4. Stop VCS on all nodes with `hastop -all -force`.
5. Identify the node which has types.cf or types.cf.previous containing SFCache (`grep -l SFCache /etc/VRTSvcs/conf/config/types.cf /etc/VRTSvcs/conf/config/types.cf.previous`).
6. If SFCache is present in types.cf, start VCS on this node first (`hastart`). Once VCS goes into RUNNING state on this node, you can proceed starting VCS on other nodes.
7. If SFCache is not present in types.cf on any node but in types.cf.previous, then copy types.cf.previous to types.cf on this node. Copy the `.previous` files for other type definition files like OracleTypes.cf, SybaseTypes.cf, Db2udbTypes.cf and so on, which are included in the main.cf. Start VCS on this node. Once VCS goes into RUNNING state on this node, check whether “Frozen =1” attribute got configured for available service groups. If present , unfreeze all service groups using `hagrp -unfeeze <group_name> -persistent` command. You can now proceed starting VCS on other nodes.

vxlustart failed due to lumount error when performing Live Upgrade to Solaris 10 Update 11 (3035982)

Live Upgrade (LU) to Solaris 10 Update 11 using `vxlustart` fails with following error:

```
# lumount -n dest.7667 /altroot.5.10
ERROR: mount point directory </altroot.5.10> is not empty
ERROR: failed to create mount point </altroot.5.10> for file system
</dev/dsk/c1t1d0s0>
ERROR: cannot mount boot environment by name <dest.7667>
ERROR: vxlustart: Failed: lumount -n dest.7667 /altroot.5.10
```

Workaround: To perform Live Upgrade to Solaris 10 Update 11, use one of the following procedures for your operating system version.

To perform Live Upgrade from Solaris 10 Update 10 to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 10 Update 9 or below to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 9 to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Install the patch 121430-72. (Do NOT patch to a higher version of 121430, such as 121430-92.)
3. Use `vxlustart` to upgrade to Solaris 10 Update 11.

vxlustart failed due to lumount error when performing Live Upgrade to Solaris 10 Update 11 (3035982)

Live Upgrade (LU) to Solaris 10 Update 11 using `vxlustart` fails with following error:

```
# lumount -n dest.7667 /altroot.5.10
ERROR: mount point directory </altroot.5.10> is not empty
```

```
ERROR: failed to create mount point </altroot.5.10> for file system
</dev/dsk/c1t1d0s0>
ERROR: cannot mount boot environment by name <dest.7667>
ERROR: vxlustart: Failed: lumount -n dest.7667 /altroot.5.10
```

Workaround: To perform Live Upgrade to Solaris 10 Update 11, use one of the following procedures for your operating system version.

To perform Live Upgrade from Solaris 10 Update 10 to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 10 Update 9 or below to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 9 to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Install the patch 121430-72. (Do NOT patch to a higher version of 121430, such as 121430-92.)
3. Use `vxlustart` to upgrade to Solaris 10 Update 11.

Operational issues for VCS

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

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

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

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the *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.

Stale legacy_run services seen when VCS is upgraded to support SMF [2431741]

If you have VCS 5.0MPx installed on a Solaris 10 system, VCS uses RC scripts to manage starting services. If you upgrade VCS to any version that supports SMF for VCS, you see stale legacy_run services for these RC scripts in addition to the SMF services.

Workaround: There are two ways to remove these legacy services:

- Open svccfg console using `svccfg -s smf/legacy_run` and delete the legacy services.
  For example:

```
svccfg -s smf/legacy_run
svc:/smf/legacy_run> listpg *
rc2_d_S70llt    framework     NONPERSISTENT
rc2_d_S92gab    framework     NONPERSISTENT
svc:/smf/legacy_run> delpg rc2_d_S70llt
svc:/smf/legacy_run> delpg rc2_d_S92gab
svc:/smf/legacy_run> exit
```

- Reboot the system.

The hastop -all command on VCS cluster node with AlternateIO resource and StorageSG having service groups may leave the node in LEAVING state

On a VCS cluster node with AlternateIO resource configured and StorageSG attribute contain service groups with Zpool, VxVM or CVMVolDG resources, `hastop -local` or `hastop -all` commands may leave the node in "LEAVING" state.

This issue is caused by lack of dependency between service group containing LDom resource and service groups containing storage resources exported to logical domain in alternate I/O domain scenarios. In this scenario VCS may attempt to stop the storage service groups before stopping logical domain which is using the resources.

Workaround: Stop the LDom service group before issuing `hastop -local` or `hastop -all` commands.

Missing characters in system messages [2334245]

You may see missing characters, especially in long system messages in response to certain commands.

Workaround: No workaround.

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.

After OS upgrade from Solaris 10 update 8 or 9 to Solaris 10 update 10 or 11, Samba server, SambaShare and NetBios agents fail to come online [3321120]

On Solaris 10 update 8 and update 9, default path of Samba binaries is `/usr/sfw/sbin/smbd` and default samba configuration file location is `/etc/sfw/smb.conf`. On Solaris 10 update 10 and update 11, the default path of Samba binaries is changed to `/usr/sbin/smbd` and default Samba configuration file location is `/etc/samba/smb.conf`. Therefore, after OS upgrade from Solaris 10 update 8 or update 9 to Solaris 10 update 10 or update 11, Samba server, SambaShare and NetBios agents are unable to locate binaries and configuration file.

Workaround: After the OS upgrade from Solaris 10 update 8 or update 9 to Solaris 10 update 10 or update 11, update the SambaTopDir and ConfFile attributes of the Samba server resources appropriately to reflect the correct location.

