Veritas Storage Foundation™ and High Availability Solutions 6.0.5 Release Notes - Solaris

6.0.5 Maintenance Release

April 2014

Symantec™

Veritas Storage Foundation and High Availability Release Notes

Product version: 6.0.5

Document version: 6.0.5 Rev 0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

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

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 Veritas Storage Foundation and High Availability Solutions

This document includes the following topics:

Introduction

This document provides information about the products in Veritas Storage Foundation and High Availability Solutions 6.0.5 Maintenance Release (6.0.5 MR).

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

http://www.symantec.com/docs/TECH164885

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

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

Before installing or upgrading Veritas Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

For instructions to install or upgrade the product see the *Veritas Storage Foundation and High Availability Solutions 6.0.5 Installation Guide* at available on the Symantec website:

http://sort.symantec.com/documents

This Maintenance Release applies to the following releases of Storage Foundation and High Availability products:

- Storage Foundation and High Availability Solutions 6.0.1
- Storage Foundation and High Availability Solutions 6.0.3

This Maintenance Release is available as 6.0.5

List of products

Apply the patches for the following Veritas Storage Foundation and High Availability products:

- Veritas Dynamic Multi-Pathing (DMP)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)
- Veritas Storage Foundation (SF)
- Veritas Cluster Server (VCS)
- Veritas Storage Foundation and High Availability (SFHA)
- Veritas Storage Foundation Cluster File System and High Availability (SFCFSHA)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
- Veritas Storage Foundation for Sybase ASE CE (SFSYBASECE)
- Symantec VirtualStore (SVS)

List of patches

This section lists the patches and packages for 6.0.5.

Table 1-1 Patches and packages for Solaris x64

Patch ID	Package Name	Products Affected	Patch Size	Solaris 10
148482-03	VRTSvxfs	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	32 MB	X
148491-07	VRTSvxvm	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS, VM	211 MB	X
148493-02	VRTSvcsag	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	6.5 MB	X
148494-02	VRTSvcs	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	230 MB	X
148495-02	VRTSamf	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	4.8 MB	X
149692-02	VRTScavf	SF Oracle RAC, SFCFSHA, SVS	922 KB	X
149693-02	VRTSvcsea	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	1.4 MB	X
149694-02	VRTSvxfen	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	2.6 MB	X
149697-02	VRTSdbed	SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	120 MB	X
149698-02	VRTSperl	DMP,FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS, VM	54 MB	X
149703-02	VRTSsfcpi601	DMP,FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS, VM	5.0 MB	X
149817-02	VRTSodm	SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	975 KB	X
150707-01	VRTSgab	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	1.8 MB	X
150718-01	VRTSllt	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	1.6 MB	X
150721-01	VRTSdbac	SF Oracle RAC	4.2 MB	X

Table 1-2 Patches and packages for Solaris SPARC

Patch ID	Package Name	Products Affected	Patch Size	Solaris 10
148481-03	VRTSvxfs	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	37 MB	X
148490-07	VRTSvxvm	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	205 MB	X
148492-02	VRTSvcs	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	221 MB	X
148496-02	VRTSvcsag	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	6.2 MB	X
148497-02	VRTSvcsea	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	1.4 MB	X
148498-02	VRTSamf	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	4.8 MB	X
149691-02	VRTScavf	SF Oracle RAC, SFCFSHA, SFSYBASECE, SVS	890 KB	X
149695-02	VRTSvxfen	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	2.7 MB	X
149696-02	VRTSdbed	SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	115 MB	X
149699-02	VRTSperl	DMP,FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS, VM	57 MB	X
149702-02	VRTSsfcpi601	DMP,FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS, VM	5.2 MB	X

Table 1-2 Patches and packages for Solaris SPARC *(continued)*

Patch ID	Package Name	Products Affected	Patch Size	Solaris 10
149816-02	VRTSodm	SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	1.1 MB	X
150706-01	VRTSgab	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	1.8 MB	X
150710-01	VRTSllt	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	1.6 MB	X
150716-01	VRTSfsadv	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	15 MB	X
150720-01	VRTSdbac	SF Oracle RAC	4.5 MB	X

Note: You can also view the list using the `installmr` command: `./installmr -listpatches`

Changes introduced in 6.0.5

This section lists the changes in 6.0.5.

Changes in documentation in 6.0.5

The following are the changes related to documentation introduced in this release:

SFHA Release Notes content now split into separate installation and release notes documents

Maintenance releases until version 6.0.5 included both release-specific and installation content in a single release notes document. Starting with 6.0.5, future maintenance releases will deliver the following documentation with the release:

Document	Description
Veritas Storage Foundation and High Availability Solutions Release Notes	This document will contain release-specific information such as system requirements, changes in the release, fixed issues in the release, known issues and limitations in the release.
Veritas Storage Foundation and High Availability Solutions Installation Guide	This document will contain instructions specific to installing, upgrading, or uninstalling the product.

Both documents will be available on the Symantec SORT web site at the time of release:

https://sort.symantec.com/welcome/documentation

Changes related to Storage Foundation and High Availability

There are no changes related to Storage Foundation and High Availability in 6.0.5.

Changes related to installing, upgrading and rolling back

The following changes are related to installing, upgrading and rolling back of the product in 6.0.5 release.

Using Install Bundles with `-base_path` option to install or upgrade to 6.0.5 in one execution.

In version 6.0.5, Symantec offers you a method to easily install or upgrade your systems directly to 6.0.5 in one step using Install Bundles with `-base_path` option. With this option, the installers have the ability to merge base releases like 6.0.1 with 6.0.5 which is a maintenance release, so that you can install or upgrade directly to 6.0.5 in one execution. You do not have to perform two install actions to install or upgrade systems to 6.0.5.

You can get base release from FileConnect that requires customer serial number. For 6.0.5, the base release version should be 6.0.1.

You can also download 6.0.5 from the SORT website.

When you want to install or upgrade to 6.0.5 using Install Bundles with `-base_path` option, the command must be executed from the 6.0.5 install script.

For example, enter the following command:

```
./installmr -base_path <path_to_base>
```

Enhancement on VRTSaslapm package upgrade and rolling back

During the upgrade, if the version of VRTSaslapm package is earlier than 6.0.5, then the installer upgrades the VRTSaslapm package with the one in 6.0.5 release.

On Solaris 10, during the rolling back, VRTSaslapm package does not revert to the earlier version that you roll back to. If you need to use VRTSaslapm package of an earlier version, then uninstall the current VRTSaslapm package and reinstall VRTSaslapm package of a specific version after rolling back.

Local installer scripts' version suffix changed

The local installer scripts' name under /opt/VRTS/install/ is changed from [un]install<prod>601 to [un]install<prod>605. This script name change does not affect any functionality.

Changes related to Veritas Volume Manager

There are no changes related to Veritas Volume Manager in 6.0.5.

Changes related to Veritas File System

There are no changes related to Veritas File System in 6.0.5.

Changes related to Veritas Cluster Server

Veritas Cluster Server includes the following changes in 6.0.5:

- New Service Management Facility (SMF) services to avoid race conditions when you add or remove LLT driver on Solaris 11
  On Solaris 11, Symantec has added two new SMF services, llt-postinstall, and llt-preremove to manage addition and removal of LLT driver. With the addition of these new SMF services, the LLT driver is added only during package installation and removed on package removal. The new SMF services avoid failure to install the LLT driver during system restart.
- New SMF services to avoid race conditions when you add or remove GAB driver on Solaris 11
  On Solaris 11, Symantec has added two new SMF services, gab-postinstall, and gab-preremove to manage addition and removal of GAB driver. With the addition of these new SMF services, the GAB driver is added only during package installation and removed on package removal. The new SMF services avoid failure to install the GAB driver during system restart.

- New SMF services to avoid race conditions when you add or remove I/O Fencing driver on Solaris 11

  On Solaris 11, Symantec has added two new SMF services, `vxfen-postinstall`, and `vxfen-preremove` to manage addition and removal of I/O Fencing driver. With the addition of these new SMF services, the I/O Fencing driver is added only during package installation and removed on package removal. The new SMF services avoid failure to install the I/O Fencing driver during system restart.

- Packaging-related SMF services on Solaris 11

  After installing packages on Solaris 11 system, the following SMF services are present in online state. These SMF services are meant for proper package operation during uninstall operation. Symantec recommends you to not disable these services.

  ```
  svc:/system/gab-preremove:default

  svc:/system/llt-preremove:default

  svc:/system/vxfen-preremove:default
  ```

Support to have multiple zone resources within a service group and facility to update zone admin password

The `hazonesetup` command now supports creation of multiple zone resources in a service group. For more details, please refer to the `hazonesetup` man page.

When you run `hazonesetup` command to create a service group which contains more than one zone, the `ResContainerInfo` attribute is displayed for each zone resource.

Note: The `ResContainerInfo` attribute should be set for a resource, when multiple zone resources are configured within a service group.

Note: While running the `hazonesetup` command on any node, ensure that all the nodes are in running state.

The `hazoneverify` command now supports verification of multiple zone resources in a service group.

New attribute `ClearClone` added to `DiskGroup` and `CVMVolDg` agents to support `-c` option to reset `clone_disk` flag during disk group import

In this release, Symantec has introduced boolean attribute `ClearClone` to `DiskGroup` and `CVMVolDg` agents. The default value of the `ClearClone` attribute is 0. If the value of `ClearClone` attribute is set to 1, then the disk group is imported with the `-c` option. While importing the disk group, this option clears the clone and the `udid_mismatch` flags from the disks of the disk groups and also updates the UDID.

You can modify the `ClearClone` attribute using the following procedure.

To enable the `ClearClone` attribute

1 Enable the write access for VCS configuration.

```
#haconf -makerw
```

2 Set the value of `ClearClone` attribute to 1.

```
#hares -modify < resource_name > ClearClone 1
```

3 Save the VCS configuration.

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

To disable the `ClearClone` attribute

1 Enable the write access for VCS configuration.

```
#haconf -makerw
```

2 Set the value of `ClearClone` attribute to 0.

```
#hares -modify < resource_name > ClearClone 0
```

3 Save the VCS configuration.

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

New command for `hacli` in `vxfenswap` utility

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

Support for Oracle Single Instance 12c

In 6.0.5 release, Veritas Cluster Server supports Oracle Single Instance 12c.

Changes related to Veritas Storage Foundation for Oracle RAC in 6.0.5

Veritas Storage Foundation for Oracle RAC includes the following changes in 6.0.5:

This release introduces script-based installer support for configuring Highly Available IP (HAIP) addresses on SF Oracle RAC nodes running Oracle RAC 11.2.0.2 and later versions.

The Oracle Network Configuration menu now displays the following options:

- **Configure private IP addresses (HAIP Configuration)** - For Oracle RAC 11.2.0.2 and later
- **Configure private IP addresses (PrivNIC Configuration)** - For Oracle RAC prior to 11.2.0.2
- **Configure private IP addresses (MultiPrivNIC Configuration)** - For Oracle RAC prior to 11.2.0.2
- **Exit** - Exit SF Oracle RAC Configuration
- **Back** - Back to the previous menu

Oracle 12c support

In 6.0.5 release, Veritas Storage Foundation for Oracle RAC supports Oracle 12c.

Enabling health check monitoring in VCS agent for Oracle with SFHA 6.0.5

In Veritas Storage Foundation High Availability 6.0.5 release, Symantec has enabled the health check monitoring feature in VCS agent for Oracle. Please refer to the following tech note for more details:

http://www.symantec.com/docs/TECH214888

Changes related to Veritas Dynamic Multi-Pathing

Veritas Dynamic Multi-Pathing 6.0.5 includes the following changes:

SCSI-3 PR I/O Fencing support from Oracle VM Server for SPARC guests with Layered DMP configuration in Oracle VM environment

From 6.0.5 release onwards, Oracle VM Server for SPARC guest can use SCSI-3 PR fencing with virtual devices backed by DMP devices in I/O domains. Please note that the primary and the alternate I/O domains must have patch VRTSvxvm 6.1HF1 or higher installed. Please contact Symantec Technical Support to obtain VRTSvxvm patch for installing 6.1HF1.

Support for Oracle VM Server for SPARC guest Live Migration with SCSI-3 PR I/O Fencing enabled on DMP devices

After live migration of virtual machine having Storage Foundation stack with data disks fencing enabled, the I/O operations fail on shared SAN devices with reservation conflict and causes service groups to fault. Live migration causes the SCSI initiator change. Hence, the I/O operations coming from the migrated servers to the shared SAN storage fails with reservation conflict.

With this release, Oracle VM Server for SPARC guest with fencing enabled on DMP devices can be live migrated. DMP automatically takes care of re-registration of PGR keys with the new initiators. As of now, this is supported for selected arrays. Please check with Symantec Technical Support for more details.

Changes related to Veritas Storage Foundation for databases (SFDB) tools

Veritas Storage Foundation for databases (SFDB) tools includes the following changes in 6.0.5:

Veritas Storage Foundation for databases (SFDB) tools support Oracle 12c release for Oracle databases.

Note: For Oracle 12c, the SFDB tools do not support the Multitenant database features, including the CDB and PDB databases.

Changes related to Symantec Virtual Store

There are no changes related to Symantec Virtual Store in 6.0.5.

System requirements

This section describes the system requirements for this release.

Supported Solaris operating systems

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

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

Table 1-3 Supported operating systems

Operating systems	Levels	Chipsets
Solaris 10	Update 8, 9, 10, and 11	SPARC
Solaris 10	Update 8, 9, 10, and 11.	x86-64
Solaris 11	Solaris 11.1 and up to Support Repository Updates (SRU) 11.1.17.5.0 Solaris 11 and up to Support Repository Updates (SRU) 11.0.13.4.0	SPARC
Solaris 11	Solaris 11.1 and up to Support Repository Updates (SRU) 11.1.17.5.0 Solaris 11 and up to Support Repository Updates (SRU) 11.0.13.4.0	x86-64

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

If necessary, upgrade Solaris before you install the Veritas products.

See http://www.symantec.com/docs/TECH202397 before upgrading to Oracle Solaris 11.1.

See http://www.symantec.com/docs/TECH206172 before upgrading to 6.0.5 on Oracle Solaris 11.

Install all the latest required Solaris patches listed in this Release Notes.

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

http://www.symantec.com/docs/TECH75362

For Storage Foundation for Oracle RAC, all nodes in the cluster must have the same operating system version and update level.

Supported database software

For the latest information on supported Oracle database versions, see the following TechNote:

http://www.symantec.com/docs/DOC5081

Support for minor database versions is also documented in the afore-mentioned TechNote.

Additionally, see the following Oracle support site for information on patches that may be required by Oracle for each release.

https://support.oracle.com

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/TECH170013

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

For information on specific HA setup requirements, see the *Veritas Cluster Server Installation Guide*.

Veritas Storage Foundation for Database features supported in database environments

Veritas Storage Foundation for Database (SFDB) features are supported for the following database environments:

Table 1-4 SFDB features database support for 6.0.5

SFDB feature	DB2	Oracle	Sybase
Oracle Disk Manager, Cached Oracle Disk Manager	No	Yes	No
Quick I/O, Cached Quick I/O	Yes	Yes	Yes
Concurrent I/O	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes

Table 1-4 SFDB features database support for 6.0.5 *(continued)*

SFDB feature	DB2	Oracle	Sybase
Database Storage Checkpoints	Yes	Yes	No
Database Flashsnap	Yes	Yes	No
SmartTier for Oracle	No	Yes	No

Review current documentation for your database to confirm the compatibility of your hardware and software.

For the most current information on Storage Foundation products and single instance Oracle versions supported, see:

http://www.symantec.com/docs/DOC4039

Veritas Storage Foundation memory requirements

Symantec recommends 2 GB of memory over the minimum requirement for the operating system.

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

Fixed issues

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

Installation and upgrades: issues fixed in 6.0.5

This section describes the incidents that are fixed in installation and upgrades in 6.0.5.

Table 1-5 Installation and upgrades fixed issues in 6.0.5

Fixed issues	Description
3128888	When validating the upgrade procedure with an encapsulated boot disk and thedisk group minoring is already set to 1000, the installer fails.
3131744	installmp CPI should not enable DMP native support ,it should ask customer if they want to enable it.

Table 1-5 Installation and upgrades fixed issues in 6.0.5 *(continued)*

Fixed issues	Description
3230218	The Common Product Installer (CPI) based installation fails on Solaris version 11 because the nslookup command is missing.
3243089	During a live upgrade the installation process takes more time than expected.
3248893	If the Solaris 10 kernel patch is Update 10, the installer fails to detect the OS update level and issues a warning message.
3295841	CPI patches 6.0.1_P4.pl and 6.0.3_P6.pl fails when ssh banner is enabled
3304955	The installer blocks upgrading the product to 6.0.1 if Veritas Cluster Server (VCS) is not configured.
3309409	Upon installation SFHA with any hotfix in Solaris 10 Update 11, installer getting the warning message about unsupported OS version.
3432524	For Oracle 12c, the installation of Clusterware's response file fails.
3472265	The installvcs script cannot set the heartbeat NIC name to be "L_101".
3448674	After upgrading from 5.0MP3RP5 to 6.0.5 using the base_path option, NFSRestart and NFS upper or lower resource cannot come online automatically.

Installation and upgrades: issues fixed in 6.0.3

This section describes the installation and upgrade issues fixed in 6.0.3.

Table 1-6 Installation and upgrades 6.0.3 fixed issues

Incident	Description
2967125	Eval injection vulnerability in the Digest module before 1.17 for Perl allows context-dependent attackers to execute arbitrary commands via the new constructor.
2880917	Veritas product services do not get enabled after phased upgrade.

Installation and upgrades: issues fixed in 6.0.1

This section describes the incidents that are fixed related to installation and upgrades in this release.