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.

System encounters multiple VCS resource timeouts and agent core dumps [3424429]

The system encounters multiple VCS resource timeouts and agent core dumps without any specific reason.

The issue pertains to a hardware errata with the Intel Xeon CPUs where a processor can go into a low power sleep mode, but takes a long time to wake up. This can cause erratic scheduling behavior, leading to unexpected delays, expired timers, or occasional freezes. For more information, see the Oracle document: https://support.oracle.com/epmos/faces/BugDisplay?id=15659645

Workaround: Add the following lines to the `/etc/system` file and reboot the system:

```
set idle_cpu_prefer_mwait = 0
set idle_cpu_no_deep_c = 1
```

Some VCS components do not work on the systems where a firewall is configured to block TCP traffic [3545338]

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.

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.

Missing host names in engine_A.log file (1919953)

The GUI does not read the `engine_A.log` file. It reads the `engine_A.ldf` file, gets the message id from it, and then queries for the message from the bmc file of the appropriate locale (Japanese or English). The bmc file does not have system names present and so they are read as missing.

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.

Character corruption observed when executing the uuidconfig.pl -clus -display -use_llthost command [2350517]

If password-less ssh/rsh is not set, the use of `uuidconfig.pl` command in non-English locale may print garbled characters instead of a non-English string representing the Password prompt.

Workaround: No workaround.

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.

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.

Oracle service group faults on secondary site during failover in a disaster recovery scenario [2653704]

Oracle service group fails to go online in the DR site when disaster strikes the primary site. This happens if the AutoFailover attribute on the Service Group is set to 1 and when the corresponding service group's FireDrill is online in the DR site. Firedrill Service group may remain ONLINE on the DR site.

Workaround: If the service group containing the Oracle (or any database) resource faults after attempting automatic DR failover while FireDrill is online in the DR site, manually offline the FireDrill Service Group. Subsequently, attempt the online of the Oracle Service Group in the DR site.

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

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

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

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

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

Workaround: No workaround.

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

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

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

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

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

GCO clusters remain in INIT state [2848006]

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

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

Workaround: Make sure that above two conditions are rectified. Refer to *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 for one of the following reasons:

- If you first use a non-root user without a home directory and then create a home directory for the same user.
- If you configure security on a cluster and then un-configure and reconfigure it.

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).

Startup trust failure messages in system logs [2721512]

If you configure a cluster with security enabled, there might be some messages logged in system message logs related to Symantec authentication. These messages can be ignored and have no effect on functionality.

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 [3380586]

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.

VCS crashes if feature tracking file is corrupt [3603291]

VCS keeps a track of some specific features used in the VCS cluster. For example, if a Global service group is brought online then the feature is logged in a specific feature tracking file. If the file however is corrupt, then VCS may dump core when attempting to write data to the file.

Workaround: Delete the corrupt feature tracking file (`/var/vx/vftrk/vcs`) and restart VCS.

RemoteGroup agent on versions lower than 6.2 reports service group status as UNKNOWN [3638347]

When the RemoteGroup agent running on a VCS version lower than 6.2 tries to monitor a service group on a 6.2 cluster, it reports the service group status as UNKNOWN.

Workaround: No workaround.

RemoteGroup agent and non-root users may fail to authenticate after a secure upgrade [3649457]

On upgrading a secure cluster to 6.2, the following issues may occur with unable to open a secure connection error:

- The RemoteGroup agent may fail to authenticate with remote cluster.

- Non-root users may fail to authenticate.

Workaround

1. Set `LC_ALL=C` on all nodes before upgrade or perform the following steps after the upgrade on all nodes of the cluster:
   - Stop HAD.
   - Set `LC_ALL=C`.
   - Start HAD using `hastart`.
2. Reset LC_ALL attribute to the previous value once the non-root users are validated.

Issues related to the bundled agents

Entry points that run inside a zone are not cancelled cleanly [1179694]

Cancelling entry points results in the cancellation of only the `zlogin` process. The script entry points that run inside a zone are forked off using the `zlogin` command. However, the `zlogin` command forks off an `sh` command, which runs in the context of the Solaris zone. This shell process and its family do not inherit the group id of the `zlogin` process, and instead get a new group id. Thus, it is difficult for the agent framework to trace the children or grand-children of the shell process, which translates to the cancellation of only the `zlogin` process.

Workaround: Oracle must provide an API or a mechanism to kill all the children of the `zlogin` process that was started to run the entry point script in the local-zone.

Solaris mount agent fails to mount Linux NFS exported directory

The Solaris mount agent mounts the mount directories. At this point, if it tries to mount a Linux NFS exported directory, the mount fails showing the following error:

```
nfs mount: mount: <MountPoint>: Not owner
```

This is due to system NFS default version mismatch between Solaris and Linux.

The workaround for this is to configure `MountOpt` attribute in mount resource and set `vers=3` for it.

Example

```
root@north $ mount -F nfs south:/test /logo/
nfs mount: mount: /logo: Not owner
root@north $
Mount nfsmount (
                MountPoint = "/logo"
                BlockDevice = "south:/test"
                FSType = nfs
                MountOpt = "vers=3"
               )
```

The zpool command runs into a loop if all storage paths from a node are disabled

The Solaris Zpool agent runs `zpool` commands to import and export zpools. If all paths to the storage are disabled, the zpool command does not respond. Instead, the zpool export command goes into a loop and attempts to export the zpool. This continues till the storage paths are restored and zpool is cleared. As a result, the offline and clean procedures of Zpool Agent fail and the service group cannot fail over to the other node.