Table 1-7 Fixed issues related to installation and upgrades

Incident	Description
2329580	Unable to stop some SFCFSHA processes.
2873102	Perl module error on completion of SFHA installation
2627076	Incorrect server names sometimes display if there is a clock synchronization issue.
2622987	sfmh discovery issue when you upgrade your Veritas product to 6.0.1
2593148	cssd agent configuration failed with CPI when have two priv IP's in setup.
2526709	DMP-OSN tunable value not get persistence after upgrade from 5.1SP1 to 6.0.
2088827	During product migration the installer overestimates disk space use.
2618482	Some agents may fail to come online after full upgrade if they were online before the upgrade.

Veritas Storage Foundation Cluster File System High Availability: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) in 6.0.5.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues

Fixed issues	Description
3259634	A Cluster File System (CFS) with blocks larger than 4GB may become corrupt.
3462694	The fsdedupadm(1M) command fails with error code 9 when it tries to mount checkpoints on a cluster.
3189562	Oracle daemons get hang with the vx_growfile() kernel function.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3214328	A mismatch is observed between the states for the Global Lock Manager (GLM) grant level and the Global Lock Manager (GLM) data in a Cluster File System (CFS) inode.
2495673	Mismatch of concurrent I/O related data in an inode is observed during communication between the nodes in a cluster.
3092114	The information output displayed by the "df -i" command may be inaccurate for cluster mounted file systems.
3224101	After you enable the optimization for updating the i_size across the cluster nodes lazily, the system panics.
2977035	A debug assert issue was encountered in vx_dircompact() function while running an internal noise test in the Cluster File System (CFS) environment
3312897	System can hang when the Cluster File System (CFS) primary node is disabled.
3274592	Internal noise test on cluster file system is unresponsive while executing the fsadm(1M) command
1949445	System is unresponsive when files were created on large directory.
2972183	The fsppadm(1M) enforce command takes a long time on the secondary nodes compared to the primary nodes.
3003679	When running the fsppadm(1M) command and removing a file with the named stream attributes (nattr) at the same time, the file system does not respond.
3072036	Read operations from secondary node in CFS can sometimes fail with the ENXIO error code.
3364312	The fsadm(1M) command is unresponsive while processing the VX_FSADM_REORGLK_MSG message.
3359200	Internal test on Veritas File System (VxFS) fsdedup(1M) feature in cluster file system environment results in a hang.
2735912	The performance of tier relocation using the fsppadm(1M)enforce command degrades while migrating a large number of files.
3153919	The fsadm (1M) command may hang when the structural file set re-organization is in progress.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3332902	While shutting down, the system running the fsclustadm(1M)command panics.
3046983	Invalid CFS node number in ".__fsppadm_fclextract", causes the DST policy enforcement failure.
3364287	Debug assert may be hit in the vx_real_unshare() function in the cluster environment.
3364301	Assert failure because of improper handling of inode lock while truncating a reorg inode.
3444775	Internal noise testing on cluster file system results in a kernel panic in function vx_fsadm_query() with an error message.
3192985	Checkpoints quota usage on Cluster File System (CFS) can be negative.
3410837	The error message has an issue when the user uses the cfsumount(1M) to unmount a mount point which has a samba parent resource.
2756779	The code is modified to improve the fix for the read and write performance concerns on Cluster File System (CFS) when it runs applications that rely onthe POSIX file-record using the fcntl lock.

Veritas Storage Foundation Cluster File System High Availability: issues fixed in 6.0.3

Table 1-9 describes the Veritas Storage Foundation Cluster File System fixed issues in 6.0.3.

Table 1-9 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues

Incident	Description
2977697	vx_idetach generated kernel core dump while filestore replication is running.
2942776	Mount fails when volumes in vset are not ready.
2923867	Internal test hits an assert "f:xted_set_msg_pri1:1".
2923105	The upgrade VRTSvxfs5.0MP4HFaf hangs at vxfs preinstall scripts.
2916691	Customer experiencing hangs when doing dedups.

Table 1-9 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues *(continued)*

Incident	Description
2906018	The vx_iread errors are displayed after successful log replay and mount of the file system.
2843635	Internal testing is having some failures.
2841059	full fsck fails to clear the corruption in attribute in ode 15.
2750860	Performance issue due to CFS fragmentation in CFS cluster.
2715175	It takes 30 minutes to shut down a 4-node cluster.

Veritas Storage Foundation Cluster File System High Availability: issues fixed in 6.0.1

This section describes the incidents that are fixed in Veritas Storage Foundation Cluster File System High Availability in this release.

Table 1-10 Veritas Storage Foundation Cluster File System High Availability fixed issues

Incident	Description
2867282	An ENOSPC error may return to the cluster file system application.
2703747	CFS failover takes up to 20 minutes due to slow log replay.
2684573	The performance of the cfsumount(1M) command for the VRTScavf package is slow when some checkpoints are deleted.

Veritas Volume Manager: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.0.5.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues

Fixed issues	Description
1942051	IO hangs on a master node after disabling the secondary paths from slave node and rebooting the slave node.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2020017	Cluster node panics when mirrored volumes are configured in the cluster.
2054606	During the DMP driver unload operation the system panics.
2101093	A system panic is observed in the dmp_signal_event() function.
2106530	The vxresize(1M) command fails on the data volume in rootdg if the file system is mounted using block device reference as 'bootdg'.
2165920	The vxrelocd(1M) daemon creates a defunct (zombie) process.
2236443	Disk group import failure should be made fencing aware, in place of VxVM vxdmp V-5-0-0 i/o error message.
2308875	vxddladm(1M) list command options (hbas, ports, targets) don't display the correct values for the state attribute.
2398954	The system panics while performing I/O on a VxFS mounted instant snapshot with the Oracle Disk Manager (ODM) SmartSync enabled.
2422535	Changes on the Veritas Volume Manager (VxVM) recovery operations are not retained after the patch or package upgrade.
2599887	The DMP device paths that are marked as "Disabled" cannot be excluded from VxVM control.
2643506	vxconfigd dumps core when LUNs from the same enclosure are presented as different types, say A/P and A/P-F.
2685230	In a Cluster Volume Replicator (CVR) environment, if the SRL is resized and the logowner is switched to and from the master node to the slave node, then there could be a SRL corruption that leads to the Rlink detach.
2735364	The "clone_disk" disk flag attribute is not cleared when a cloned disk group is removed by the "vxdg destroy dg-name" command.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2746907	The vxconfigd(1M) daemon can hang under the heavy I/O load on the master node during the reconfiguration.
2790864	For OTHER_DISKS enclosure, the vxdmpadm config reset CLIfails while trying to reset IO Policy value.
2804326	In the Veritas Volume Replicator (VVR) environment, secondary logging is seen ineffect even if Storage Replicator Log (SRL) size mismatch is seen across primary and secondary.
2812161	In a Veritas Volume Replicator (VVR) environment, after the Rlink is detached, the vxconfigd(1M) daemon on the secondary host may hang.
2825102	CVM reconfiguration and VxVM transaction code paths can simultaneously access volume device list resulting in data corruption.
2845383	The site gets detached if the plex detach operation is performed with the site-consistency set to off.
2856579	When a disk is resized from less than 1TB to greater than 1TB, "EFI PART" is missing in the primary label.
2860230	In a Cluster Volume Manager (CVM) environment, the shared disk remains as opaque after execution of vxdiskunsetup(1M) command on a master node.
2861011	The "vxdisk -g <dgname> resize diskname" command fails with an error for the Cross-platform Data Sharing(CDS) formatted disk.
2866299	The NEEDSYNC flag set on volumes in a Replicated Volume Group (RVG) not getting cleared after the vxrecover command is run.
2869514	In the clustered environment with large Logical unit number(LUN) configuration, the node join process takes long time.
2882312	If an SRL fault occurs in the middle of an I/O load, and you immediately issue a read operation on data written during the SRL fault, the system returns old data.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2882412	The 'vxdisk destroy' command uninitializes a VxVM disk which belongs to a deported disk group.
2882566	You can successfully add a disk which is removed from a disk group using thevxdg rmdisk -k command to another disk group without any error messages.
2893530	With no VVR configuration, when system is rebooted, it panicked.
2898324	UMR errors reported by Purify tool in "vradmind migrate" command.
2909668	In case of multiple sets of the cloned disks of the same source disk group, the import operation on the second set of the clone disk fails, if the first set of the clone disks were imported with "updateid".
2910367	When SRL on the secondary site disabled, secondary panics.
2912263	On Solaris LDOMs, the "vxdmpadm exclude" command fails to exclude a controller.
2916911	The vxconfigd(1M) daemon sends a VOL_DIO_READ request before the device is open. This may result in a scenario where the open operation fails but the disk read or write operations proceeds.
2921816	System panics while starting replication after disabling the DCM volumes.
2925746	In the cluster volume manager (CVM) environment, cluster-wide vxconfigd may hang during CVM reconfiguration.
2929206	When turning on the dmp_native_support tunable with Solaris10 U10 and onwards,the Zettabyte File System (ZFS) pools are seen on the OS device paths but not on the dynamic multipathing (DMP) devices.
2932214	"vxdisk resize" operation may cause the disk goes into "online invalid" state.
2933476	The vxdisk(1M) command resize fails with a generic error message. Failure messages need to be more informative.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2933688	When the 'Data corruption protection' check is activated by Dynamic Mult-Pathing (DMP), the device- discovery operation aborts, but the I/O to the affected devices continues, this results in data corruption.
2938710	The vxassist(1M) command dumps core during the relayout operation .
2950624	vradmind fails to operate on the new master when a node leaves the cluster.
2952403	Shared disk group fails to destroy if master has lost storage.
2952553	Refresh of a snapshot should not be allowed from a different source volume without force option.
2954455	During Dynamic Reconfiguration Operations in vxdiskadm, when a pattern is specified to match a range of LUNs for removal, the pattern is matched erroneously.
2957555	The vxconfigd(1M) daemon on the CVM master node hangs in the userland during the vxsnap(1M) restore operation.
2957556	The vxdisksetup(1M) command fails when the tpdmode attribute is set to native and the enclosure-based naming scheme is on.
2958983	Memory leak is observed during the reminor operations.
2959333	The Cross-platform Data Sharing (CDS) flag is not listed for disabled CDS disk groups.
2959733	Handling the device path reconfiguration in case the device paths are moved across LUNs or enclosures to prevent the vxconfigd(1M) daemon coredump.
2962010	The replication hangs when the Storage Replicator Log (SRL) is resized.
2966990	In a Veritas Volume Replicator (VVR) environment, the I/O hangs at the primary side after multiple cluster reconfigurations are triggered in parallel.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2969335	The node that leaves the cluster node while the instant operation is in progress, hangs in the kernel and cannot join back to the cluster node unless it is rebooted.
2969844	The device discovery failure should not cause the DMP database to be destroyed completely.
2972513	In CVM, PGR keys from shared data disks are not removed after stopping VCS.
2973786	Dynamic Reconfiguration (DR) tool pre-check fails for "Is OS& VM Device Tree in Sync" item if dmp_native_support is on.
2979824	The vxdiskadm(1M) utility bug results in the exclusion of the unintended paths.
2980955	Disk group (dg) goes into disabled state if vxconfigd(1M) is restarted on new master after master switch.
2986596	The disk groups imported with mix of standard and clone logical unit numbers(LUNs) may lead to data corruption.
2992667	When new disks are added to the SAN framework of the Virtual Intelligent System (VIS) appliance and the Fibre Channel (FC) switcher is changed to the direct connection, the "vxdisk list" command does not show the newly added disks even after the "vxdisk scandisks" command is executed.
2993667	Veritas Volume Manager (VxVM) allows setting the Cross-platform Data Sharing (CDS) attribute for a disk group even when a disk is missing, because it experienced I/O errors.
2996142	Data is corrupted or lost if the mapping from disk access (DA) to Data Module (DM) of a disk is incorrect.
2996443	In a cluster volume replication (CVR) environment, log ownername mismatch configuration error is seen on Slave nodes after it brings downthe master node.
2999871	The vxinstall(1M) command gets into a hung state when it is invoked through Secure Shell (SSH) remote execution.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3002498	When a disk is initialized with the "vxdisk -f init DANAME" command,vxconfigd(1M) dumps core.
3003991	The vxdg adddisk command hangs when paths for all the disks in the disk group are disabled.
3006245	While executing a snapshot operation on a volume which has 'snappoints' configured, the system panics in frequently.
3010191	Previously excluded paths are not excluded after upgrade to VxVM 5.1SP1RP3.
3011405	Execution of "vxtune -o export" command fails and displays an error message.
3012929	The vxconfigbackup(1M) command gives errors when disk names are changed.
3015181	I/O hangs on both the nodes of the cluster when the disk array is disabled.
3031796	Snapshot reattach operation fails if any other snapshot of the primary volume is not accessible.
3038382	The vxlufinish(1M) command runs 'fuser -k' on non-root filesystems, which is unexpected.
3038684	The restore daemon enables the paths of Business Continuance Volumes-Not Ready (BCV-NR) devices.
3041014	Beautify error messages seen during relayout operation.
3045033	"vxdg init" should not create a disk group on clone disk that was previously part of a disk group.
3046560	The vradmin syncrvg command fails with the error: "VxVM VVRvxrsync ERROR V-5-52-2009 Could not open device [devicename] due to: DKIOCGVTOCioctl to raw character volume failed".
3049633	In Veritas Volume Replicator (VVR) environment, the VxVMconfiguration daemon vxconfigd(1M) hangs on secondary node when all disk paths are disabled on secondary node.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3052770	The vradmin syncrvg operation with a volume set fails to synchronize the secondary RVG with the primary RVG.
3052879	Auto import of the cloned disk group fails after reboot even when source disk group is not present.
3053073	Dynamic Reconfiguration (DR) Tool doesn't pick thin LUNs in "online invalid" state for disk remove operation.
3060327	The vradmin repstatus(1M) shows "dcm contains 0 kbytes" during the Smart Autosync.
3065072	Data loss occurs during the import of a clone disk group, when some of the disks are missing and the import "useclonedev" and "updateid" options are specified.
3067452	If new LUNs are added in the cluster, and its naming scheme has the avid set option set to 'no', then DR (Dynamic Reconfiguration) Tool changes the mapping between dmpnode and disk record.
3067784	The grow and shrink operations by the vxresize(1M) utility may dump core in vfprintf() function.
3074579	The "vxdmpadm config show" CLI does not display the configuration file name which is present under the root(/) directory.
3076093	The patch upgrade script "installrp" can panic the system while doing a patch upgrade.
3084449	The shared flag sets during the import of private disk group because a shared disk group fails to clear due to minor number conflict error during the import abort operation.
3085519	Missing disks are permanently detached from the disk group because -o updateid and tagname options are used to import partial disks.
3088907	A node in a Cluster Volume Manager can panic while destroying a shared disk group.
3090488	Memory leaks occur in the device discovery code path ofVeritas Volume Manager (VxVM).

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3091916	The Small Computer System Interface (SCSI) I/O errors overflow the syslog.
3098559	Cluster File System (CFS) data corrupted due to cloned copy of logical unit numbers (LUNs) that is imported with volume asymmetry.
3099796	The vxevac command fails on volumes having Data Change Object (DCO) log. Theerror message "volume is not using the specified disk name" is displayed.
3101419	In CVR environment, I/Os to the data volumes in an RVG experience may temporary hang during the SRL overflow with the heavy I/O load.
3102114	A system crash during the 'vxsnap restore' operation can cause the vxconfigd(1M) daemon to dump core after the system reboots.
3111062	When diffsync is executed, vxrsync gets the following error in lossy networks: VxVM VVR vxrsync ERROR V-5-52-2074 Error opening socket between[HOST1] and [HOST2] -- [Connection timed out]
3114134	The Smart (sync) Autosync feature fails to work and instead replicates the entire volume size for larger sized volumes.
3116990	When you run Veritas Volume Manager (VxVM) commands such as'vxdisk scandisks', 'vxdctl enable', and 'vxconfigd restart' on write protected hardware mirror LUNs, error messages are displayed in the console logs.
3120458	In cluster volume replication (CVR) in data change map (DCM) mode, cluster-widevxconfigd hang is seen when one of the nodes is stopped.
3121380	I/O of replicated volume group (RVG) hangs after one data volume is disabled.
3122828	Dynamic Reconfiguration (DR) tool lists the disks which are tagged with Logical Volume Manager (LVM), for removal or replacement.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3125631	Snapshot creation on volume sets may fail with the error: "vxsnap ERRORV-5-1-6433 Component volume has changed".
3127543	Non-labeled disks go into udid_mismatch after vxconfigd restart.
3130353	Continuous disable or enable path messages are seen on the console forEMC Not Ready (NR) devices.
3131071	VxVM patch installation in Solaris Alternate Boot Environment(ABE) results in data corruption.
3133908	DR (Dynamic Reconfiguration) Tool throws cfgadm(1M) usagemessage while adding the LUN.
3136272	The disk group import operation with the "-o noreonline" option takes additional import time.
3138849	"ERROR: Configuration daemon is not accessible" message is displayed during the boot process.
3142315	Disk is misidentified as clone disk with udid_mismatch flag.
3143622	Initializing power-path managed disks through VxVM generates error message.
3144781	In the Veritas Volume Replicator (VVR) environment, execution of the vxrlinkpause command causes a hang on the secondary node if the rlink disconnect is already in progress.
3146955	Remote disks (lfailed or lmissing disks) go into the "ONLINE INVALID LFAILED" or"ONLINE INVALID LMISSING" state after the disk loses global disk connectivity.
3152274	The dd command to SRDF-R2 (write disable) device hangs, which causes the vm command hangs for a long time. But no issues with the Operating System (OS)devices.
3162418	The vxconfigd(1M) command dumps core due to wrong check in ddl_find_cdevno() function.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3162987	The disk has a UDID_MISMATCH flag in the vxdisk list output.
3163549	vxconfigd(1M) hangs on master node if slave node joins the master having disks which are missing on master.
3163970	The "vxsnap -g disk group syncstart volume" command isunresponsive on the Veritas Volume Replicator (VVR) DR site.
3178029	When you synchronize a replicated volume group (RVG), the diff string is over 100%.
3178182	During a master take over task, shared disk group re-import operation fails due to false serial split brain (SSB) detection.
3185471	If there are iSCSI LUNs visible to the host, then VM imports all the disk groups which are visible to the host, regardless of the noautoimport flag.
3188154	The vxconfigd(1M) daemon does not come up after enabling the native support and rebooting the host.
3199056	Veritas Volume Replicator (VVR) primary system panics in the vol_cmn_errfunction due to the VVR corrupted queue.
3222707	Dynamic Reconfiguration (DR) tool does not permit the removal of disks associated with a deported diskgroup(dg).
3225660	The Dynamic Reconfiguration (DR) tool does not list thin provisioned LUNs during a LUN removal operation.
3238397	Dynamic Reconfiguration (DR) Tool's Remove LUNs option does not restart the vxattachd daemon.
3239521	When you do the PowerPath pre-check, the DynamicReconfiguration (DR) tool displays the following error message: 'Unable to runcommand [/sbin/powermt display]' and exits.
3240858	The /etc/vx/vxesd/.udev_lock file may have different permissions at different instances.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3244217	Cannot reset the clone_disk flag during vxdg import.
3247040	vxdisk scandisks enables the PowerPath (PP) enclosure which was disabled previously.
3247094	The DR (Dynamic Reconfiguration) tool is unable to apply SMI label for newly added devices which had EFI label.
3254311	System panics when reattaching site to a site-consistent diskgroup having volumelarger than 1.05 TB
3259732	In a CVR environment, rebooting the primary slave followed by connect-disconnectin loop causes rlink to detach.
3261485	The vxcdsconvert(1M) utility failed with the error "Unable to initialize the disk as a CDS disk".
3261601	System panics when dmp_destroy_dmpnode() attempts to free an already free virtual address.
3268905	After reboot, the non-root zpools created using DMP device go into FAULTED state and DMP device state is shown as UNAVAIL.
3271595	Veritas Volume Manager (VxVM) should prevent the disk reclaim flag from getting turned off, when there are pending reclaims on the disk.
3271985	In Cluster Volume Replication (CVR), with synchronous replication, aborting a slave node from the Cluster Volume Manager (CVM) cluster makes the slave node panic.
3279932	The vxdisksetup and vxdiskunsetup utilities were failing on disk which is part of a deported disk group (DG), even if "-f" option is specified.
3280830	Multiple vxresize operations on a layered volume fail with error message "There are other recovery activities. Cannot grow volume"
3287880	In a clustered environment, if a node doesn't have storage connectivity to clone disks, then the vxconfigd on the node may dump core during the clone disk group import.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3289202	Handle KMSG_EPURGE error in CVM disk connectivity protocols.
3323548	In the Cluster Volume Replicator (CVR) environment, a cluster-wide vxconfigd hang occurs on primary when you start the cache object.
3368361	When site consistency is configured within a private disk group and CVM is up,the reattach operation of a detached site fails.
3373142	Updates to vxassist and vxedit man pages for behavioral changes after 6.0.
3376725	EMC VNX and VPLEX devices managed by PowerPath(PP) are not shown properly in thevxdisk list output.
3385753	Replication to the Disaster Recovery (DR) site hangs eventhough Replication links (Rlinks) are in the connected state.
3399131	For Point Patch (PP) enclosure, both DA_TPD and DA_COEXIST_TPD flags are set.
3400504	Upon disabling the host side Host Bus Adapter (HBA) port,extended attributes of some devices are not seen anymore.
3408320	Thin reclamation fails for EMC 5875 arrays.
3409612	The value of reclaim_on_delete_start_time cannot be set to values outside the range: 22:00-03:59
3415188	I/O hangs during replication in Veritas Volume Replicator (VVR).
3416622	The hot-relocation feature fails for a corrupted disk in the CVM environment.
3417044	System becomes unresponsive while creating a VVR TCP connection.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3424798	Veritas Volume Manager (VxVM) mirror attach operations(e.g., plex attach, vxassist mirror, and third-mirror break-off snapshot resynchronization) may take longer time under heavy application I/O load.
3435041	The "vxdmpadm settune dmp_native_support=on" CLI fails with script errors.
3435225	In a given CVR setup, rebooting the master node causes one ofthe slaves to panic.
3435475	The vxcdsconvert(1M) conversion process gets aborted for athin LUN formatted as a simple disk with Extensible Firmware Interface (EFI) format.
3441356	Pre-check of the upgrade_start.sh script fails on Solaris.
3451625	Operating System fails to boot after encapsulation of rootdisk in LDom (Logical Domain) guest.
3271315	The vxdiskunsetup command with the shred option fails to shred sliced or simple disks on Solaris X86 platform.
3258276	The system panics when the Dynamic Multi-Pathing (DMP) cache open parameter is enabled.
3250450	In the presence of a linked volume, running the vxdisk(1M) command with the -o thin, fssize list option causes the system to panic.
3250369	Execution of vxdisk scandisks command causes endless I/O error messages in syslog.
3249264	Veritas Volume Manager (VxVM) thin disk reclamation functionality causes disk label loss, private region corruption and data corruption.
3237503	System hangs after creating space-optimized snapshot with large size cache volume.
3236773	Multiple error messages of the same format are displayed during setting or getting the failover mode for EMC Asymmetric Logical Unit Access (ALUA disk array.
3235350	I/O on grown region of a volume leads to system panic if the volume has instant snapshot.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3230148	Clustered Volume Manager (CVM) hangs during split brain testing.
3218013	Dynamic Reconfiguration (DR) Tool does not delete the stale OS (Operating System device handles.
3199398	Output of the command "vxdmpadm pgrrereg" depends on the order of DMP node list where the terminal output depends on the last LUN (DMP node).
3194358	The continuous messages displayed in the syslog file with EMC not-ready (NR) LUNs.
3194305	In the Veritas Volume Replicator (VVR) environment, replication status goes in a paused state.
3189830	When you use the 'Mirror volumes on a disk' option of the vxdiskadm(1M) command for a root disk, you get an error.
3182350	VxVM volume creation or size increase hangs.
3182175	vxdisk -o thin,fssize list command can report incorrect File System usage data.
3160973	vxlist hangs while detecting foreign disk format on EFI disk.
3158320	VxVM (Veritas Volume Manager) command "vxdisk -px REPLICATED list (disk)" displays wrong output.
3156295	When DMP native support is enabled for Oracle Automatic Storage Management (ASM devices, the permission and ownership of /dev/raw/raw# devices goes wrong after reboot.
3152769	DMP path failover takes time in the Solaris LDOM environment when 1 I/O domain is down.
3147241	The pkgchk(1M) command on VRTSvxvm fails on Solaris x86.
3146715	Rlinks' do not connect with NAT configurations on Little Endian Architecture.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3139983	Failed I/Os from SCSI are retried only on very few paths to a LUN instead of utilizing all the available paths, and may result in DMP sending I/O failures to the application bounded by the recovery option tunable.
3125711	When the secondary node is restarted and the reclaim operation is going on the primary node, the system panics.
3119102	Support LDOM Live Migration with fencing enabled.
3115206	When ZFS root and the 'dmp_native_support' tunable are enabled, system panics along with a stack trace.
3107741	The vxrvg snapdestroy command fails with the "Transaction aborted waiting for io drain" error message.
3107699	VxDMP causes system panic after a shutdown/reboot.
3090667	The system panics or hangs while executing the "vxdisk -o thin,fssize list" command as part of Veritas Operations Manager (VOM) Storage Foundation (SF) discovery.
3087893	EMC TPD emcpower names are changing every reboot with VxVM.
3086627	The "vxdisk -o thin,fssize list" command fails with error message V-5-1-16282.
3081410	Dynamic Reconfiguration (DR) tool fails to pick up any disk for LUNs removal operation.
3063378	VM commands are slow when Read Only disks are presented.
3057554	The VxVM (Veritas Volume Manager) command "vxdiskunsetup -o shred =1" failed for EFI (Extensible Firmware Interface) disks on Solaris X86 system.
3056311	For release < 5.1 SP1, allow disk initialization with CDS format using raw geometry.
3026977	The Dynamic Reconfiguration (DR) option with vxdiskadm(1M) removes Logical Unit Numbers (LUNs) which even are not in the Failing or Unusable state.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3021970	A secondary node panics due to NULL pointer dereference when the system frees an interlock.
3019684	I/O hang is observed when SRL is about to overflow after the logowner switches from slave to master.
2994976	System panics during mirror break-off snapshot creation or plex detach operation in vol_mv_pldet_callback(function.
2959325	The vxconfigd(1M) daemon dumps core while performing the disk group move operation.
2957645	When the vxconfigd daemon/command is restarted, the terminal gets flooded with error messages.
2921147	udid_mismatch flag is absent on a clone disk when source disk is unavailable.
2859470	The Symmetrix Remote Data Facility R2 (SRDF-R2) with the Extensible Firmware Interface (EFI label is not recognized by Veritas Volume Manager (VxVM) and goes in an error state.
2857044	System crash on voldco_getalloffset when trying to resize filesystem.
2824977	The Command Line Interface (CLI) "vxdmpadm setattr enclosure <enclrname> failovermode" which is meant for Asymmetric Logical Unit Access ALUA) type of arrays fails with an error on certain arrays without providing an appropriate reason for the failure.
2724067	Enhance vxdisksetup CLI to specify disk label content to label all corresponding paths for DMP device.
2665425	The vxdisk -px "attribute" list(1M) Command Line Interface (CLI) does not support some basic VxVM attributes.
2567618	The VRTSexplorer dumps core in vxcheckhbaapi/print_target_map_entry.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2152830	A diskgroup (DG) import fails with a non-descriptive error message when multiple copies (clones) of the same device exist and the original devices are either offline or not available.
2091520	The ability to move the configdb placement from one disk to another using "vxdisk set <disk> keepmeta=[always|skip|default]" command.
1783763	In a Veritas Volume Replicator (VVR) environment, the vxconfigd(1M) daemon may hang during a configuration change operation.
3377383	The vxconfigd crashes when a disk under Dynamic Multi-pathing (DMP) reports device failure.
3423316	The vxconfigd(1M) daemon observes a core dump while executing the vxdisk(1M) scandisks command.
3325371	Panic occurs in the vol_multistepsio_read_source() function when snapshots are used.
3373208	DMP wrongly sends the SCSI PR OUT command with APTPL bit value as A0A to arrays.
3327842	In the Cluster Volume Replication (CVR) environment, with IO load on Primary and replication going on, if the user runs the vradmin resizevol(1M) command on Primary, often these operations terminate with error message "vradmin ERROR Lost connection to host".
3301470	All cluster volume replication (CVR) nodes panic repeatedly due to null pointer dereference in vxio.
3300418	VxVM volume operations on shared volumes cause unnecessary read I/Os.
3283525	The vxconfigd(1M) daemon hangs due to Data Change Object (DCO) corruption after volume resize.
3325122	In a Clustered Volume Replicator (CVR) environment, when you create stripe-mirror volumes with logtype=dcm, creation may fail.
3312162	Data corruption may occur on the Secondary Symantec Volume Replicator (VVR) Disaster Recovery (DR) Site.

Table 1-11 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3326964	VxVM hangs in Clustered Volume Manager (CVM) environments in the presence of FMR operations.
3332796	Getting message: VxVM vxisasm INFO V-5-1-0 seeking block #... while initializing disk that is not ASM disk.
3353211	A. After EMC Symmetrix BCV (Business Continuance Volume) device switches to read-write mode, continuous vxdmp (Veritas Dynamic Multi Pathing) error messages flood syslog. B. DMP metanode/path under DMP metanode gets disabled unexpectedly.
3372724	When the user installs VxVM, the system panics with a warning.

Veritas Volume Manager: issues fixed in 6.0.3

Table 1-12 describes the incidents that are fixed in Veritas Volume Manager in 6.0.3.

Table 1-12 Veritas Volume Manager 6.0.3 fixed issues

Incident	Description
3002770	Accessing NULL pointer in dmp_aa_recv_inquiry() caused system panic
2970368	Enhancing handling of SRDF-R2 WD devices in DMP.
2965910	vxassist dump core with the `-o ordered` option.
2964547	About DMP message - cannot load module 'misc/ted'.
2962262	Uninstallation of DMP fails in presence of other multi-pathing solutions.
2948172	Executing the `vxdisk -o thin,fssize list` command can result in panic.
2942609	Message displayed when user quits from Dynamic Reconfiguration Operations is shown as error message.
2940446	Full fsck hangs on I/O in VxVM when cache object size is very large
2935771	In the VVR environment, RLINK disconnects after master switch.
2933138	panic in voldco_update_itemq_chunk() due to accessing invalid buffer

Table 1-12 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2930569	The LUNs in 'error' state in output of 'vxdisk list' cannot be removed through DR (Dynamic Reconfiguration) Tool.
2930396	vxdmpasm command can not work on the solaris10_x64 platform
2919720	vxconfigd core in rec_lock1_5()
2919714	exit code from vxevac is zero when migrating on thin luns but FS is not mounted
2919627	Dynamic Reconfiguration tool should be enhanced to remove LUNs feasibly in bulk.
2919318	The I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing.
2916094	Enhancements have been made to the Dynamic Reconfiguration Tool(DR Tool) to create a separate log file every time DR Tool is started, display a message if a command takes longer time, and not to list the devices controlled by TPD (Third Party Driver) in 'Remove Luns' option of DR Tool.
2915751	Solaris machine panics during dynamic lun expansion of a CDS disk.
2915063	Rebooting VIS array having mirror volumes, master node panicked and other nodes CVM FAULTED
2911040	Restore from a cascaded snapshot when its source is DETACHED leaves the volume in unusable state
2910043	Avoid order 8 allocation by vxconfigd in node reconfig.
2907823	If the user removes the lun at the storage layer and at VxVM layer beforehand, DMP DR tool is unable to clean-up cfgadm (leadville) stack.
2899173	vxconfigd hang after executing the `vradmin stoprep` comand.
2898547	vradmind on VVR Secondary Site dumps core, when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2892983	vxvol dumps core if new links are added while the operation is in progress.
2886402	vxconfigd hang while executing tc ./scripts/ddl/dmpapm.tc#11
2886333	The `vxdg(1M) join` command should not allow mixing clone and non-clone disks in a DiskGroup.

Table 1-12 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2878876	vxconfigd dumps core in vol_cbr_dolog() due to race between two threads processing requests from the same client.
2875962	During the upgrade of VRTSaslapm package, a conflict is encountered with VRTSvxvm package because an APM binary is included in VRTSvxvm package which is already installed
2869594	Master node panics due to corruption if space optimized snapshots are refreshed and 'vxclustadm setmaster' is used to select master.
2866997	Initializing using vxdisksetup -i and Solaris patch 147440-20 gives error VxVM vxdisksetup ERROR V-5-2-43 : Invalid disk device for vxdisksetup
2866059	Improving error messages hit during the `vxdisk resize` operation.
2859470	SRDF R2 with EFI label is not recognized by VxVM and showing in error state
2858853	vxconfigd coredumps in dbf_fmt_tbl on the slave node after a Master Switch if you try to remove a disk from the DG
2851403	The `vxportal` and `vxfs` processes are failed to stop during first phase of rolling upgrade.
2851085	DMP doesn't detect implicit LUN ownership changes for some of the dmpnodes
2837717	The `vxdisk(1M) resize` command fails if `da name` is specified.
2836798	Prevent DLE on simple/sliced disk with EFI label
2836528	vxdisk resize is failing with error: "resize failed: New geometry makes partition unaligned"
2834046	NFS migration failed due to device reminoring.
2833498	vxconfigd hangs while reclaim operation is in progress on volumes having instant snapshots
2826125	VxVM script daemon is terminated abnormally when it is invoking with exact the same process id of the last invocation.
2815517	vxdg adddisk should not allow mixing clone & non-clone disks in a DiskGroup
2807158	On Solaris platform, sometimes system can hang during VM upgrade or patch installation.
2801962	Grow of a volume takes significantly large time when the volume has version 20 DCO (Data Change Object) attached to it

Table 1-12 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2798673	System panics in voldco_alloc_layout() while creating volume with instant DCO
2779580	Secondary node gives configuration error (no Primary RVG) after reboot of master node on Primary site.
2753954	At cable disconnect on port1 of dual-port FC HBA, paths via port2 are also marked SUSPECT
2744004	vxconfigd is hung on the VVR secondary node during VVR configuration.
2742706	Panic due to mutex not being released in vxlo_open
2715129	vxconfigd hangs during Master takeover in a CVM (Clustered Volume Manager) environment.
2692012	The vxevac move error message needs to be enhanced to be less generic and give clear message for failure.
2619600	Live migration of virtual machine having SFHA/SFCFSHA stack with data disks fencing enabled, causes service groups configured on virtual machine to fault.
2567618	VRTSexplorer coredumps in vxcheckhbaapi/print_target_map_entry
2510928	Extended attributes for SRDF luns reported as Mirror with EMC (VMAX array)
2398416	vxassist dumps core while creating volume after adding attribute "wantmirror=ctlr" in default vxassist rulefile
2273190	Incorrect setting of UNDISCOVERED flag can lead to database inconsistency
2149922	Record the diskgroup import and deport events in syslog
2000585	The `vxrecover -s` command does not start any volumes if a volume is removed whilst it is running.
1982965	vxdg import fails if da-name is based on naming scheme which is different from the prevailing naming scheme on the host
1973983	vxunreloc fails when dco plex is in DISABLED state
1903700	vxassist remove mirror does not work if nmirror and alloc is specified on VxVM 3.5
1901838	Incorrect setting of Nolicense flag can lead to dmp database inconsistency.

Table 1-12 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
1859018	The `link detached from volume` warnings are displayed when a linked-breakoff snapshot is created.
1856733	Support for FusionIO on Solaris x64
1725593	The `vxdmpadm listctlr` command has to be enhanced to print the count of device paths seen through the controller.
1289985	vxconfigd core dumps upon running the `vxdctl enable` command.