Workaround: You must restore the storage paths and run the zpool clear command for all the pending commands to succeed. This will cause the service group to fail over to another node.

Zone remains stuck in down state if tried to halt with file system mounted from global zone [2326105]

If zone halts without unmounting the file system, the zone goes to down state and does not halt with the `zoneadm` commands.

Workaround: Unmount the file system manually from global zone and then halt the zone. For VxFS, use following commands to unmount the file system from global zone.

To unmount when VxFSMountLock is 1

```
umount -o mntunlock=VCS <zone root path>/<Mount Point>
```

To forcefully unmount when VxFSMountLock is 1:

```
# umount -f -o mntunlock=VCS <zone root path>/<Mount Point>
```

To unmount when VxFSMountLock is 0:

```
# umount <zone root path>/<Mount Point>
```

To forcefully unmount when VxFSMountLock is 0:

```
# umount -f <zone root path>/<Mount Point>
```

To halt the zone, use following command:

```
# zoneadm -z <zone_name>  halt
```

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

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

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

The zpool commands hang and remain in memory till reboot if storage connectivity is lost [2368017]

If the FailMode attribute of `zpool` is set to continue or wait and the underlying storage is not available, the `zpool` commands hang and remain in memory until the next reboot.

This happens when storage connectivity to the disk is lost, the `zpool` commands hang and they cannot be stopped or killed. The zpool commands run by the monitor entry point remains in the memory.

Workaround: There is no recommended workaround for this issue.

Offline of zone resource may fail if `zoneadm` is invoked simultaneously [2353541]

Offline of zone EP uses `zoneadm` command to offline a zone. Therefore, if `zoneadm` is invoked simultaneously for multiple zones, the command may fail. This is due to Oracle bug 6757506 that causes a race condition between multiple instances of `zoneadm` command and displays the following message:

```
zoneadm: failed to get zone name: Invalid argument
```

Workaround: No workaround.

Password changed while using `hazonesetup` script does not apply to all zones [2332349]

If you use the same user name for multiple zones, updating password for one zone does not updated the password of other zones.

Workaround: While updating password for VCS user which is used for multiple zones, update password for all the zones.

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 local Zone on Solaris

Workaround: No workaround.

Share resource goes offline unexpectedly causing service group failover [1939398]

Share resource goes offline unexpectedly and causes a failover when NFSRestart resource goes offline and UseSMF attribute is set to 1 (one).

When NFSRestart resource goes offline, NFS daemons are stopped. When UseSMF attribute is set to 1, the exported file systems become unavailable, hence Share resource unexpectedly goes offline.

Workaround: Set the value of ToleranceLimit of Share resource to a value more than 1.

Mount agent does not support all scenarios of loopback mounts

For a mount point under VCS control, you can create loop back mounts for the mount point. For example, mount point /mntpt is mounted on /a as loop back mount and /a is mounted on /b as loop back mount, then offline and online of the mount resource fails.

Workaround: Mount the mount point /mntpt on /b as loop back mount.

Invalid Netmask value may display code errors [2583313]

If you specify invalid Netmask value for the IP resource attribute, you may see the code errors similar to the following when you try to online the resource.

```
==================================
Illegal hexadecimal digit 'x' ignored at
/opt/VRTSperl/lib/site_perl/5.12.2/Net/Netmask.pm line 78.
ifconfig: <Netmask_value>: bad address
==========================================
```

Workaround: Make sure you specify a valid Netmask value.

Zone root configured on ZFS with ForceAttach attribute enabled causes zone boot failure (2695415)

On Solaris 11 system, attaching zone with `-F` option may result in zone boot failure if zone root is configured on ZFS.

Workaround: Change the ForceAttach attribute of Zone resource from 1 to 0. With this configuration, you are recommended to keep the default value of DetachZonePath as 1.

Error message is seen for Apache resource when zone is in transient state [2703707]

If the Apache resource is probed when the zone is getting started, the following error message is logged:

```
Argument "VCS ERROR V-16-1-10600 Cannot connect to VCS engine\n"
isn't numeric in numeric ge (>=) at /opt/VRTSvcs/bin/Apache/Apache.pm
line 452.
VCS ERROR V-16-1-10600 Cannot connect to VCS engine
LogInt(halog call failed):TAG:E:20314 <Apache::ArgsValid> SecondLevel
MonitorTimeOut must be less than MonitorTimeOut.
```

Workaround: You can ignore this message. When the zone is started completely, the `halog` command does not fail and Apache agent monitor runs successfully.

Monitor falsely reports NIC resource as offline when zone is shutting down (2683680)

If a NIC resource is configured for an Exclusive IP zone, the NIC resource is monitored inside the zone when the zone is functional. If the NIC monitor program is invoked when the zone is shutting down, the monitor program may falsely report the NIC resource as offline. This may happen if some of the networking services are offline but the zone is not completely shut down. Such reports can be avoided if you override and set the ToleranceLimit value to a non-zero value.

Workaround: When a NIC resource is configured for an Exclusive IP zone, you are recommended to set the ToleranceLimit attribute to a non-zero value.

Calculate the ToleranceLimit value as follows:

Time taken by a zone to completely shut down must be less than or equal to NIC resource's MonitorInterval value + (MonitorInterval value x ToleranceLimit value).