Veritas Volume Manager: issues fixed in 6.0.1

This section describes the incidents that are fixed in Veritas Volume Manager in this release. This list includes Veritas Volume Replicator fixed issues.

This section describes the incidents that are fixed in Veritas Volume Manager in this release. This list includes Veritas Volume Replicator and Cluster Volume Manager fixed issues.

Table 1-13 Veritas Volume Manager fixed issues

Incident	Description
2838059	VVR Secondary panic in vol_rv_update_expected_pos.
2832784	ESX panicked after applying a template file from GUI.
2826958	The pwwn number is not displayed in the output of command `vxdmpadm list dmpnode dmpnodename=`*`dmpnode name`*.
2818840	Enhance the `vxdmpraw` utility to support permission and "root:non-system" ownership to be set and make it persistent.
2794625	Unable to configure ASM to use DMP native block device path.
2792242	I/O hang after performing zone remove/add operations.
2774406	The svol_flush_srl_to_dv_start fails to start.
2771452	IO hung because of hung port deletion.
2763206	The `vxdisk rm` command core dumps when list of disknames is very long.
2756059	Panic in voldco_or_drl_to_pvm when volume started at boot.

Table 1-13 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2754819	Live deadlock seen during disk group rebuild when the disk group contains cache object.
2751278	The vxconfigd daemon hung on all cluster nodes during `vxsnap` operation.
2743926	DMP `restored` daemon fails to restart during system boot.
2741240	The `vxdg join` transaction failed and did not rollback to the sourcedg.
2739709	Disk group rebuild related issues.
2739601	VVR: repstatus output occasionally reports abnormal timestamp.
2737420	The `vxconfigd` daemon dumps core while onlining of the disk.
2729501	Exclude path not working properly and can cause system hang while coming up after enabling native support.
2711167	While starting replication, `vradmin` fails with error 'Cannot start command execution on Secondary'.
2710579	Do not write backup labels for CDS disk - irrespective of disk size.
2710147	Node panics in dmp_pr_do_reg during key registration with fencing enabled.
2703858	Site failure (storage and all nodes including master node) led to 'configuration daemon not accessible' error on all the sites.
2700792	SEGV in `vxconfigd` daemon during CVM startup.
2700486	The `vradmind` daemon coredumps when Primary and Secondary have the same hostname and an active Stats session exists on Primary.
2700086	EMC BCV (NR) established devices are resulting in multiple DMP events messages (paths being disabled/enabled).
2698860	The `vxassist mirror` command failed for thin LUN because statvfs failed.
2689845	After upgrade, some VxVM disks changed to error status and the disk group import failed.
2688747	Logowner local sequential I/Os starved with heavy I/O load on logclient.
2688308	Do not disable other disk groups when a re-import of a disk group fails during master take-over.
2684558	The `vxesd` daemon dumps core on startup in `libc`.

Table 1-13 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2683300	The hosts repeatedly reboot after 6.0x build was installed.
2680482	Empty `vx.*` directories are left in the `/tmp` directory.
2680343	Node panic during cur pri path update in cluster while running I/O shipping.
2679917	Corrupt space optimized snapshot after a refresh with CVM master switching.
2675538	The `vxdisk resize` command may cause data corruption.
2672401	Not able to initialize EFI disks which are larger than 1TB in size.
2664825	Disk group import fails when disk contains no valid UDID tag on config copy and config copy is disabled.
2656803	Race between `vxnetd start` and `stop` operations causes panic.
2653143	System panic while loading `vxdmp` driver during installation.
2652485	Inactive snapshot LUNs cause trespassing.
2648176	Performance difference on Master versus Slave during recovery with Data Change Object (DCO).
2645196	Campus Cluster + Hot Relocation: When a disk failure is detected, the associated disks for that site are detached and ALL disks as marked as RLOC.
2643634	Message enhancement for a mixed (non-cloned and cloned) disk group import.
2641510	Site consistency: When a disk failure is detected the associated disks for that site are detached and ALL disks as marked as RLOC.
2627126	Lots of I/Os and paths are stuck in dmp_delayq and dmp_path_delayq respectively. DMP daemon did not wake up to process them.
2626199	The `vxdmpadm list dmpnode` printing incorrect path type.
2620555	I0 hang due to SRL overflow & CVM reconfig.
2617336	Solaris patch 147440-04 panics in `vxioioctl`.
2605444	The `vxdmpadm disable/enable` operation of primary path (EFI labelled) in A/PF array results in all paths getting disabled.
2580393	Removal of SAN storage cable on any node brings Oracle Application Groups down on all nodes.
2566174	Null pointer dereference in `volcvm_msg_rel_gslock()`.

Table 1-13 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2564092	Automate the LUN provisioning (addition) / removal steps using `vxdiskadm`.
2553729	Status of the EMC Clariion disk changed to "online clone_disk" after upgrade.
2441283	The `vxsnap addmir` command sometimes fails under heavy I/O load.
2427894	Opaque disk support for VIS appliance.
2249445	Develop a tool to get the disk-related attributes like geometry, label, media capacity, partition info etc.
2240056	The `vxdg move` transaction not completing and backups fail.
2227678	The second rlink gets detached and does not connect back when overflowed in a multiple-secondaries environment.
2123677	Expanding a LUN to a size greater than 1 TB fails to show correct expanded size.
1675482	The `vxdg list` *`dgname`* command gives error 'state=new failed'.
1190117	vxdisk -f init can overwrite some of the public region contents.
2698035	Tunable values do not change as per the values applied through vxtune.

Dynamic Multi-Pathing: issues fixed in 6.0.1

This section describes the incidents that are fixed for Dynamic Multi-Pathing in this release.

Table 1-14 Veritas Dynamic Multi-Pathing fixed issues

Incident	Description
2826958	pwwn no is not displayed in the output of command "vxdmpadm list dmpnode dmpnodename=".
2818840	Enhance the vxdmpraw utility to support permission and root:non-system ownership be set and make it persistent.
2792242	I/O hang after performing zone remove/add operations.
2743926	DMP restored fails to restart during system boot in 6.0.
2729501	exclude path not working properly and can cause system hang while coming up after enabling native support.

Table 1-14 Veritas Dynamic Multi-Pathing fixed issues *(continued)*

Incident	Description
2700086	EMC BCV (NR) established devices are resulting in multiple dmp events messages (paths being disabled/enabled).
2684558	vxesd dumps core on startup in libc.
2653143	VxVM: System panic while loading vxdmp driver during installation.
2652485	Inactive snapshot luns cause trespassing.
2626199	vxdmpadm list dmpnode printing incorrect path-type.
2605444	vxdmpadm disable/enable of primary path (EFI labelled) in A/PF array results in all paths getting disabled.
2564092	[VxVM][Usability]Automate the lun provisioning (addition) / removal steps using vxdiskadm /or new VxVM CLI command.
2556467	DMP-ASM: disable all paths and reboot host cause /etc/vx/.vxdmprawdev records losing.

Veritas File System: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas File System (VxFS) in 6.0.5

Table 1-15 Veritas File System 6.0.5 fixed issues

Fixed issues	Description
2059611	The system panics due to a NULL pointer dereference while flushing bitmaps to the disk.
2439261	When the vx_fiostats_tunable value is changed from zero to non-zero, the system panics.
2667658	The 'fscdsconv endian' conversion operation fails because of a macro overflow.
2833450	The fstyp(1M) command displays a negative value for ninode on file systems more than 2 Tera Byte (TB).
2834192	You are unable to mount the file system after the full fsck(1M) utility is run.
2839871	On a system with DELICACHE enabled, several file system operations may hang.

Table 1-15 Veritas File System 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2908391	It takes a long time to remove checkpoints from the VxFS file system, when there are a large number of files present.
2926684	In rare cases, the system may panic while performing a logged write.
2963763	When the thin_friendly_alloc() and deliache_enable() functionality is enabled,VxFS may enter a deadlock.
2966277	Systems with high file system activity like read/write/open/lookup may panic the system.
2977828	The file system is marked bad after an inode table overflow error.
2982157	During internal testing, the f:vx_trancommit:4 debug asset was hit when the available transaction space is lesser than required.
2983248	The vxrepquota(1M) command dumps core.
2999493	The file system check validation fails with an error message after a successful full fsck operation during internal testing.
2999560	The 'fsvoladm'(1M) command fails to clear the 'metadataok' flag on a volume.
3031901	The 'vxtunefs(1M)' command accepts the garbage value for the 'max_buf_dat_size' tunable.
3042485	During internal stress testing, the f:vx_purge_nattr:1 assert fails.
3096834	Intermittent vx_disable messages are displayed in the system log.
3121933	The pwrite(2) function fails with the EOPNOTSUPP error.
3164418	Internal stress test on locally mounted VxFS filesytem results in data corruption in no space on device scenario while doing spilt on Zero Fill-On-Demand (ZFOD) extent
3194635	The internal stress test on a locally mounted file system exited with an error message.
3197901	Prevent duplicate symbol in VxFS libvxfspriv.a and vxfspriv.so
3214816	With the DELICACHE feature enabled, frequent creation and deletion of the inodes of a user may result in corruption of the user quota file.
3233284	FSCK binary hangs while checking Reference Count Table (RCT).

Table 1-15 Veritas File System 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3252983	On a high-end system greater than or equal to 48 CPUs, some file system operations may hang.
3253210	File system hangs when it reaches the space limitation.
3261462	File system with size greater than 16TB corrupts with vx_mapbad messages in the system log.
3291635	Internal testing found debug assert vx_freeze_block_threads_all:7c on locally mounted file systems while processing preambles for transactions.
3297840	A metadata corruption is found during the file removal process.
3298041	With the delayed allocation feature enabled on a locally mounted file system, observable performance degradation might be experienced when writing to a file and extending the file size.
3308673	A fragmented file system is disabled when delayed allocations feature is enabled.
3310755	Internal testing hits a debug assert vx_rcq_badrecord:9:corruptfs.
3323866	Some ODM operations may fail with "ODM ERROR V-41-4-1-328-22 Invalid argument"
3331045	Kernel Oops in unlock code of map while referring freed mlink due to a race with iodone routine for delayed writes.
3331093	Issue with MountAgent Process for vxfs. While doing repeated switchover on HP-UX, MountAgent got stuck.
3331109	The full fsck does not repair the corrupted reference count queue (RCQ) record.
3331419	System panic because of kernel stack overflow.
3335272	The mkfs (make file system) command dumps core when the log size provided is not aligned.
3364282	The fsck(1M) command fails to correct inode list file
3364290	The kernel may panic in Veritas File System (VxFS) when it is internally working on reference count queue (RCQ) record.
3364306	Stack overflow seen in extent allocation code path.

Table 1-15 Veritas File System 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3394803	A panic is observed in VxFS routine vx_upgrade7() functionwhile running the vxupgrade command(1M).
3426511	Unloading the VxFS modules may fail on Solaris 11 even after successful uninstallation of the VxFS package.
3436699	An assert failure occurs because of a race condition between clone mount thread and directory removal thread while pushing data on clone.
3115066	The multi-user services fail to run on a non-global zone.

Veritas File System: issues fixed in 6.0.3

Table 1-16 describes the incidents that are fixed in Veritas File System in 6.0.3.

Table 1-16 Veritas File System 6.0.3 fixed issues

Incident	Description
3018869	On Solaris 11 update 1, the fsadm command shows that the mountpoint is not a vxfs file system.
3013950	Internal assert "f:vx_info_init:2" is hit during Solaris 11 update 1 validation.
2905820	process hangs during file system access
2895743	Accessing named attributes for some files seems to be slow.
2885592	vxdump to the vxcompress file system is aborted.
2881211	File ACLs not preserved in checkpoints properly if file has hardlink.
2858683	Reserve extent attributes changed after vxrestore, only for files greater than 8192bytes.
2857751	The internal testing hits the assert "f:vx_cbdnlc_enter:1a".
2857629	File system corruption can occur requiring a full fsck of the system.
2827751	When ODM is used with non-VxVM devices high kernel memory allocation is seen.
2806466	fsadm -R resulting in panic at LVM layer due to vx_ts.ts_length set to 2GB.
2624262	fsdedup.bin hit oops at vx_bc_do_brelse.

Veritas File System: issues fixed in 6.0.1

This section describes the incidents that are fixed in Veritas File System in this release.

Table 1-17 Veritas File System fixed issues

Incident	Description
2838471	Need to add rstchown mount option to support customer use case.
2764861	Uncompress by vxcompress ignores quota limitation.
2753944	The file creation threads can hang.
2735912	The performance of tier relocation using fsppadm enforce is poor when moving a large amount of files.
2712392	Threads hung in VxFS.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2674639	The cp(1) command with the –p option may fail on a file system whose File Change Log (FCL) feature is enabled. The following error messages are displayed: cp: setting permissions for 'file_name': Input/output error cp: preserving permissions for 'file_name': No data available.
2670022	Duplicate file names can be seen in a directory.
2655788	Using cross-platform data sharing to convert a file system that has more than 32k nlinks does not update the vx_maxlink and maxlink_enable tunables.
2651922	ls -l command on local VxFS file system is running slow and high CPU usage is seen.
2600168	The -p option of cp_vxfs command does not work correctly on Solaris.
2597347	fsck should not coredump when only one of the device record has been corrupted and the replica is intact.
2566875	The write(2) operation exceeding the quota limit fails with an EDQUOT error (Disc quota exceeded) before the user quota limit is reached.
2559450	Command fsck_vxfs(1m) may core-dump with SEGV_ACCERR error.
2536130	fscdsconv fails to convert FS between specific platforms if FCL is enabled.
2272072	GAB panics the box because VCS engine HAD did not respond. The lobolt wraps around.
2086902	Spinlock held too long on vxfs spinlock, and there is high contention for it.

Table 1-17 Veritas File System fixed issues *(continued)*

Incident	Description
1529708	Formatting issue with the output of vxrepquota.

Veritas Cluster Server: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Cluster Server (VCS) in 6.0.5

Table 1-18 Veritas Cluster Server 6.0.5 fixed issues

Fixed issues	Description
1919203	Add Health check monitoring for Oracle Agent.
2848020	When IP is unplumbed or network cable is pulled, the SambaShare agent fails to detect the fault.
2979745	The MultiNICA agent is unable to detect the loss of network connectivity.
3028760	The NFSRestart resource does not start the NFS processes such as statd and lockd, during the online or offline operation.
3042450	A parent service group which is frozen and configured with online local hard dependency is brought offline when its child service group faults.
3076468	Incorrect configuration for one zone results in failure report for other zone resources.
3079893	The value of LastSuccess attribute of the service group equals the GlobalCounter value of the cluster if the resource faults while you online the service group.Hence the service group fails to come online.
3088915	The Oracle resource fails to come online in the container after VCS is upgraded from 6.0.1 to 6.0.3 version. Additionally, the monitor fails to detect the processes.
3090229	The Asynchronous Monitoring Framework (AMF) driver panics the node when the vxconfigd daemon is unresponsive.
3090710	The High Availability Daemon (HAD) starts and stops before the VxFEN driver configuration completes.
3097342	The Asynchronous Monitoring Framework (AMF) driver causes a panic in the node when AMF is being stopped.
3101761	The vcsauthserver process dumps core due to issues in VxAT library.

Table 1-18 Veritas Cluster Server 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3104071	The service group online propagate operation fails without giving proper error message.
3106493	If for some reason, kernel components of the Veritas Cluster Server (VCS) software stack are stopped and restarted in quick succession, then during a restart, the cluster communication may fail.
3112608	Resource fails to become online after switch operations fails for a service group.
3117829	A very high memory load on a system may cause a panic in the Cluster File System (CFS) node.
3125918	The Asynchronous Monitoring Framework (AMF) driver causes a panic in the node when the vxconfigd process is unresponsive.
3140359	Global Atomic Broadcast (GAB) fails to start when the gabconfig -c and gabconfig -cx commands are executed simultaneously on the same system.
3145047	The Asynchronous Monitoring Framework (AMF) driver causes a panic in the node after VXFS driver is unloaded.
3153987	In the Application agent, the clean operation is reported successful even when the CleanProgram returns a non-zero value.
3154104	For Application agent, an error message is logged when the StartProgram or StopProgram returns a non-zero value. This gives incorrect implication of the program failure.
3207663	Incorrect user privileges are set in case of incorrect use of the '-group' option in command "hauser -addprive".
3211834	CurrentLimits attribute value is not updated correctly when a service group faults.
3233895	Error message does not specify the source directory for the missing detailed monitoring script of the Db2udb agent.
3240209	During the Oracle online operation, due to an incorrect pattern match, the Oracle agent unnecessarily tries to back up the database.
3246141	The vxfenswap(1M) utility does not work in the clusters where rsh/ssh logins to other nodes are disabled.

Table 1-18 Veritas Cluster Server 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3259682	If vxconfigd daemon hangs, then the registration thread of IMFD daemon trying to get disk group status from vxconfigd daemon also hangs.Therefore, the amfregister command waiting for IMFD daemon gets stuck.
3302091	VCS and SFRAC fail to start after a reboot on Solaris version 11.
3318764	Unexpected deletion of temporary files causes the VCS agents to report an incorrect state.
3326591	The IPMultiNICB agent delays bringing IPv4 VIP online on Solaris version 10 by 5 seconds.
3338946	The Process resource fails to register for offline monitoring with the AMF kernel driver.
3341320	The "Cannot delete event (rid %d) in reaper" error message is repeatedly logged in the Syslog file.
3347536	The Application agent may dump a core while registering the resource with Asynchronous Monitoring Framework (AMF).
3362108	The system panics if LLT receives a corrupt packet from the network.
3389647	The hazonesetup and hazoneverify utilities do not support multiple zone resources in the same service group.
3409593	The ASMDG agent shows offline before volumes are released and service group fail-over will be delayed because volumes won't stop.
3422904	The Zone agent does not handle a zone in unavailable state.
3443419	VCS does not authenticate non-root users using halogin or Java-based GUI.

Veritas Cluster Server: issues fixed in 6.0.3

Table 1-19describes the incidents that are fixed in Veritas Cluster Server in 6.0.3.

Table 1-19 Veritas Cluster Server 6.0.3 fixed issues

Incident	Description
3025931	Sometimes gab module does not get loaded after system reboot automatically. We have to remove the driver manually and add it back to make it work.
3013962	Monitor fails to detect DB2 resource online for DB2 version 10.1.

Table 1-19 Veritas Cluster Server 6.0.3 fixed issues *(continued)*

Incident	Description
3013940	If DB2 is installed in NON-MPP mode and UseDB2Start attribute is set to 0, we still use db2start command to start the DB2 process instead of using db2gcf command.
2980002	High availability support for logical domain when control domain is down with multiple I/O domains in Oracle VM (OVM) for SPARC environment.
2979745	`MultiNICA` is unable to detect Network connectivity lost.
2967536	MonitorProgram of Application does not work correctly with C shell of User environment.
2964772	If you take an NFSRestart resource offline, the NFSRestart agent may unexpectedly stop NFS processes in a local container Zones.
2941155	Group is not marked offline on faulted cluster in a GCO environment after a cluster failure is declared.
2937673	AMF driver panics the machine when amfstat is executed.
2861253	In vxfen driver log statement, jeopardy membership is printed as garbage.
2848009	AMF panicks the machine when an Agent is exiting.
2848005	VCS SMF service goes into maintenance state on a two-node cluster.
2737653	Incorrect descriptions about the RVGPrimary online script .
2736627	REMOTE CLUSTER STATE remains in INIT state and Icmp heartbeat status is UNKNOWN.

Veritas Cluster Server: Issues fixed in 6.0.1

This section describes Veritas Cluster Server fixed issues in 6.0.1.

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

Table 1-20 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-20 LLT, GAB, and I/O fencing fixed issues

Incident	Description
2845244	`vxfen` startup script gives error `grep: can't open /etc/vxfen.d/data/cp_uid_db`. The error comes because vxfen startup script tries to read a file that might not be present. This error is typically seen when starting vxfen for the very first time after installation.
2554167	Setting peerinact value to 0 in the `/etc/llttab` file floods the system log file with large number of log messages.
2699308	Vxfenswap fails when *LANG* is set to a value other than 'C'. The vxfenswap utility internally uses the `tr` command. If the *LANG* environment variable is set to something other than C, it may cause improper functioning of the vxfenswap utility.
2726341	On Solaris 11, `vxfen` startup scripts do not report the correct status of fencing .
2850926	Fencing may start up with an error message `open failed for device: /dev/vxfen` in the log. It happens when the fencing startup script tries to access the driver that is still loading into memory. However, fencing comes up seamlessly in spite of the error message.
2699291	Logs report errors related to `mv` command when the `vxfen` service is disabled.
2762660	Post-install script of the VRTSllt package reports error while attempting to disable the SMF service system/llt.
2762660	Post-install script of the VRTSvxfen package reports error while attempting to disable the SMF service system/vxfen.
2852863	Veritas Cluster Server may not come up after rebooting the first node in phased upgrade on Oracle Solaris 11.

Bundled agents fixed issues in 6.0.1

Table 1-21 lists the fixed issues for bundled agents.

Table 1-21 Bundled agents fixed issues

Incident	Description
2509227	You can configure the MultiNICB agent to perform link-based detection and probe-based detection by setting the LinkTestRatio attribute to 1 (one) and the IgnoreLinkStatus attribute to 0 (zero). However, when you set these values, the agent may fail to send out ICMP requests to determine the resource state. As a result, the agent may report an erroneous resource state and probe-based detection may fail.
2850904	Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and PanicSystemOnDGLoss is set to 0.
2730451	On an Oracle Solaris 11 system, when you configure IP resource in a shared IP zone of type solaris brand, the IP resource does not go online.
2850923	In cases where the Zpool is corrupted and is unable to be imported cleanly even after using the `-f` and `-F` option the Zpool agent reports as ONLINE. The health of the Zpool is degraded and the filesystem reports I/O errors.
2639181	The clean entry point for the Mount agent fails to unmount a resource of type NFS after a halt.
2850924	Handle errors seen when `zfs list` command fails in Zpool agent.
2850925	The `pkg verify` command displays error messages.
2794175	Fire drill for Mount agent does not accurately portray a failover scenario.
2728802	If the 'httpd' binary or the 'ab' binary is not present at the location that you specified in the 'httpdDir' attribute, the Apache agent cannot perform detail monitoring or start the HTTP server.
2703707	Whenever Apache resource monitor runs and if the zone is still booting, warning messages may be seen due to stale credentials. It is safe to ignore these messages.
2509227	You may be unable to configure the MultiNICB agent to perform probe-based detection of resource state.
2680428	When you configure an IPMultiNICB resource for a Solaris zone, agent fails to plumb the options.
2730979	In IPMP mode, when `if_mpadm` command to disable interface fails, IPMultiNICB agent may report resource as faulted.

Table 1-21 Bundled agents fixed issues *(continued)*

Incident	Description
2714464	Using only spaces in an attribute value may cause issues with the related VCS agent.
2850905	IMF registration for Mount resource for file systems type other than VxFS and NFS should be blocked.
2850916	Mount resource does not get registered with IMF if the attributes BlockDevice and/or MountPoint have a trailing slash in their values.
2822920	DNSAgent goes to UNKNOWN state if the Top Level Domain (TLD) is more than 4 characters in length.
2846389	In releases prior to VCS 6.0.1, the upper bound value of FaultTolerance attribute of the CoordPoint agent was the one less than the number of coordination points. If the majority number of coordination points fault, the entire cluster panicked under network partition scenario. Therefore, the upper bound value of the FaultTolerance attribute of CoordPoint agent had to be set to less than the majority of the coordination points. Subsequent to VCS 6.0.1, the FaultTolerance attribute of CoordPoint agent is less than the majority of coordination points.

VCS engine fixed issues in 6.0.1

Table 1-22 lists the fixed issues for VCS engine.

Table 1-22 VCS engine fixed issues

Incident	Description
2879413	You may see two instances of CmdServer running on a node. One of these using IPv4 and the other IPv6.
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility `hagrp` gives incorrect output with the "-clear", "-flush", "-state" options.
2741299	CmdSlave gets stuck in a tight loop when it gets an EBADF on a file descriptor(fd). The CmdSlave process keeps retrying on the FD and eventually dumps core.
2850906	If a group is auto-enabled, the engine clears the Start attribute even if the resource is online.
2692173	Engine does not check whether remote parent is online when –nopre option is selected.

Table 1-22 VCS engine fixed issues *(continued)*

Incident	Description
2684818	If the following attributes are specified before SystemList attribute in main.cf, then the value got rejected when HAD started: ■ PreOnline ■ ContainerInfo ■ TriggersEnabled ■ SystemZones
2696056	Memory leak occurs in the engine when haclus –status <cluster> command is run.
2746802	When failover group is probed, VCS engine clears the MigrateQ and TargetCount.
2746816	The syslog call used in gab_heartbeat_alarm_handler and gabsim_heartbeat_alarm_handler functions is not async signal safe.

Enterprise agents fixed issues in 6.0.1

Table 1-23 lists the fixed issues for enterprise agents.

Table 1-23 Enterprise agents fixed issues

Incident	Description
1985093	Ensure that the ohasd process has an entry in the init scripts so that when the process is killed or the machine is rebooted, this automatically restarts the process.
2831044	Sybase agent script entry points must handle large process command line.
2850920	Uninstallation of VRTSvcsea package must not fail even if /opt/VRTS/messages is missing due to uninstallation of CCSMH 5.2 or 5.1 version.

Agent framework fixed issues in 6.0.1

Table 1-24 lists the fixed issues for agent framework.

Table 1-24 Agent framework fixed issues

Incident	Description
1786742	The output of `hares -action` is displayed in English text and not in your configured locale.
2660011	Resource moves to FAULTED state even if value of ManageFaults attribute is set to NONE at service group level. This will cause service group to fault if the resource is Critical.

Fixed issues related to AMF in 6.0.1

Table 1-25 AMF fixed issues

Incident	Description
2632569	Even after AMF module is stopped, it gets loaded automatically when accessed.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 6.0.5

This section describes the incidents fixed in Veritas Storage Foundation for Oracle RAC in 6.0.5.

Table 1-26 Veritas Storage Foundation for Oracle RAC 6.0.5 fixed issues

Fixed issues	Description
3090447	The CRSResource agent does not support the C shell (csh) environment.

Veritas Storage Foundation for Oracle RAC 6.0.3 fixed issues

There are no issues fixed in SF Oracle RAC 6.0.3.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 6.0.1

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 6.0.1.

Issues fixed in 6.0.1

Table 1-27 lists the issues fixed in 6.0.1.

Table 1-27 Issues fixed in 6.0.1

Incident	Description
2585899	The SF Oracle RAC installer does not support the use of fully qualified domain names (FQDN). Specifying the fully qualified domain name of a system results in the following error: `The node sys1 doesn't seem to be part of the cluster, or CVM is not running on the node sys1.`
2329580	If you install and start SFHA, but later configure SFHA using `installvcs`, some drivers may not stop successfully when the installer attempts to stop and restart the SFHA drivers and processes. The reason the drivers do not stop is because some dependent SFHA processes may be in the running state.
2392741	Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool. If the cardinality of a policy-managed Oracle RAC database is set to a number lesser than the number of nodes in the server pool, and if the Oracle agent tries to bring the database online on all the nodes in the server pool, the operation fails on some of the nodes in the server pool. The resource on respective nodes move to the faulted state.
2749412	Setting the `UseVirtualIP` attribute to 1 overwrites the IP address of the virtual interface on some nodes in the cluster.
2757032	The PrivNIC/MultiPrivNIC agents fail to match the exact IP address configured in the agent configuration with the IP address configured on the system. As a result, the agent detects the wrong interface as the active interface resulting in a resource fault.
2580393	Removal of SAN cable from any node in a global cluster setup takes application service groups offline on all nodes. In a replicated global cluster setup, the removal of SAN cable from any node in the cluster causes the CFS mount points to fault. As a result, dependent application groups are taken offline and replication to the secondary site is adversely affected.
2734745	The PrivNIC resource faults after the `UseVirtualIP` attribute is set to 1.
2740150	The SF Oracle RAC installer fails to set the value of the CSSD resource attribute `OfflineWaitLimit` to 3.

Table 1-27 Issues fixed in 6.0.1 *(continued)*

Incident	Description
2746948	Some drivers fail to add to the system. Sometimes during bootup, some of the drivers fail to add in the system because of add_drv/rem_drv race between our modules which are independent of each other.
2729806	The MultiPrivNIC agent plumbs network interfaces on the global zone.

Veritas Storage Foundation for databases (SFDB) tools: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Storage Foundation for databases (SFDB) tools in 6.0.5

Table 1-28 Veritas Storage Foundation for databases (SFDB) tools 6.0.5 fixed issues

Fixed issues	Description
2715323	The DBED operations may not work with the non-standard Oracle database character sets like ZHS16GBK.
3237852	Oracle 12c database is not supported. SYMPTOM: Oracle 12c database is not supported.
3290416	Some DBED operations may fail with the following error message: "ORA-01406: fetched column value was truncated".
3211388	While cloning a Veritas Database Edition (DBED) instant checkpoint, if you enable the Block Change Tracking feature, the error message ORA-00600 is displayed.

Veritas Storage Foundation for Databases (SFDB) tools: issues fixed in 6.0.3

Table 1-29describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 6.0.3.

Table 1-29 Veritas Storage Foundation for Databases (SFDB) tools 6.0.3 fixed issues

Incident	Description
3030663	`dbed_vmclonedb` does not read pfile supplied by `-p 'pfile_modification_file'` option.

Veritas Storage Foundation for databases (SFDB) tools: Issues fixed in 6.0.1

This section describes Veritas Storage Foundation for databases (SFDB) tools fixed issues in 6.0.1.

Storage Foundation for Databases (SFDB) tools: issues fixed in 6.0.1

Table 1-30 describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-30 SFDB tools fixed issues

Incident	Description
2585643	If you provide an incorrect host name with the `-r` option of `vxsfadm`, the command fails with an error message similar to one of the following: `FSM Error: Can't use string ("") as a HASH ref while "strict refs" in use at /opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 776. SFDB vxsfadm ERROR V-81-0609 Repository location is invalid.` The error messages are unclear.
2703881 (2534422)	The FlashSnap validation operation fails with the following error if the mirrors for data volumes and archive log volumes share the same set of disks: `SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.`
2582694 (2580318)	After you have done FlashSnap cloning using a snapplan, any further attempts to create a clone from the same snapplan using the `dbed_vmclonedb` continue to use the original clone SID, rather than the new SID specified using the *new_sid* parameter. This issue is also observed when you resynchronize the snapplan, take a snapshot again without specifying the new clone SID, and then try to clone with the new SID.

Table 1-30 SFDB tools fixed issues *(continued)*

Incident	Description
2579929	The `sfae_auth_op -o auth_user` command, used for authorizing users, fails with the following error message: `SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username>` The authentication setup might have been run with a strict umask value, which results in the required files and directories being inaccessible to the non-root users.

Symantec Virtual Store: Issues fixed in 6.0.5

There are no fixed issues for Symantec Virtual Store in this release.

Symantec Virtual Store: issues fixed in 6.0.3

There are no Symantec Virtual Store fixed issues in 6.0.3.

Symantec Virtual Store: issues fixed in 6.0.1

There are no Symantec Virtual Store fixed issues in 6.0.1.

Known issues

This section covers the known issues in this release.

- Issues related to installation and upgrade
- Veritas Volume Manager known issues
- Veritas Cluster Server known issues
- Veritas Dynamic Multi-pathing known issues
- Veritas Storage Foundation known issues
- Veritas Storage Foundation and High Availability known issues
- Veritas Storage Foundation Cluster File System High Availability known issues
- Veritas Storage Foundation for Oracle RAC known issues
- Veritas Storage Foundation for Sybase ASE CE known issues
- Symantec VirtualStore known issues

Issues related to installation and upgrade

This section describes the known issues during installation and upgrade.

The installer fails to unload all modules after uninstallation of SF packages [3348829]

During uninstallation of SF packages, some SF modules fail to get unloaded, which causes the installer to error out with the following error messages:

```
Symantec Storage Foundation Cluster File System HA Shutdown did not complete

fdd failed to stop on machine_name
vxfs failed to stop on machine_name
```

The issue is seen with the recent versions (SRUs, updates) of Solaris 11 OS, such as Solaris 11 Update 1 + SRU 11.4 or 12.5. Earlier OS versions don't have the issue.

Workaround: Restart the system.

VRTSvxfs verification reports error after upgrading to 6.0.5 [3463479]

Upgraded to 6.0.5, the VRTSvxfs package cannot pass the verification check with the pkg verify VRTSvxfs command. You can see error messages similar to the following:

```
# pkg verify VRTSvxfs
PACKAGE                                         STATUS
pkg://Symantec/VRTSvxfs                         ERROR
        driver: vxfs
                etc/name_to_major: 'vxfs' entry not present
```

Workaround: Use the following command to fix this issue:

```
# pkg fix VRTSvxfs
```

After Live Upgrade to Solaris 10 Update 10/Update 11, boot from 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 Reboot the system.

During Live Upgrade, installer displays incorrect message about VRTSaa package removal (1710504)

If you use Live Upgrade to upgrade SFHA 5.0MP1 to SFHA 6.0.1, the installer may display a message that the VRTSaa package failed to uninstall.

Workaround: Verify whether the `VRTSaa` package was removed correctly from the alternate boot disk.

```
# pkginfo -R alternate_root_path VRTSaa
```

For example, run the following command

```
# pkginfo -R /altroot.5.10 VRTSaa
```

If the VRTSaa package was removed, you can ignore this error.

If the VRTSaa package was not removed, remove the package manually:

```
# pkgrm -R alternate_root_path VRTSaa
```

For example, run the following command

```
# pkgrm -R /altroot.5.10 VRTSaa
```

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.

The `vxlufinish` command fails during upgrade to Solaris 10 Update 11 [2939321]

The `luumount` command and the `luactivate` command are internally used by`vxlufinish` for Live upgrade on Solaris. The behavior of these commands changed in Solaris 10 Udate 11, resulting in Live Upgrade failure with following error message:

```
Generating boot-sign for ABE <dest.num3>
/bin/rmdir: directory "/tmp/.liveupgrade.num1.num2/.alt.luactivate": \
Directory is a mount point or in use Generating partition and slice \
information for ABE <dest.num3>
/tmp/.liveupgrade.num1.num2/.alt.luactivate
ERROR: The target boot environment <dest.num3> root device \
<device-path> is already mounted.
ERROR: The root slice <> of the target boot environment <dest.num3>\
is not available.
rm: Unable to remove directory /tmp/.liveupgrade.num1.num2/.\
alt.luactivate:
Device busy
rm: Unable to remove directory /tmp/.liveupgrade.num1.num2: File exists
ERROR: vxlufinish Failed: luactivate dest.num3
```

Workaround: This is an issue with Oracle Solaris 10 Update 11 (x86). To resolve the issue, install the operating system patch 121430-88 or later.

Live Upgrade to 6.0.1 on Solaris 10 with dmp_native_support enabled fails (2632422)

During Live Upgrade to 6.0.1 on Solaris 10, the `vxlustart` command fails if `dmp_native_support` is enabled. Veritas Dynamic Multi-Pathing (DMP) support for native devices requires that the naming scheme be set to enclosure-based

naming (EBN). DMP 6.0.1 does not allow changing the naming scheme from EBN when support for native devices is enabled.

Due to a bug in DMP 5.1 Service Pack 1 (5.1SP1), the naming scheme could be set to operating-system based naming (OSN). However, this is not a supported configuration. With the naming scheme set to OSN, the `vxlustart` command fails.