For example, if a zone take 90 seconds to shut down and the MonitorInterval for NIC agent is set to 60 seconds (default value), set the ToleranceLimit value to 1.

Apache resource does not come online if the directory containing Apache pid file gests deleted when a node or zone restarts (2680661)

The directory in which Apache http server creates PidFile may get deleted when a node or zone restarts. Typically the PidFile is located at `/var/run/apache2/httpd.pid`. When the zone reboots, the `/var/run/apache2` directory may get removed and hence the http server startup may fail.

Workaround: Make sure that Apache http server writes the PidFile to an accessible location. You can update the PidFile location in the Apache http configuration file (For example: `/etc/apache2/httpd.conf`).

Online of LDom resource may fail due to incompatibility of LDom configuration file with host OVM version (2814991)

If you have a cluster running LDom with different OVM versions on the hosts, then the LDom configuration file generated on one host may display error messages when it is imported on the other host with a different OVM version. Thus, the online of LDom resource may also fail.

For example, if you have a cluster running LDom with OVM versions 2.2 on one and OVM 2.1 on the other node, the using XML configuration generated on the host with OVM 2.2 may display errors when the configuration is imported on the host with OVM 2.1. Thus, the online of LDom resource fails.

The following error message is displayed:

```
ldm add-domain failed with error Failed to add device
/ldom1/ldom1 as ld1_disk1@primary-vds0 because this device
is already  exported on LDom primary. Volume ld1_disk1
already exists in vds primary-vds0.
```

Workaround: If the CfgFile attribute is specified, ensure that the XML configuration generated is compatible with the OVM version installed on the nodes.

Online of IP or IPMultiNICB resource may fail if its IP address specified does not fit within the values specified in the allowed-address property (2729505)

While configuring an IP or IPMultiNICB resource to be run in a zone, if the IP address specified for the resource does not match the values specified in the **allowed-address** property of the zone configuration, then the online of IP resource may fail. This behavior is seen only on Solaris 11 platform.

Workaround: Ensure that the IP address is added to **allowed-address** property of the zone configuration.

Application resource running in a container with PidFiles attribute reports offline on upgrade to VCS 6.0 or later [2850927]

Application resource configured to run in a container configured with PidFiles attribute reports state as offline after upgrade to VCS 6.0 or later versions.

When you upgrade VCS from lower versions to 6.0 or later, if application resources are configured to run in a container with monitoring method set to PidFiles, then upgrade may cause the state of the resources to be reported as offline. This is due to changes introduced in the Application agent where if the resource is configured to run in a container and has PidFiles configured for monitoring the resource then the value expected for this attribute is the pathname of the PID file relative to the zone root.

In releases prior to VCS 6.0, the value expected for the attribute was the pathname of the PID file including the zone root.

For example, a configuration extract of an application resource configured in VCS 5.0MP3 to run in a container would appear as follows:

```
Application apptest (
User = root
StartProgram = "/ApplicationTest/app_test_start"
StopProgram = "/ApplicationTest/app_test_stop"
PidFiles = {
      "/zones/testzone/root/var/tmp/apptest.pid" }
ContainerName = testzone
)
```

Whereas, the same resource if configured in VCS 6.0 and later releases would be configured as follows:

```
Application apptest (
User = root
StartProgram = "/ApplicationTest/app_test_start"
StopProgram = "/ApplicationTest/app_test_stop"
PidFiles = {
      "/var/tmp/apptest.pid" }
)
```

Note: The container information is set at the service group level.

Workaround: Modify the PidFiles pathname to be relative to the zone root as shown in the latter part of the example.

```
# hares -modify apptest PidFiles /var/tmp/apptest.pid
```

NIC resource may fault during group offline or failover on Solaris 11 [2754172]

When NIC resource is configured with exclusive IP zone, NIC resource may fault during group offline or failover. This issue is observed as zone takes long time in shutdown on Solaris 11. If NIC monitor is invoked during this window, NIC agent may treat this as fault.

Workaround: Increase ToleranceLimit for NIC resource when it is configured for exclusive IP zone.

NFS client reports error when server is brought down using `shutdown` command [2872741]

On Solaris 11, when the VCS cluster node having the NFS share service group is brought down using `shutdown` command, NFS clients may report "Stale NFS file handle" error. During shutdown, the SMF service svc:/network/shares un-shares all the shared paths before taking down the virtual IP. Thus, the NFS clients accessing this path get stale file handle error.

Workaround: Before you shutdown the VCS cluster node, disable the svc:/network/shares SMF service, so that only VCS controls the un-sharing of the shared paths during the shutdown operation.

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>
```

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.

IP Agent fails to detect the online state for the resource in an exclusive-IP zone [3592683]

IP Agent does not detect the online state for the resource inside an exclusive-IP zone monitoring an IPv6 address if the link-local address is down.

Workaround: Bring the link-local address of the device up for the IP agent to detect the IPv6 address state properly.

SFCache Agent fails to enable caching if cache area is offline [3644424]

SFCache agent cannot enable caching if cache area associate with this particular object is in offline state. User need to manually online the cache area to make sure that caching can be enabled/disabled.

Workaround: Online the cache area using `sfcache` command

```
# sfcache online <cache_area_name>
```

RemoteGroup agent may stop working on upgrading the remote cluster in secure mode [3648886]

RemoteGroup agent may report the resource state as UNKNOWN if the remote cluster is upgraded to VCS 6.2 in secure mode.

Workaround: Restart the RemoteGroup agent.

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.

ASMInstance resource monitoring offline resource configured with OHASD as application resource logs error messages in VCS logs [2846945]

When the Oracle High Availability Services Daemon (OHASD) is configured as an application resource to be monitored under VCS and if this resource is offline on the failover node then the ASMInstance resource in the offline monitor logs the following error messages in the VCS logs:

```
ASMInst:asminst:monitor:Cluster Synchronization Service
process is not running.
```

Workaround: Configure the application in a separate parallel service group and ensure that the resource is online.

Oracle agent fails to offline pluggable database (PDB) resource with PDB in backup mode [3592142]

If the PDB is in backup mode and if you attempt to offline the corresponding PDB resource, this will cause PDB resource to go into "Unable to Offline" state.

Workaround: Manually remove the PDB from the backup mode before attempting to take the PDB resource offline.

Clean succeeds for PDB even as PDB staus is UNABLE to OFFLINE [3609351]

Oracle does not allow any operation on a PDB when the PDB is in backup mode. This is an expected behavior of Oracle. Therefore, a shutdown fails when it is initiated on a PDB in backup mode and returns an UNABLE TO OFFLINE status for the PDB. If PDB is removed from the backup mode using the SQL script, the agent framework is unable to change the UNABLE TO OFFLINE status of the PDB

as clean is called. Since Oracle does not differntiate between clean and offline for PDB, clean succeeds for the PDB in spite of being in UNABLE TO OFFLINE state.

Workaround: No workaround.

Second level monitoring fails if user and table names are identical [3594962]

If the table inside CDB has same name as the user name, second level monitoring fails and Oracle agent fails to update the table. For example, if user name is `c##pdbuser1` and table is created as `c##pdbuser1.vcs`, then Oracle agent is unable to update it.

Workaround: Avoid having identical user and CDB table names.

Monitor entry point times out for Oracle PDB resources when CDB is moved to suspended state in Oracle 12.1.0.2 [3643582]

In Oracle-12.1.0.2.0, when CDB is in SUSPENDED mode, then the SQL command for PDB view (`v$pdbs`) hangs. Due to this, the monitor entry point in PDB gets timed out and there is no issue found in oracle-12.1.0.1.0 .

Workaround: No workaround.

Oracle agent fails to online and monitor Oracle instance if threaded_execution parameter is set to true [3644425]

In Oracle 12c, the threaded execution feature is enabled. The multithreaded Oracle Database model enables Oracle processes to execute as operating system threads in separate address spaces. If Oracle Database 12c is installed, the database runs in the process mode. If you set a parameter to run the database in threaded mode, some background processes on UNIX and Linux run with each process containing one thread, whereas the remaining Oracle processes run as threads within the processes.

When you enable this parameter, Oracle agent is unable to check smon (mandatory process check) and lgwr (optional process check) processes which where tradtionally used for monitoring and which now run as threads.

Workaround: Disable the threaded execution feature as it is no supported on Oracle 12C.

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.

Logs from the script executed other than the agent entry point goes into the engine logs [3547329]

The agent logs of C-based and script-based entry points get logged in the agent log when the attribute value of LogViaHalog is set to 1 (one). To restore to the older logging behavior in which C-based entry point logs were logged in agent logs and script-based entry point logs were logged in engine logs, you can set the LogViaHalog value as 0 (zero). However, it is observed that some C-based entry point logs continue to appear in the engine logs even when LogViaHalog is set to 1 (one). This issue is observed on all the database agents.

Workaround: No workaround.

Issues related to Live Upgrade

Live Upgrade to Solaris 10 Update 10 fails in the presence of zones (2521348)

SFCFSHA Live Upgrade from Solaris 10 Update 7 5.1SP1 to Solaris 10 Update 10 using the `vxlustart` commands fails in the presence of zones with the following error message:

```
ERROR: Installation of the packages from this media of the media failed;
pfinstall returned these diagnostics:
Processing default locales
        - Specifying default locale (en_US.ISO8859-1)
Processing profile
ERROR: This slice can't be upgraded because of missing usr packages for
the following zones:
ERROR:    zone1
ERROR:    zone1
ERROR: This slice cannot be upgraded because of missing usr packages for
one or more zones.
The Solaris upgrade of the boot environment <dest.27152> failed.
```

This is a known issue with the Solaris `luupgrade` command.

Workaround: Check with Oracle for possible workarounds for this issue.

Issues related to VCS in Japanese locales

This section covers the issues that apply to VCS 6.2 in a Japanese locale.

The hares -action command displays output in English [1786742]

The `hares -action` command incorrectly displays output in English.

Character corruption issue

Character corruption occurs if installer is run with HIASCII option on French locale. [1539754, 1539747]

Workaround: No workaround.

Messages inside the zone are not localized

Locale is not set correctly for Solaris zone. Therefore, you may not see localized messages inside the zone.

Workaround: No workaround.

System messages having localized characters viewed using `hamsg` may not be displayed correctly

If you use `hamsg` to view system messages, the messages containing a mix of English and localized characters may not be displayed correctly. [2405416]

Workaround: No workaround. However, you can view English messages in the VCS log file.

Standalone utilities display output in English [2848012]

The following utilities display output in English:

- `-haping`
- `-hamultinicb`
- `-haipswitch`

Workaround: No workaround.

English error messages displayed by the gcoconfig wizard [3018221]

Whenever the gcoconfig wizard calls a command internally, the messages from that command are displayed in English.

Workaround: No workaround.

Issues related to global clusters

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

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

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

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

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

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

LLT known issues

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

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

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

Cannot configure LLT if full device path is not used in the lltab file (2858159)

(Oracle Solaris 11) On virtual machines ensure that you use the full path of the devices corresponding to the links in lltab. For example, use /dev/net/net1 instead of /dev/net/net:1 in the lltab file, otherwise you cannot configure LLT.

Fast link failure detection is not supported on Solaris 11 (2954267)

Fast link failure detection is not supported on Solaris 11 operating system because the operating system cannot provide notification calls to LLT when a link failure occurs. If the operating system kernel notifies LLT about the link failure, LLT can detect a link failure much earlier than the regular link failure detection cycle. As Solaris 11 does not notify LLT about link failures, failure detection cannot happen before the regular detection cycle.

Workaround: None

GAB known issues

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

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

After you unregister the 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.

GAB may fail to stop during a phased upgrade on Oracle Solaris 11 (2858157)

While performing a phased upgrade on Oracle Solaris 11 systems, GAB may fail to stop. However, CPI gives a warning and continues with stopping the stack.

Workaround: Reboot the node after the installer completes the upgrade.

Cannot run pfiles or truss files on gablogd (2292294)

When pfiles or truss is run on gablogd, a signal is issued to gablogd. gablogd is blocked since it has called an gab ioctl and is waiting for events. As a result, the pfiles command hangs.

Workaround: None.

(Oracle Solaris 11) On virtual machines, sometimes the common product installer (CPI) may report that GAB failed to start and may exit (2879262)

GAB startup script may take longer than expected to start up. The delay in start up can cause the CPI to report that GAB failed and exits.

Workaround: Manually start GAB and all dependent services.

The Installer fails to unload GAB module while installation of SF packages [3560458]

The Installer succeeds to upgrade SF package from 6.0.1 to 6.0.5 or from 6.1 to 6.1.1, but GAB module (for 6.0.1 or 6.1) fails to unload and remains in loaded state. The issue is seen with the recent updates of Solaris OS 11U1 (SRU 8). During un-installation of VCS packages, unloading of GAB fails.

`Workaround:` Restart the system. Restarting the system will unload the module successfully.

I/O fencing known issues

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

Delay in rebooting Solaris 10 nodes due to vxfen service timeout issues (1897449)

When you reboot the nodes using the `shutdown -i6 -g0 -y` command, the following error messages may appear:

```
svc:/system/vxfen:default:Method or service exit
timed out. Killing contract 142
svc:/system/vxfen:default:Method "/lib/svc/method/vxfen stop"
failed due to signal Kill.
```

This error occurs because the vxfen client is still active when VCS attempts to stop I/O fencing. As a result, the vxfen stop service times out and delays the system reboot.

Workaround: Perform the following steps to avoid this vxfen stop service timeout error.

To avoid the vxfen stop service timeout error

1 Stop VCS. On any node in the cluster, run the following command:

```
# hastop -all
```

2 Reboot the systems:

```
# shutdown -i6 -g0 -y
```

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.

When I/O fencing is not up, the svcs command shows VxFEN as online (2492874)

Solaris 10 SMF marks the service status based on the exit code of the start method for that service. The VxFEN start method executes the vxfen-startup script in the background and exits with code 0. Hence, if the vxfen-startup script subsequently exits with failure then this change is not propagated to SMF. This behavior causes the `svcs` command to show incorrect status for VxFEN.

Workaround: Use the `vxfenadm` command to verify that I/O fencing is running.

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 package 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 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.

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.

The `vxfenconfig -l` command output does not list Coordinator disks that are removed using the `vxdmpadm exclude dmpnodename=<dmp_disk/node>` command [3644431]

After you remove a Coordinator disk used by fencing or fencing disk group by running the `vxdmpadm exclude dmpnodename=<dmp_disk/node>` command, the removed disk is not listed in the `vxfenconfig -l` command output.

In case of a split brain, the `vxfen` program cannot use the removed disk as a coordination point in the subsequent fencing race.

Workaround: Run the `vxdmpadm include dmpnodename=<dmp_disk/node>` command to again enable the dmp disk. This disk will show up in subsequent `vxfencondig -l` output.

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 fault zones if zones are in ready or down state [2290883]

IMF does not fault zones if zones are in ready or down state.

IMF does not detect if zones are in ready or down state. In Ready state, there are no services running inside the running zones.

Workaround: Offline the zones and then restart.

IMF does not detect the zone state when the zone goes into a maintenance state [2535733]

IMF does not detect the change in state. However, the change in state is detected by Zone monitor in the next cycle.

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.

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.

ProPCV fails to prevent a script from running if it is run with relative path [3617014]

If the absolute path is registered with AMF for prevention and the script is run with the relative path, AMF fails to prevent the script from running.

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.

Unable to log on to secure VCS clusters on Solaris 11 using Java GUI (2718943)

Connecting to secure clusters deployed on Solaris 11 systems using VCS Java GUI is not supported in VCS 6.0PR1. The system displays the following error when you attempt to use the Java GUI:

```
Incorrect username/password
```

Workaround: No workaround.

Issues related to live migration

Following are the issues related to live migration.

Operating system in guest domain with multiple IO services hangs when guest migrates back [3127470]

Operating system inside the guest domain hangs when the guest domain is provided with IO services from multiple IO domains but not from primary domain and guest domain is migrated to another node and back to the source node.

Workaround: Make sure that firmware of the physical system is upgraded to latest version.

Issues related to virtualization

Locale message displayed on Solaris 11 system for solaris10 brand zones

When you run the `zlogin` command on a Solaris 11 system, the system logs the following error message:

```
Could not set locale correctly.
```

The default locale for Solaris 11 is en_US.UTF-8 and that of Solaris 10 is C. With solaris10 brand zone, en_US.UTF-8 is not installed inside the zone by default. Therefore, the error message is logged.

Workaround: This message can be safely ignored as there is no functionality issue. To avoid this message, install en_US.UTF-8 locale on solaris10 brand zone.

VCS Cluster Configuration wizard issues

IPv6 verification fails while configuring generic application using VCS Cluster Configuration wizard [3614680]

The VCS Cluster Configuration wizard fails to check whether IPv6 IP is already plumbed while configuring a generic application through the Virtual IP page. The wizard does neither displays a warning if IPv6 IP is already plumbed elsewhere nor indicates whether it is reachable through a ping.

Workaround: Manually ensure that IPv6 is not plumbed elsewhere on the network before configuring the generic application through the wizard.

Browser shows 404 error and wizard fails to launch when VCS is installed with Jumpstart or upgraded with Live upgrade [3626253]

On Solaris 10 systems, when ApplicationHA or VCS is installed through Jumpstart or Live upgrade mechanism, the wizards cannot be launched. The browser displays the 404 – page not found error because VCS namespace values are not set in the `xprtld` configuration.

Workaround:

1. Boot the system to the newly created boot environment.
2. Ensure xprtld service is in online state

   ```
   # svcs /system/xprtld
   ```

3. Run the following commands:

   For VCS:

   ```
   # /opt/VRTSvcs/portal/admin/conf/configGen.pl
   ```

   For ApplicationHA

   ```
   # /opt/VRTSvcs/portal/admin/plugins/unix/conf/configGen.pl
   ```

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 93.

Limitations related to bundled agents

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

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

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

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

/etc/nsswitch.conf

Volume agent clean may forcibly stop volume resources

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

False concurrency violation when using PidFiles to monitor application resources

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

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

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

Online for LDom resource fails [2517350]

Online of LDom resource fails when the boot disk configured in the guest domain that is a part of the virtual disk multi-pathing group (mpgroup) and also the primary path to the virtual disk is not available.

This is due to the limitations in Oracle VM Server that do not allow retrying of other device paths that exist for the virtual disks, which are part of a virtual disk multi-pathing group, when booting a guest domain.

Workaround: None.

Zone agent registered to IMF for Directory Online event

The Directory Online event monitors the Zone root directory. If the parent directory of the Zone root directory is deleted or moved to another location, AMF does not

provide notification to the Zone agent. In the next cycle of the zone monitor, it detects the change and reports the state of the resource as offline.

LDom resource calls clean entry point when primary domain is gracefully shut down

LDom agent sets failure policy of the guest domain to stop when primary domain stops. Thus when primary domain is shut down, guest domain is stopped. Moreover, when primary domain is shutdown, ldmd daemon is stopped abruptly and LDom configuration cannot be read. These operations are not under VCS control and VCS may call clean entry point.

Workaround: No workaround.

Application agent limitations

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

Interface object name must match net<*x*>/v4static for VCS network reconfiguration script in Solaris 11 guest domain [2840193]

If the Solaris 11 guest domain is configured for DR and its interface object name does not match the `net<x>/v4static` pattern then the VCS guest network reconfiguration script (VRTSvcsnr) running inside the guest domain adds a new interface object and the existing entry will remain as is.

Share agent limitation (2717636)

If the Share resource is configured with VCS to share a system directory (for example, /usr) or Oracle Solaris 11 which gets mounted at boot time, the VCS share resource detects it online once VCS starts on the node after a panic or halt. This can lead to a concurrency violation if the share resource is a part of a failover service group, and the group has failed over to another node in the cluster. VCS brings down the Share resource subsequently. This is due to the share command behavior or Oracle Solaris 11, where a directory shared with share command remains persistently on the system across reboots.

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.

On Solaris 10, the online operation of IP resource may fail if `ifconfig -a` returns an error [3609861]

The IP agent uses the output of `ifconfig -a` to determine the next alias of free NIC to plumb IP. In rare and specific scenarios, the `ifconfig -a` command may return an error if it does not find an interface at the time of listing the interface. The IP resource online operation is affected by this and the resource may fault.

Workaround: Increase OnlineRetryLimit to a value higher than the default value.

Mount agent reports resource state as OFFLINE if the configured mount point does not exist [3435266]

If a configured mount point does not exist on a node, then the Mount agent reports the resource state as OFFLINE instead of UNKNOWN on that particular node. If an attempt is made for onlining the resource, it fails on that node as the mount point does not exist.

Workaround: Make sure that configured mount point exists on all nodes of the cluster or alternatively set the CreateMntPt attribute value of Mount agent to 1. This will ensure that if a mount point does not exist then it will create while onlining the resource.

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.

Symantec cluster configuration wizard limitations

Environment variable used to change log directory cannot redefine the log path of the wizard [3609791]

By default, the Symantec cluster configuration wizard writes the logs in `/var/VRTSvcs/log` directory. VCS provides a way to change the log directory through environment variable VCS_LOG, but this does not apply to the logs of VCS wizards.

Workaround: No workaround.

Cluster configuration wizard takes long time to configure a cluster on Solaris systems [3582495]

Some times the VCS cluster configuration wizard takes a long time (10 to 15 minutes) to configure a VCS cluster on Solaris systems. The wizard may appear stuck but it completes the configuration in some time.

Workaround: No workaround.

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.

Sybase agent does not perform qrmutil based checks if Quorum_dev is not set (2724848)

If you do not set the Quorum_dev attribute for Sybase Cluster Edition, the Sybase agent does not perform the qrmutil-based checks. This error in configuration may lead to undesirable results. For example, if qrmutil returns failure pending, the agent does not panic the system. Thus, the Sybase agent does not perform qrmutil-based checks because the Quorum_dev attribute is not set.

Therefore, setting Quorum_Dev attribute is mandatory for Sybase cluster edition.

Pluggable database (PDB) online may timeout when started after container database (CDB) [3549506]

PDB may take long time to start when it is started for the first time after starting CDB. As a result, the PDB online initiated using VCS may cause ONLINE timeout and the PDB online process may get cancelled.

Workaround: Increase the OnlineTimeout attribute value of the Oracle type resource.

Engine hangs when you perform a global cluster upgrade from 5.0MP3 in mixed-stack environments [1820327]

If you try to upgrade a mixed stack VCS environment (where IPv4 and IPv6 are in use), from 5.0MP3 to 5.1SP1, HAD may hang.

Workaround: When you perform an upgrade from 5.0MP3, make sure no IPv6 addresses are plumbed on the system..

Systems in a cluster must have same system locale setting

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

Limitations with DiskGroupSnap agent [1919329]

The DiskGroupSnap agent has the following limitations:

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

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

Cluster Manager (Java console) limitations

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

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

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

See the *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.

Limitations related to LLT

This section covers LLT-related software limitations.

Limitation of LLT support over UDP using alias IP [3622175]

When configuring the VCS cluster, if alias IP addresses are configured on the LLT links as the IP addresses for LLT over UDP, LLT may not work properly.

Workaround: Do not use alias IP addresses for LLT over UDP.

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

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

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *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 package, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm package is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

Node may panic if HAD process is stopped by force and then node is shut down or restarted [3640007]

A node may panic if the HAD process running on it is stopped by force and then it is shut down or restarted.This limitation is observed when you perform the following steps on a cluster node:

1 Stop the HAD process with the `force` flag.