Workaround: Disable `dmp_native_support` on all nodes.

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 SFHA 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.

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.

Installed 5.0 MP3 without configuration, then upgrade to 6.0.1, installer cannot continue (2016346)

If you install the 5.0 MP3 release without configuring the product, then you cannot upgrade to the 6.0.1 release. This upgrade path is not supported.

Workaround: Uninstall 5.0 MP3, and then install 6.0.1.

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: Install Oracle patch 113280-10,121430-72 or higher before running `vxlustart`.

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
```

Flash Archive installation not supported if the target system's root disk is encapsulated

Symantec does not support SFHA installation using Flash Archive if the target system's root disk is encapsulated.

Make sure that the target system's root disk is unencapsulated before starting the installation.

After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon you change the locale of nodes that use it. The vxconfig daemon starts at boot. If you have changed locale, you need to restart the daemon.

Workaround: See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more information on vxconfigd daemon recovery.

After performing a manual rolling upgrade, make sure the CVM is online on all nodes without errors (2595441)

Make sure that the CVM is online on all nodes without errors after you perform the first phase of a manual rolling upgrade. The CVM protocol version will not upgrade successfully on the nodes where CVM is offline or has errors.

If the CVM protocol version does note upgrade successfully, upgrade the CVM protocol on the CVM master node.

To upgrade the CVM protocol on the CVM master node

1 Find out which node is the CVM master:

```
# vxdctl -c mode
```

2 On the CVM master node, upgrade the CVM protocol:

```
# vxdctl upgrade
```

Verification of Oracle binaries incorrectly reports as failed during Oracle Grid Infrastructure installation

The verification of Oracle binaries may incorrectly report as failed during the Oracle Grid Infrastructure installation using the SFHA installer. The message is erroneously reported due to a break in passwordless SSH communication. The SSH communication fails because execution of the `root.sh` script changes the owner of the operating system root directory to the grid user directory.

The Configure Sybase ASE CE Instance in VCS option creates duplicate service groups for Sybase binary mount points (2560188)

The CPI installer does not check to see if Sybase binary mount points are already configured on systems, nor does it give an error message. It creates a duplicate service group for Sybase binary mount points.

This issue will be resolved in a later release.

Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

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

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

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.

2. Set the product level to *NONE* with the command:

   ```
   # vxkeyless set NONE
   ```

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

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

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

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

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

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

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

Instaling 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.

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.

VCS Zone users must be added after upgrade

If you upgrade your configuration containing Zone resources 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, 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 *Veritas Storage Foundation and High Availability Solutions Virtualization Guide* for Solaris for more information on `hazonesetup` utility and see the *Veritas Storage Foundation and High Availability Solutions Virtualization Guide* for Solaris.

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.

Lack of install script to configure product after installing SFCFSHA 6.0.1 (2978290)

For Solaris11 U1, there are no install scripts in the `/opt/VRTS/install` directory to configure the product.

Workaround:

On Solaris 11 U1, install 6.0.1, upgrade to 6.0.5 and then configure the product using the newly generated `install` scripts under `/opt/VRTS/install`.

The svc:/system/VRTSperl-runonce:default SMF service in the maintenance state (3089938)

After upgrade from 6.0.1 to 6.0.5 on Solaris 11 SPARC, the svc:/system/VRTSperl-runonce:default SMF service may be in the maintenance state.

Workaround:

If you find the svc:/system/VRTSperl-runonce:default SMF service in the maintenance state, perform the following steps:

1. `# svcadm disable svc:/system/VRTSperl-runonce:default`

2. `# svccfg delete -f svc:/system/VRTSperl-runonce:default`

Also you can ignore the following message from the console:

```
svc.startd[11]: svc:/system/VRTSperl-runonce:default: Method
"/opt/VRTSperl/bin/runonce" failed with exit status 127
```

SF failed to upgrade when Application HA is installed (3088810)

If you have installed both ApplicationHA 6.0 and SF 6.0.1, the installer can't upgrade SF 6.0.1 to 6.0.5. The following error message is displayed:

```
CPI ERROR V-9-30-1303 SFHA 6.0.100 does not appear to be installed
on <your system>
```

Workaround:

Use the following command to specify the exact product for the upgrade:

```
# ./installmr -prod SF60
```

Ignore the warning message that SFHA is already installed after the pre-check and continue the upgrade.

File system check daemon fails to restart after abnormal termination (2689195)

The file system check daemon (`vxfsckd`) fails to update the `vxfsckd-pid` file with the new process ID (pid) of the vxfsckd process after abnormal termination. As a result, the CFSfsckd agent fails to detect the status of the vxfsckd daemon.

Workaround: Perform the following steps to resolve the issue on the node where the `vxfsckd` resource faults:

1. Log into the node as the root user.
2. Kill all `vxfsckd` processes:

   ```
   # kill -9 `ps -ef|grep vxfsckd|awk '{print $2}'`
   ```

3. Remove the `vxfsckd-pid` file:

   ```
   # rm /var/adm/cfs/vxfsckd-pid
   ```

4. Bring the `vxfsckd` resource online:

   ```
   # hares -online vxfsckd_resname -sys node_name
   ```

Partial upgrade causes “installmr –version” to detect wrong product version [3438634]

On Solaris 10, if any packages of the product are upgraded to higher version by installer or manually, the `./installmr -version` command detects the product as a 6.0.5 product.

Workaround: No workaround.

Veritas Volume Manager known issues

This section describes the known issues in this release of Veritas Volume Manager (VxVM).

Vradmin verifydata reports differences in case of mirrored or layered volumes with SmartMove enabled [3426434]

When the SmartMove utility is enabled, mirrored or layered volumes plexes are not fully synced. The `vradmin verifydata` command compares the checksum block wise, and reports differences on mirrored or layered volumes. The following error message is displayed:

```
VxVM VVR vxrsync INFO V-5-52-10190 \
Verification of the remote volumes found differences.
```

Workaround: No workaround. Since it does not relate any data corruption, it is safe to ignore the message. You may want to use file checksum to verify whether the volumes are same.

Reclamation of storage fails with error (3422185)

Reclamation of storage fails with the following error:

```
Reclaiming storage on:
Disk <disk_name>: Failed. Failed to reclaim /<fs_mnt>, /<fs_mnt> and
<fs_mnt>.
```

Where *<fs_mnt>* is the file system mount point.

<disk_name> is the name of the disk.

The error occurs in the following scenario: A disk group is created of thin disks and some volumes are created of different layouts on this disk group. A volume set is created of these volumes and a file system is mounted over it. When some volumes of this volume set are removed, reclamation fails.

Workaround: There is no workaround for this issue.

Duplicate disk access (da) entries on vxdisk list (2705055)

If there is a change in the naming scheme and some disks in the disk group are not accessible, then duplicate disk access (da) entries will be visible on the same node.

Workaround: Perform the following steps to resolve the issue:

1. Remove the duplicate disk entry.

```
# vxdisk rm duplicate_da_name
```

2. Verify the disk.

```
# vxdisk scandisks
```

The VVR vradmin addsec command fails on Solaris 11 [3418222]

On Oracle Solaris 11, the `/etc/hosts` file is not updated to map the system's nodename to one of the system's non-loopback IP address. Instead, the host name is mapped to a system's IPv4 and IPv6 addresses. For example:

```
::1 foobar localhost
127.0.0.1 foobar loghost localhost
```

This is a Solaris 11 issue. For details, see http://docs.oracle.com/cd/E23824_01/html/E24456/gllcs.html.

Workaround: Manually modify the `/etc/hosts` file as follows:

```
:1 localhost
127.0.0.1 loghost localhost
129.148.174.232 foobar
```

The Dynamic Reconfiguration tool is not supported inside the guest domain for Oracle VM Server for SPARC (3405223)

SFHA provides a Dynamic Reconfiguration tool to simplify online dynamic reconfiguration of a LUN. The Dynamic Reconfiguration tool is not supported if SFHA is running inside the guest domain for Oracle VM Server for SPARC.

Dynamic LUN expansion is not supported for EFI disks in simple or sliced formats (2836798)

Dynamic LUN expansion is not supported for EFI (Extensible Firmware Interface) disks in simple or sliced formats. The recommended format is the Cross-platform Data Sharing (CDS) disk format.

Workaround:

Convert the disk format to CDS using the `vxcdsconvert` utility.

The vxrecover command does not handle RAID5 volumes correctly (2715124)

The `vxrecover` command calls the recovery process for the top-level volume, which internally takes care of recovering its subvolumes. The `vxrecover` command does not handle RAID5 volumes correctly. The recovery process fails to recover the subvolumes, which remain in the NEEDSYNC state.

Workaround:

Manually recover the RAID5 volumes using the `vxvol` utility, as follows:

```
# vxvol -g diskgroup resync volume
```

The vxcdsconvert utility is supported only on the master node (2616422)

The `vxcdsconvert` utility should be run only from the master node, not from the slave nodes of the cluster.

Veritas Volume Manager (VxVM) might report false serial split brain under certain scenarios (1834513)

VxVM might detect and report a false serial split brain when all of the following conditions are met:

- One or more arrays that provide the shared storage for the cluster are being powered off
- At the same time when the arrays are being powered off, an operation that requires an internal transaction is initiated (such as VxVM configuration commands)

In such a scenario, disk group import will fail with a split brain error and the vxsplitlines output will show 0 or 1 pools.

Workaround:

To recover from this situation

1 Retrieve the disk media identifier (dm_id) from the configuration copy:

```
# /etc/vx/diag.d/vxprivutil dumpconfig device-path
```

The dm_id is also the serial split brain id (ssbid)

2 Use the dm_id in the following command to recover from the situation:

```
# /etc/vx/diag.d/vxprivutil set device-path ssbid=dm_id
```

Probing vxio with DTrace fails on SPARC machines. (2180635)

This issue exists because of inability of DTrace to load a module whose text size is greater than 2MB on SPARC machines. While trying to load `vxio` with DTrace you may see following warning messages on console:

```
dtrace: WARNING: couldn't allocate SDT table for module vxio
fbt: WARNING: couldn't allocate FBT table for module vxio
```

Workaround:

There is no workaround for this issue.

The "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set (2354560)

In Cluster Volume Manager environment, "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set. This can happen on master/slave nodes.

Workaround:

Administrator has to run "vxdisk scandisks" or "vxdisk -o alldgs list" followed by "vxdg listclone" to get all the disks containing "clone_disk" or "udid_mismatch" flag on respective host.

The vxsnap print command shows incorrect value for percentage dirty (2360780)

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In SFHA 6.0, if this command is run while the volumes are online and being actively used, the shown %dirty may lag from actual percentage dirty for instant snap data cache object (DCO) volumes. That is, the command output may show less %dirty than actual.

Issues with the disk state on the CVM slave node when vxconfigd is restarted on all nodes (2615680)

When a CVM master node and a slave node have lost storage access, and `vxconfigd` is restarted on all nodes, the disk state on the CVM slave node shows as invalid.

Workaround:

To work around this issue

1. Restore storage connectivity.
2. Deport the disk group.
3. Import the disk group.

vxmirror to SAN destination failing when 5 partition layout is present: for example, root, swap, home, var, usr (2815311)

The `vxmirror` command may fail with following error on a Solaris 10 host, for a thin LUN, if more than one partition excluding root and swap is present.

```
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
```

Example:

```
# /etc/vx/bin/vxmirror" -f -g rootdg_17_23_49 rootdisk01 \
rootdisk02
! vxassist -g rootdg_17_23_49 mirror swapvol rootdisk02
! vxassist -g rootdg_17_23_49 mirror rootvol rootdisk02
! vxassist -g rootdg_17_23_49 mirror usr rootdisk02
! vxassist -g rootdg_17_23_49 mirror var rootdisk02
! vxassist -g rootdg_17_23_49 mirror home rootdisk02
! vxbootsetup -g rootdg_17_23_49
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'home_dcl'
because no free partitions are available on disk 'disk_0'.
Either remove the volume or make a partition available
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'usr_dcl'
because no free partitions are available on disk 'disk_0'.
Either remove the volume or make a partition available
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'var_dcl' because
no free partitions are available on disk 'disk_0'.
```

```
Either remove the volume or make a partition available
/usr/lib/vxvm/bin/vxmksdpart: 3pardata0_2492: is not an identifier
```

Diskgroup import of BCV luns using -o updateid and -o useclonedev options is not supported if the diskgroup has mirrored volumes with DCO or has snapshots. (2831658)

VxVM uses guid stored in configuration to uniquely identify all objects. The DCO volume stores the guid of mirrors and snapshots. If the diskgroup is imported with -o updateid and -o useclonedev, it changes the guid of objects in VxVM configuration database and the guids stored in DCO volume are not updated. So the operations involving DCO will not be able to find objects with the stored guid and this could lead to failure of certain operations involving DCO or could lead to unexpected behaviour.

Workaround:

No workaround available.

Disks on the Oracle VM Server for SPARC guest are claimed under other_disks category (2354005)

The disks on the Oracle VM Server for SPARC guest are claimed under "other_disks" enclosure, because these disks are not capable of being multi-pathed by DMP. This is expected because these devices represent VxVM volumes in the host. By design, devices under other_disks enclosure have their name based on underlying OS path regardless of the DDL naming scheme.

Upgrading from Veritas Storage Foundation and High Availability Solutions 5.x to 6.0.1 may fail for IBM XIV Series arrays (2715119)

Starting in the Veritas Storage Foundation and High Availability Solutions 5.1 SP1 release, the Array Support Library (ASL) for the IBM XIV enclosures converts the LUN Serial Number from Hexadecimal to Decimal. Because of this change, the enclosure names differ from releases prior to the 5.1 SP1 releases. When you upgrade Veritas Storage Foundation and High Availability Solutions from a release prior to that release to the current 6.0.1 release, XIV LUNs may go into an error state. Note that the latest RPs on 5.1/5.1SP1 are already modified to use the same logic for enclosure naming.

Workaround:

After the upgrade, run `vxddladm assign names`.

Cannot grow Veritas Volume Manager (VxVM) disk using the vxdisk resize command during Dynamic LUN Expansion operation (2064510)

The following error message is displayed during the Dynamic LUN Expansion operation of a LUN with the SIMPLE format:

```
VxVM vxdisk ERROR V-5-1-8643 Device <device name>: resize failed:
Invalid data in request
```

The `vxdisk resize` command keeps the cylinder size (number of the heads * total number of the sectors per track) constant before and after the resize operation, unless the number of cylinders go beyond 2^16-1 (65535) . Because of the VTOC limitation of storing geometry values only till 2^16 -1, if the number of cylinders increases beyond the limit, `vxdisk resize` increases the cylinder size. If this happens, the private region will overlap with the public region data and corrupt the user data.

As a result of this LUN geometry change, VxVM is unable to complete `vxdisk resize` on simple format disks. VxVM was not designed to handle such geometry changes during Dynamic LUN Expansion operations on simple disks.

Workaround:

The VxVM `vxdisk resize` command behaves differently depending on whether the disk is simple, sliced, or CDS format.

The problem shown above only occurs on simple disk configurations. As a result of this difference in behavior, if the geometry changes during a Dynamic LUN Expansion operation at the LUN level, you can convert the disk to a CDS format disk. Use the `vxcdsconvert` command on the disk. Then you can issue the `vxdisk resize` command.

See http://www.symantec.com/docs/TECH136240 for more information.

1 TB luns goes in error state with Solaris x86 (2706776)

If you label a disk device as EFI using format on a subset of the paths or on the DMP device, Solaris will not be able to propagate the label to all the other paths of the LUN. This will lead the device to appear in the error state under 'vxdisk list'.

Workaround:

There is no workaround for this issue.

After devices that are managed by EMC PowerPath lose access to storage, Veritas Volume Manager commands are delayed (2757198)

In an enviroment which includes devices that are managed by EMC PowerPath, a storage loss causes Veritas Volume Manager commands to be delayed. In the event of storage loss, VxVM sends SCSI inquiry from each LUN path to check the health of path, which are delayed by the presence of EMC PowerPath.

Performance impact when a large number of disks are reconnected (2802698)

If the storage connectivity is lost to part of the storage, the disk group configuration copy is rebalanced to the disks that have connectivity. For example, if the storage for an entire enclosure is removed from a disk group with muliple enclosures. The rebalancing process takes time, during which time the `vxconfigd` daemon is busy and does not respond to commands.

Plex synchronization is not completed after resuming synchronization on a new master when the original master lost connectivity (2788077)

When you run `vxrecover -o force`, it recovers only one subvolume and it cannot detect that the rest of the volume needs recovery.

When you run the `vxassist mirror` command, you run the `vxplex att` command serially on each subvolume. If the failure happens before you start the `attach` operation (need to mark the concerned plex as the attach operation is in progress), `vxrecover` will not redo the attach operation because it cannot find any record of the attach operation in progress.

Workaround:

Run the following command on each subvolume to manually recover the complete volume:

```
# /usr/lib/vxvm/type/fsgen/vxplex -U fsgen -g diskgroup \
 -o force useopt att volume plex
```

CVMVolDg agent may fail to deport CVM disk group

The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the `CVMDeportOnOffline` attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline.

If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.

Workaround: If multiple CVMVolDg resources are configured for a shared disk group, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

vxlustart failed due to lumount error when liveupgrade to Solaris 10 U11 (3035982)

Live Upgrade (LU) to Solaris 10 U11 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:

Use the Solaris 10 U10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 U11 LU packages. Then use `vxlustart` to upgrade to Solaris 10 U11.

The vxdisksetup command fails to initialize disks in cdsdisk format for Oracle VM Server for SPARC disks greater than 1 TB (2557072)

The `vxdisksetup` command fails to initialize disks in `cdsdisk` format for Oracle VM Server for SPARC disks greater than 1 TB. This issue is due to an Oracle VM Server for SPARC operating system command which fails when the number of partitions in the GUID partition table (GPT) label is greater than 9. The `cdsdisk` format requires at least 128 partitions to be compatible with Linux systems.

Workaround: There is no workaround for this issue.

Veritas Cluster Server known issues

This section describes the known issues in this release of Veritas Cluster Server (VCS).

- Operational issues for VCS
- Issues related to the VCS engine
- Issues related to the bundled agents

Operational issues for VCS

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

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

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

Workaround:

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

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 Oracle VM Server for SPARC 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 Oracle VM Server for SPARC 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.

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
```

ClearClone attribute is not honored by DiskGroup resources on Solaris 11 systems when publisher is not set (3484821)

On Solaris 11 systems, if the `ClearClone` attribute for the DiskGroup resource is set to 1 and no publisher is set for the system, the DiskGroup agent fails to correctly identify the VRTSvxvm package version and resets the ClearClone attribute to 0. Thus the disk group will not be imported with the `-c` option.