```
# hastop -local -force
```

or

```
# hastop -all -force
```

2 Restart or shut down the node.

The node panics because forcefully stopping VCS on the node leaves all the applications, file systems, CVM, and other process online on that node. If the same node is restarted in this starte, VCS triggers a fencing race to avoid data curruption. However, the restarted node loses the fencing race and panics.

Workaround: No workaround.

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.
- Configuring Veritas Volume Replicator for Zone Disaster Recovery is not supported for zone root replication. Oracle Solaris 11 supports zone root only on ZFS file system.
- Configuring a cluster of mixed nodes such as a cluster between systems running on Solaris 10 and Solaris 11 versions is not supported in VCS 6.2. The configuration is not supported through manual as well as CPI configuration.

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.

Note: The GNOME PDF Viewer is unable to view Symantec documentation. You must use Adobe Acrobat to view the documentation.

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_62_sol.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_62_sol.pdf	Provides information required to install the product.
Symantec Cluster Server Administrator's Guide	vcs_admin_62_sol.pdf	Provides information required for administering the product.
Symantec Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_62_sol.pdf	Provides information about bundled agents, their resources and attributes, and more related information.
Symantec Cluster Server Agent Developer's Guide (This document is available online only.)	vcs_agent_dev_62_unix.pdf	Provides information about the various Symantec agents and procedures for developing custom agents.
Symantec Cluster Server Application Note: Dynamic Reconfiguration for Oracle Servers (This document is available online only.)	vcs_dynamic_reconfig_62_sol.pdf	Provides information on how to perform dynamic reconfiguration operations on VCS clustered system domains of Oracle servers.

Table 1-12 Symantec Cluster Server documentation *(continued)*

Title	File name	Description
Symantec Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_62_sol.pdf	Provides notes for installing and configuring the DB2 agent.
Symantec Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_62_sol.pdf	Provides notes for installing and configuring the Oracle agent.
Symantec Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_62_sol.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 (This document is available online.)	sfhas_whats_new_62_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 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_62_sol.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.

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

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions Virtualization Guide (This document is available online.)	sfhas_virtualization_62_sol.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 SmartIO for Solid State Drives Solutions Guide	sfhas_smartio_solutions_62_sol.pdf	Provides information on using and administering SmartIO with SFHA solutions. Also includes troubleshooting and command reference sheet for SmartIO.
Symantec Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide (This document is available online.)	sfhas_dr_impl_62_sol.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_62_sol.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_62_ldom_sol.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_62_ldom_sol.pdf	Describes the steps for installing and configuring Symantec Cluster Server. Some of the most common troubleshooting steps are also documented in this guide.

Table 1-14 Symantec ApplicationHA documentation *(continued)*

Document title	File name	Description
Symantec ApplicationHA User's Guide	applicationha_users_62_ldom_sol.pdf	Provides information about configuring and managing Symantec Cluster Server in Oracle VM Server for SPARC (OVM) 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_62_ldom_sol.pdf	Describes how to configure application monitoring for Oracle.
Symantec ApplicationHA Generic Agent Configuration Guide	applicationha_gen_agent_62_ldom_sol.pdf	Describes how to configure application monitoring for a generic application.
Symantec Cluster Server Agent for Apache HTTP Server Configuration Guide	applicationha_apache_agent_62_ldom_sol.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