Workaround: Use the below command to check the publisher information:

```
# /usr/bin/pkg publisher
```

If the output does not contains any publisher information, set a solaris publisher or any other available publisher for the system:

```
# /usr/bin/pkg set-publisher -g \
http://pkg.oracle.com/solaris/release/ solaris
```

Issues related to the VCS engine

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.

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.

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.

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

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

Workaround

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

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

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

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

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

or

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

VRTSvcs package may give error messages for package verification on Solaris 11 [2858192]

VRTSvcs package may give error messages for package verification on Solaris 11. This is because some of the VCS configuration files are modified as part of product configuration. This error can be ignored.

Workaround: No workaround.

GCO clusters remain in INIT state [2848006]

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

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

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

Verification with `hacf -verify` shows no errors when ContainerInfo and ResContainerInfo are set with service groups having multiple zones [2859495]

When defining ResContainerInfo for any resource in group in main.cf, if you also define ContainerInfo the corresponding group in main.cf then ContainerInfo is used instead of ResContainerInfo when main.cf is loaded. Verification of the configuration file using `hacf -verify` at this stage shows no errors.

Workaround: Make sure ContainerInfo is not defined while defining ResContainerInfo for corresponding group in main.cf .

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.

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.

types.cf error after rollback 6.0.5 to 6.0.1(3067510)

On Solaris 10, if you rollback Storage Foundation and High Availability Solutions 6.0.5 to 6.0.1, the DomainFailurePolicy attribute, which is a new attribute in 6.0.5, cannot be automatically removed in the VCS configuration files `types.cf` and `main.cf`. This causes `types.cf` verification to fail and `had` not to start after the uninstallation because the DomainFailurePolicy attribute is not identified by SFHA 6.0.1.

Workaround:

1. Stop the cluster.
2. Remove all the occurrences of the DomainFailurePolicy attribute in the `types.cf` and `main.cf` files at `/etc/VRTSvcs/conf/config`.
3. Start the cluster from the node where the configuration files are modified.
4. Start VCS on all other nodes.

If secure and non-secure WAC are connected the engine_A.log receives logs every 5 seconds [2653695]

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

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

The hazonesetup utility can not recognize comma in arguments [3389647]

The hazonesetup utility uses comma as a separator for the arguments of the same type. If a comma is part of an argument, the argument is recognized as two different

parameters. Therefore, the hazonesetup utility fails to recognize comma in service group name, resource name, zone name, user name, password, and system list.

Workaround: Do not include comma in service group name, resource name, zone name, user name, password, and system list.

Issues related to the bundled agents

Solaris mount agent fails to mount Linux NFS exported directory (2100188)

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 = "-o vers=3"
               )
```

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.

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

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

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

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

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

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

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.

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.

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 [2938108]

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 Oracle VM Server for SPARC resource may fail due to incompatibility of Oracle VM Server for SPARC configuration file with host OVM version (2814991)

If you have a cluster running Oracle VM Server for SPARC with different OVM versions on the hosts, then the Oracle VM Server for SPARC 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 Oracle VM Server for SPARC resource may also fail.

For example, if you have a cluster running Oracle VM Server for SPARC 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 Oracle VM Server for SPARC 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 SFHA 6.0 or later versions.

When you upgrade SFHA 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 SFHA 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 SFHA 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 SFHA 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
```

Mounting a / file system of NFS server is not supported [2847999]

Mount agent do not support BlockDevice attribute with / file system of NFS server for NFS file system.

Workaround: No workaround.

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

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

Workaround: No workaround.

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
```

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 fails to come online if "-o" is not provided [3401616]

In the Options attribute for Share resource, Share resource fails to come online if "-o" is not provided.

Workaround: Provide "-o" preceding the options in the Options attribute.

Network File System (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 resources, are still online on the failing node. The disk group on the failing node may also get disabled, but IP resources on the same node stays online. As the result, I/O error occurs.

Workaround:

Configure the pre-online trigger for the service groups containing DiskGroup resources 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/`*`filename`* as T0preonline_ipc:

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

2 Enable T0preonline_ipc, the pre-online trigger for the service group:

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

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.

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

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

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

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

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

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.

Issues related to the agent framework

Issues with configuration of resource values (1718043)

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

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

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

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

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.

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.

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: No workaround.

Delayed response to VCS commands observed on nodes with several resources and system has high CPU usage or high swap usage [3432749]

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.

Issues related to VCS in Japanese locales

This section covers the issues that apply to SFHA 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.

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.

Second secondary cluster cannot take over the primary role when primary and 1st-secondary clusters panic [2858187]

If there are three clusters(clus1, clus2, and clus3) in a GCO without a steward, when clus1 loses connection to clus2, it will send the inquiry to clus3 to check the state of clus2:

- If it is able to confirm that clus2 is down, it will mark clus2 as FAULTED.
- If it is not able to send the inquiry to clus3, it will assume that a network disconnect has happened and mark clus2 as UNKNOWN. In this case, automatic failover will not take place even if the ClusterFailoverPolicy is set to Auto. If this happens, users would need to manually failover the global service groups.

Workaround:

Configure the steward at a location geographically distinct from those of the three clusters above.

LLT known issues

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

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.

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.

GAB known issues

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

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.

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

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

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly 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.

Fail to stop GAB when upgrading the second sub cluster (3066737)

On Solaris 11, in the phased upgrade scenario which includes the Veritas Cluster Server components, the GAB kernel driver may fail to unload on some nodes in some cases. The script based installer will display the following message in such situation:

```
gab failed to stop on sys1
It is recommended to reboot the systems sys1 to resolve
the failures and then retry. If issues persist after reboot,
contact Symantec technical support or refer to
installation guide for further troubleshooting.
Do you want to continue? [y,n,q] (n) y
```

Workaround:

Continue to upgrade the second sub cluster, then execute `/usr/sbin/shutdown -y -i6 -g0` to restart the nodes when the whole upgrade is completed.

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
```

Installer is unable to split a cluster that is registered with one or more CP servers (2110148)

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

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

However, if you use the installer to reconfigure the Veritas Storage Foundation HA, the installer retains the same cluster UUID of *clus1* in both *clus1A* and *clus1B*. If both *clus1A* and *clus1B* use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: 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.

Stale .vxfendargs file lets hashadow restart vxfend in Sybase mode (2554886)

When I/O fencing is configured in customized mode, vxfend, the user mode daemon of I/O fencing, creates the `/opt/VRTSvcs/lock/.vxfendargs` file. VCS uses this file to restart the vxfend daemon when it gets killed. However, VCS does not use this file when I/O fencing is configured in Sybase mode. This file is not removed from the system when I/O fencing is unconfigured.

If user configures I/O fencing in Sybase mode and an old `/opt/VRTSvcs/lock/.vxfendargs` file is present in the system from an earlier configuration of I/O fencing in customized mode, then VCS attempts to restart the

vxfend daemon every time it is killed. This interferes with the functioning of I/O fencing in the Sybase mode.

Workaround: Before you configure I/O fencing in Sybase mode, delete the `/opt/VRTSvcs/lock/.vxfendargs` file if it is present in the system.

After upgrading coordination point server in secure mode the cpsadm command may fail with errror - Bus error (core dumped) (2846727)

After upgrading the coordination point server from SFHA 5.0 to the next version on the client system, if you do not remove the VRTSat package that were installed on the system, the cpsadm command fails. The command fails because it loads old security libraries present on the system. The cpsadm command is also run on the coordination point server to add or upgrade client clusters. The command also fails on the server because it loads old security libraries present on the system.

Workaround: Perform the following steps on all the nodes on the coordination point server:

1 Rename cpsadm to cpsadmbin

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

2 Create the `/opt/VRTScps/bin/cpsadm` file with the following details.

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

3 Give executable permissions to the new file.

```
# chmod 775 /opt/VRTScps/bin/cpsadm
```

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.

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.

CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

See the *Veritas Storage Foundation and High Availability 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.

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.

NIC resource gets created with incorrect name while configuring CPSSG with the configure_cps.pl script (2585229)

The name of the NIC resource created by the `configure_cps.pl` script does not come out correct when, for example, m^{th} VIP is mapped to n^{th} NIC and every m is

not equal to n. In this case, although CPSSG continues to function without any problem, when you unconfigure CPSSG using `configure_cps.pl`, it fails.

Workaround: To unconfigure CPSSG, you must remove the CPSSG configuration from the VCS configuration.

CP server configuration fails while setting up secure credentials for CP server hosted on an SFHA cluster (2621029)

When you configure CP server using the `configure_cps.pl` utility, the configuration fails while setting up secure credentials for CP server that is hosted on an SFHA cluster. You may see the following error:

```
Creating softlink to credential directory /etc/VRTScps/db/CPSERVER
on node nodename.
Unable to connect to node nodename using /usr/bin/ssh.
Please configure ssh communication and retry. Exiting.
```

Workaround: You can use any of the following options:

- Before running the `configure_cps.pl` utility, change the default shell for root user to either KSH or bash.
- Perform the following steps after running the `configure_cps.pl` utility on each node of the cluster:
  - Manually remove the old credential directory or softlink. For example:

    ```
    # rm -rf /var/VRTSvcs/vcsauth/data/CPSERVER
    ```

  - Create a new soft-link to the shared location of the credential directory:

    ```
    # ln -s path_of_CP_server_credential_directory \
     /var/VRTSvcs/vcsauth/data/CPSERVER
    ```

  - Start the CPSSG service group:

    ```
    # hagrp -online CPSSG -any
    ```

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.

Cannot run the vxfentsthdw utility directly from the install media if VRTSvxfen package is not installed on the system (2858190)

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

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

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

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

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

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

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

Workaround:

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

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.

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.

VRTSvxfen fails to install in local zones [3407325]

After VCS 6.0.3 upgrades to VCS 6.0.5 on Solaris 10, if a non-global zone is created or attached with the `-U` option, the following error message is displayed:

```
Installation of these packages generated errors: VRTSvxfen
```

Workaround: There is no workaround for this issue. Since the VRTSvxfen package only installs metadata inside non-global zones, it is safe to ignore the error.

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

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

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

Issues related to Intelligent Monitoring Framework (IMF)

Registration error while creating a Firedrill setup [2564350]

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

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

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

Workaround: No workaround.

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.

Engine log gets flooded with messages proportionate to the number of mount offline registration with AMF [2619778]

In a certain error condition, all mount offline events registered with AMF are notified simultaneously. This causes the following message to get printed in the engine log for each registered mount offline event:

```
<Date> <Time> VCS INFO V-16-2-13717
(vcsnode001) Output of the completed operation
(imf_getnotification)
=============================================
Cannot continue monitoring event
Got notification for group: cfsmount221

=============================================
```

This is an expected behavior for this error condition. Apart from the messages there will be no impact on the functionality of the VCS solution.

Workaround: No workaround.

Perl errors seen while using haimfconfig command

Perl errors seen while using `haimfconfig` command:

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

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

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

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

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

It should be replaced as follows in main.cf:

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

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

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

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

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

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

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

Error messages displayed during reboot cycles [2847950]

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

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

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

Workaround: No workaround.

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

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

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

Workaround: No Workaround.

Error message seen during system shutdown [2954309]

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

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

The system continues to proceed with the shutdown.

Workaround: No workaround.

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.

Engine log shows Asynchronous Monitoring Framework (AMF) error message on using the `cfsshare` command [3235274]

When you use the cfsshare command, the IMF plugin may report the following error in the engine log:

```
AMF libvxamf ERROR V-292-2-139 This event is already registered
```

Workaround: This message is harmless and can be ignored. There is no workaround of the issue.

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.

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. The system displays the following error when you attempt to use the Java GUI:

```
Incorrect username/password
```

Workaround: No workaround.

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.

Issues related to ACCLIB

WebLogic agent fails to bring the resource ONLINE or OFFLINE [3383446]

The VRTSacclib packaged with VCS 6.0.1 has an incorrect check in the `Proc.pm` file, which makes the WebLogic agent unable to bring the resource ONLINE or OFFLINE. You can find the following error message in the engine log:

```
Proc:GetLongProcessListHash:EnvFlag [1] not supported on OS [SOLARIS]
```

Workaround:

Symantec recommends the customer to use the VRTSacclib that is available on SORT at https://sort.symantec.com/agents instead of the VRTSacclib that is packaged with VCS 6.0.1.

Veritas Dynamic Multi-pathing known issues

This section describes the known issues in this release of Veritas Dynamic Multi-pathing.

Creating a zpool fails with a incorrect disk size error (2277875)

When the tunable parameter dmp_native_support is turned on, creating a zpool on DMP devices may fail with the following error:

```
one or more devices is less than the minimum size (64 M)
```

This error may occur even if the device size is greater than the required minimum size.

Workaround:

To resolve this issue, use one of the following commands:

- `# vxdisk scandisks`
- `# format -e dmp_device`

DMP aggregates EFI labelled LUNS to a 0_0 disk (2558408)

While performing `vxdiskunsetup` of some luns, if you format and label the disks as EFI, all the EFI labelled luns are aggregated to a 0_0 disk.

Workaround:

When changing the label of a disk from SMI to EFI, or vice-versa, Symantec recommends that the label be changed on all accessible paths to a disk. That is, use the `format -e` command to stamp the new label on all accessible paths. For

Active/Passive (A/P) class of arrays, this should be done only on the active paths. For other arrays, all paths should be labeled.

Symantec also recommends the installation of the patch provided by Oracle for EFI label issues (IDR144101-01 or IDR144249-01 or release kernel patch 142909-17). If this patch is installed, you can run the `format -e` command only on one path. After that, perform a read operation (such as `dd if=/dev/rdsk/<path> of=/dev/null count=1`) on the other accessible paths to propagate the label.

Splitting a mirror from a zpool causes a core dump (2273367)

The following operation to split a mirror from a zpool fails:

```
# zpool split my_pool new_pool mirror
```

This issue is an Oracle issue with zpool. This issue occurs whether DMP is controlling the devices or not. That is, whether the dmp_native_support tunable is on or off.

Suppressing the primary path of an encapsulated SAN boot disk from Veritas Volume Manager causes the system reboot to fail (1933631)

If you suppress the primary path of an array from VxVM control and then reboot the system, the system boot fails.

If you have an encapsulated SAN boot device with multiple primary paths, the issue occurs when you suppress the first primary path. When you configure a SAN boot device, the primary path is set as a boot device. In general, the first path of the SAN boot device corresponds to the first configured path during SAN boot. Even if another primary path is configured as a boot device, suppressing the first device from VxVM causes the boot to fail.

Workaround:

When the boot device is suppressed from VxVM, change the OS boot device sequencing accordingly.

For Solaris SPARC system, use the `eeprom` *`boot-device`* command to set the boot device sequencing.

For Solaris x86-64 systems, use the `eeprom` *`bootpath`* command to set the boot device sequencing.

Adding a DMP device or its OS device path as a foreign disk is not supported (2062230)

When DMP native support is enable, adding a DMP device or its OS device path as a foreign disk using the `vxddladm addforeign` command is not supported. Using this command can lead to unexplained behavior.

ZFS pool creation on a DMP device fails when the LUN size is between 1 TB and 2TB (2010919)

Creating a ZFS pool on a DMP device using the whole disk of size > 1TB and < 2TB that contains a SMI SUN label fails. The issue is that zpool create on a whole disk changes the device label from SMI to EFI. This causes confusion between the OS device paths of the same DMP device due to a bug in the Oracle SCSI layer. This is due to Oracle BugID: 6912703.

After changing the preferred path from the array side, the secondary path becomes active (2490012)

For EVA arrays, DMP requires that the prefer bit is static. If the prefer bit is not static, issues like the following may occur. After changing the prefer path of LUN from the array side, and performing a disk discovery (`vxdisk scandisks`) from the host, the secondary path becomes active for the LUN.

Workaround:

To work around this issue

1. Set the pref bit for the LUN.
2. Perform disk discovery again:

```
# vxdisk scandisks
```

Continuous trespass loop when a CLARiiON LUN is mapped to a different host than its snapshot (2761567)

If a CLARiiON LUN is mapped to a different host than its snapshot, a trespass on one of them could cause a trespass on the other. This behavior could result in a loop for these LUNs, as DMP tries to fail back the LUNs if the primary paths are available.

Workaround:

To avoid this issue, turn off the `dmp_monitor_ownership` tunable:

```
# vxdmpadm settune dmp_monitor_ownership=off
```

After excluding devices managed by PowerPath from VxVM, the devices still show as DMP devices (2494632)

The issue happens after EMC PowerPath is installed and all devices are under PowerPath control. If you want to maintain the devices under PowerPath control, you use the following command to exclude the device that is managed by PowerPath from VxVM:

```
# vxdmpadm exclude dmpnodename=PowerPath_device_name
```

After system reboot, the PowerPath device still shows as a DMP device, although the device is managed by EMC PowerPath.

Workaround:

This issue is seen only during the first bootup discovery after reboot. To resolve the issue, manually trigger DMP device discovery:

```
# vxdisk scandisks
```

The system may hang with Solaris 11 SRU1 (2876211)

When running Solaris 11 SRU1, the system may hang due to an Oracle bug. The Oracle Bug ID is 7105131 deadman panic.

Workaround: SRU1 for Solaris 11 should be updated to SRU2a. The bug is fixed in SRU2a: Oracle Solaris 11 Support Repository Updates (SRU) Index (Doc ID 1372094.1)

System panics because dmp_signal_event() called psignal() with incorrect vxesd proc pointer (3041167)

On Solaris 11 SPARC SRU1, system panics after executing the `vxrecover` operation on the master node.

Workaround:

There is no workaround for this issue in 6.0.5.

Changes in enclosure attributes are not persistent after an upgrade to VxVM 6.0.5 (2082414)

The Veritas Volume Manager (VxVM) 6.0.5 includes several array names that differ from the array names in releases prior to release 5.1SP1. Therefore, if you upgrade from a previous release to VxVM 6.0.5, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0.5. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-31 shows the Hitachi arrays that have new array names.

Table 1-31 Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC
TagmaStoreUSPV	Hitachi_USP-V
TagmaStoreUSPVM	Hitachi_USP-VM
<New Addition>	Hitachi_R700
Hitachi AMS2300 Series arrays	New array names are based on the Model Number 8x. For example, AMS_100, AMS_2100, AMS_2300, AMS_2500, etc.

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0.5. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

Veritas Storage Foundation known issues

This section describes the known issues in this release of Veritas Storage Foundation.

- Veritas Storage Foundation and High Availability Solutions known issues
- Veritas File System known issues
- Replication known issues
- Veritas Storage Foundation for Databases (SFDB) tools known issues

Veritas Storage Foundation and High Availability Solutions known issues

This section describes the known issues in this release of Veritas Storage Foundation and High Availability Solutions (SFHA).

In an IPv6 environment, db2icrt and db2idrop commands return a segmentation fault error during instance creation and instance removal (1602444)

When using IBM DB2 `db2icrt` command to create a DB2 database instance on a pure IPv6 environment, the `db2icrt` command returns segmentation fault error message. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2icrt -a server -u db2fen1 db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 4700: 26182 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -addfcm -i ${INSTNAME?}
```

The `db2idrop` command also returns segmentation fault, but the instance is removed successfully after the `db2idrop` command is issued. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2idrop db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 3599:  7350 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -remove -s DB2_${INSTNAME?} 2> /dev/null

DBI1070I  Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

This issue has been identified as an IBM issue. Once IBM has fixed this issue, then IBM will provide a hotfix for this segmentation problem.

At this time, you can communicate in a dual-stack to avoid the segmentation fault error message until IBM provides a hotfix.

To communicate in a dual-stack environment

- Add an IPv6 hostname as an IPv4 loopback address to the `/etc/hosts` file. For example:

  ```
  127.0.0.1 swlx20-v6
  ```

  Or

  ```
  127.0.0.1 swlx20-v6.punipv6.com
  ```

  `127.0.0.1` is the IPv4 loopback address.

  `swlx20-v6` and `swlx20-v6.punipv6.com` are the IPv6 hostnames.

Boot fails after installing or removing SFHA packages from a Solaris 9 system to a remote Solaris 10 system (1747640)

The following issue occurs if you install or remove a Storage Foundation package or patch from a Sparc Solaris 9 system to a remote Solaris 10 system, using the `-R rootpath` option of the `pkgadd`, `patchadd`, `pkgrm` or `patchrm` commands.

Generally, when you install or remove a SFHA package on a Solaris 10 system, the package scripts update the boot archive. However if the local system is Solaris 9 and the remote system is Solaris 10, the scripts fail to update the boot archive on the Solaris 10 system.

Note: The boot archive is synchronized correctly when you upgrade SFHA using Solaris Live Upgrade.

Workaround:

The workaround is to manually clear the boot archive when you boot the alternate. The Oracle boot process detects that the boot archive is out sync and displays instructions for how to correct the situation.

For example:

```
WARNING: The following files in / differ from the boot archive:

     stale //kernel/drv/sparcv9/vxportal
     stale //kernel/drv/vxportal.conf
     stale //kernel/fs/sparcv9/vxfs
     ...
     new     /kernel/drv/vxlo.SunOS_5.10
     new     /kernel/drv/vxlo.conf
     changed /kernel/drv/vxspec.SunOS_5.9
     changed /kernel/drv/vxspec.conf

 The recommended action is to reboot to the failsafe archive to correct
 the above inconsistency. To accomplish this, on a GRUB-based platform,
 reboot and select the "Solaris failsafe" option from the boot menu.
 On an OBP-based platform, reboot then type "boot -F failsafe". Then
 follow the prompts to update the boot archive. Alternately, to continue
 booting at your own risk, you may clear the service by running:
 "svcadm clear system/boot-archive"
```

Oracle 11gR1 may not work on pure IPv6 environment [1819585]

There is problem running Oracle 11gR1 on a pure IPv6 environment.

Tools like `dbca` may hang during database creation.

Workaround: There is no workaround for this, as Oracle 11gR1 does not fully support pure IPv6 environment.

Sybase ASE version 15.0.3 causes segmentation fault on some Solaris version (1819595)

Sybase ASE 15.0.3 produces segmentation fault on Solaris SPARC 10 Update 6 in a pure IPv6 environment. However, Sybase ASE 15.0.3 works on Solaris SPARC 10 Update 5.

When running Sybase ASE 15.0.3 GA on a pure IPv6 environment on Solaris SPARC 10 Update 6, you may receive a segmentation fault message. For example:

```
Building Adaptive Server 'CDGV240AIPV6':
Writing entry into directory services...
Directory services entry complete.
Building master device...
Segmentation Fault - core dumped
Task failed
Server 'CDGV240AIPV6' was not created.
```

This is a Sybase known issue. You should use Sybase Adaptive Server Enterprise Suite version 15.0.3 ESD 1 that supports Solaris 10 Update 6 or later. For details, refer to the Sybase Product Download Center regarding ESD 1.

Not all the objects are visible in the VOM GUI (1821803)

After upgrading SF stack from 5.0MP3RP2 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the VOM GUI.

Workaround:

To resolve this known issue

- On each manage host where `VRTSsfmh` 2.1 is installed, run:

  ```
  # /opt/VRTSsfmh/adm/dclisetup.sh -U
  ```

DB2 databases are not visible from the VOM Web console (1850100)

If you upgraded to SF 5.1, DB2 databases will be not visible from the VOM web console.

This will be fixed in the SF 5.1 Patch 1 release.

Workaround: Reinstall is required for VOM DB2-Hotfix (`HF020008500-06.sfa`), if the host is upgraded to SF 5.1. Use the deployment framework and reinstall the hotfix for DB2 (`HF020008500-06.sfa`) on the managed host.

To resolve this issue

1. In the Web GUI, go to **Settings** > **Deployment**.
2. Select **HF020008500-06 hotfix**.
3. Click **Install**.
4. Check the **force** option while reinstalling the hotfix.

A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a SmartTier placement policy (1880081)

A volume's placement class tags are not visible in the Veritas Enterprise Administrator (VEA) GUI when you are creating a SmartTier placement policy if you do not tag the volume with the placement classes prior to constructing a volume set for the volume.

Workaround:

To see the placement class tags in the VEA GUI, you must tag the volumes prior to constructing the volume set. If you already constructed the volume set before tagging the volumes, restart `vxsvc` to make the tags visible in the GUI.

NULL pointer dereference panic with Solaris 10 Update 10 on x86 and Hitachi Data Systems storage (2616044)

Due to a limitation with Solaris 10 Update 10 on x86, when the server is connected to Hitachi Data storage, the system panics due to a NULL pointer deference during the boot cycle with the following stack trace:

```
fffffe8000988570 unix:die+da ()
fffffe8000988650 unix:trap+5e6 ()
fffffe8000988660 unix:cmntrap+140 ()
fffffe8000988870 scsi_vhci:hds_sym_path_get_opinfo+62 ()
fffffe8000988920 scsi_vhci:vhci_update_pathinfo+5b ()
fffffe80009889a0 scsi_vhci:vhci_pathinfo_online+2df ()
fffffe8000988a10 scsi_vhci:vhci_pathinfo_state_change+202 ()
fffffe8000988a70 genunix:i_mdi_pi_state_change+148 ()
fffffe8000988ab0 genunix:mdi_pi_online+32 ()
fffffe8000988b20 fcp:ssfcp_online_child+ff ()
fffffe8000988b90 fcp:ssfcp_trigger_lun+2b0 ()
fffffe8000988bc0 fcp:ssfcp_hp_task+88 ()
fffffe8000988c40 genunix:taskq_thread+295 ()
fffffe8000988c50 unix:thread_start+8 ()
```

For more information, see Oracle bug ID 7079724.

Workaround:

Disable Solaris I/O multi-pathing on the server to avoid the system panic.

To disable Solaris I/O multi-pathing on the server

1 Disable Solaris I/O multi-pathing:

```
# stmsboot -d
```

2 Reboot the server:

```
# reboot
```

Veritas File System known issues

This section describes the known issues in this release of Veritas File System (VxFS).

The multi-user services fail to run on a non-global zone when zone is created after patch is applied [3115066]

The multi-user services fail to run on a non-global zone in case the zone is created after patch is applied.

Workaround: Use the `svcadm enable vxfsldlic` command to enable the service inside non-global zone.

Taking a FileSnap over NFS multiple times with the same target name can result in the 'File exists' error (2353352)

The "File exists" error occurs as a result of the caching behavior of the NFS client. Because the link operation is successful, the NFS client assumes that a file with the specified target name, such as `file2::snap:vxfs:`, was created.. As a result, the NFS client caches a file with this name.

Workaround: Remove the target file after a snapshot is created. This forces the NFS client to remove the name from the cache. For example:

```
# ln file1 file2::snap:vxfs:
# rm file2::snap:vxfs:
```

Enabling delayed allocation on a small file system sometimes disables the file system (2389318)

When you enable delayed allocation on a small file system, such as around 100 MB, the file system can get disabled. In this case, the following error message ,displays in the system console log:

```
mesg 001: V-2-1: vx_nospace - file_system file system full
(size block extent)
```

Workaround:

Use the `vxtunefs` command to turn off delayed allocation for the file system.

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system nears 100% usage even if other volumes have free space (2438368)

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system is nearing 100% usage even if other volumes in the file system have free space.

Workaround:

After sufficient space is freed from the volume, delayed allocation automatically resumes.

Deduplication can fail with error 110 (2591473)

In some cases, data deduplication fails with a message similar to the following example:

```
Saving    Status    Node            Type        Filesystem
-------------------------------------------------------------------------
00%       FAILED    node01          MANUAL      /data/fs1
        2011/10/26 01:38:58 End full scan with error
```

In addition, the deduplication log contains an error similar to the following example:

```
2011/10/26 01:35:09 DEDUP_ERROR AddBlock failed. Error = 110
```

These errors indicate that the deduplication process is running low on space and needs more free space to complete.

Workaround:

Make more space available on the file system.

vxresize fails while shrinking a file system with the "blocks are currently in use" error (2437138)

The `vxresize` shrink operation may fail when active I/Os are in progress on the file system and the file system is being shrunk to a size closer to its current usage. You see a message similar to the following example:

```
UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink /dev/vx/rdsk/dg1/vol1 -
blocks are currently in use.
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume
vol1, in diskgroup dg1
```

Workaround:

Rerun the shrink operation after stopping the I/Os.

Warning message sometimes appear in the console during system startup (2354829)

During system startup, following messages sometimes appear in system console:

```
WARNING: couldn't allocate SDT table for module vxfs
WARNING: couldn't allocate FBT table for module vxfs
Loading smf(5) service descriptions: 2/2
```

These warnings indicate that the SDT and FBT DTrace probes might not be available for the VxFS module. The VxFS module still loads and works correctly. Dtrace SDT/FBT has limits on the size of module that it can support. Since the VxFS module exceeds the size that Dtrace can support, SDT and FBT Dtrace probes might not work for VxFS.

Workaround:

There is no workaround for this issue.

Possible assertion failure in vx_freeze_block_threads_all() (2244932)

There is a possible assertion failure in the `vx_freeze_block_threads_all`() call when the `pdir_threshold` tunable is set to 1.

Workaround:

There is no workaround for this issue.

Severe impact in read performance (sequential and random) on compressed files compared to uncompressed files (2609152)

The read throughput is highly degraded for compressed files. The difference is seen for sequential I/O and random I/O. For sequential reads, the degrdataion is visbile even when the amount of data read compressed files is one-third of the uncompressed files (compression ratio).

Workaround:

There is no workaround for this issue.

fsppadm operations issued on multi-volume file system fail if there are other mounted file systems with a disk layout Version less than 6 (2909206, 2909203)

The `fsppadm` command checks all mounted file systems, and if it finds any file systems with a disk layout Version that is less than 6, then it exits with the following error message:

```
# fsppadm assign /dst_vset /tmp/pol_test.xml

UX:vxfs fsppadm: ERROR: V-3-26510: Low level Volume enumeration failure
on / with message  Function not implemented
```

This error occurs because the `fsppadm` command functionality is not supported on a disk layout Version that is less than 6.

Workaround:

There is no workaround for this issue.

Internal tests hit assert "f:vx_putpage1:1a" on Solaris11 U1 SRU2.5 (3060829)

The assert failure will not cause any outage but may lead to performance issues on systems with large memory and swap file configurations. This is due to a known defect on Solaris 11 update 1, refer to the Bug ID below for more details on the Oracle support site:

Bug ID: 15813035

SUNBT7194962 pageout no longer pushes pages asynchronously.

Workaround:

There is no workaround for this issue.

Replication known issues

This section describes the replication known issues in this release of Veritas Storage Foundation and High Availability Solutions.

vradmin syncvol command compatibility with IPv6 addresses (2075307)

The `vradmin syncvol` command does not work with the compressed form of IPv6 addresses if the target disk group and volume names are not specified.

Workaround:

In IPv6 environments, if you run the `vradmin syncvol` command and identify the target host using the compressed form of the IPv6 address, then you also need to specify the target disk group and volume names.

RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2054804)

The RVGPrimary agent initiated operation to start replication between the original Primary and the bunker fails during failback – when migrating back to the original Primary after disaster recovery – with the error message:

```
VxVM VVR vxrlink ERROR V-5-1-5282 Error getting information from
remote host. Internal Error.
```

The issue applies to global clustering with a bunker configuration, where the bunker replication is configured using storage protocol. It occurs when the Primary comes back even before the bunker disk group is imported on the bunker host to initialize the bunker replay by the RVGPrimary agent in the Secondary cluster.

Workaround:

To resolve this issue

1. Before failback, make sure that bunker replay is either completed or aborted.
2. After failback, deport and import the bunker disk group on the original Primary.
3. Try the start replication operation from outside of VCS control.

Bunker replay did not occur when the Application Service Group was configured on some of the systems in the Primary cluster, and ClusterFailoverPolicy is set to "AUTO" (2047724)

The time that it takes for a global cluster to fail over an application service group can sometimes be smaller than the time that it takes for VVR to detect the configuration change associated with the primary fault. This can occur in a bunkered, globally clustered configuration when the value of the `ClusterFailoverPolicy` attribute is `Auto` and the `AppGroup` is configured on a subset of nodes of the primary cluster.

This causes the RVGPrimary online at the failover site to fail. The following messages appear in the VCS engine log:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname could not be
imported on bunker host hostname. Operation failed with error 256
and message VxVM VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote
server unreachable... Timestamp VCS ERROR V-16-2-13066 (hostname)
Agent is calling clean for resource(RVGPrimary) because the resource
is not up even after online completed.
```

Workaround:

To resolve this issue

- When the configuration includes a bunker node, set the value of the `OnlineRetryLimit` attribute of the RVGPrimary resource to a non-zero value.

The RVGPrimary agent may fail to bring the application service group online on the new Primary site because of a previous primary-elect operation not being run or not completing successfully (2043831)

In a primary-elect configuration, the RVGPrimary agent may fail to bring the application service groups online on the new Primary site, due to the existence of previously-created instant snapshots. This may happen if you do not run the `ElectPrimary` command to elect the new Primary or if the previous `ElectPrimary` command did not complete successfully.

Workaround:

Destroy the instant snapshots manually using the `vxrvg -g` *`dg`* `-P` *`snap_prefix`* `snapdestroy` *`rvg`* command. Clear the application service group and bring it back online manually.

A snapshot volume created on the Secondary, containing a VxFS file system may not mount in read-write mode and performing a read-write mount of the VxFS file systems on the new Primary after a global clustering site failover may fail (1558257)

Issue 1:

When the `vradmin ibc` command is used to take a snapshot of a replicated data volume containing a VxFS file system on the Secondary, mounting the snapshot volume in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/snapshot_volume
is corrupted. needs checking
```

This happens because the file system may not be quiesced before running the `vradmin ibc` command and therefore, the snapshot volume containing the file system may not be fully consistent.

Issue 2:

After a global clustering site failover, mounting a replicated data volume containing a VxFS file system on the new Primary site in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/data_volume
is corrupted. needs checking
```

This usually happens because the file system was not quiesced on the original Primary site prior to the global clustering site failover and therefore, the file systems on the new Primary site may not be fully consistent.

Workaround:

The following workarounds resolve these issues.

For issue 1, run the `fsck` command on the snapshot volume on the Secondary, to restore the consistency of the file system residing on the snapshot.

For example:

```
# fsck -F vxfs /dev/vx/dsk/dg/snapshot_volume
```

For issue 2, run the `fsck` command on the replicated data volumes on the new Primary site, to restore the consistency of the file system residing on the data volume.

For example:

```
# fsck -F vxfs /dev/vx/dsk/dg/data_volume
```

In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon (1672410, 1672417, 1825031)

Issue: After upgrading VVR to an IPv6-only environment in 6.0 release, vradmin commands may not work when a colon is specified in the RVG, data volume(s) and/or SRL name. It is also possible that after upgrading VVR to an IPv6-only environment, vradmin createpri may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

Workaround:

Make sure that colons are not specified in the volume, SRL and RVG names in the VVR configuration

vradmin commands might fail on non-logowner node after logowner change (1810827)

When VVR is used for replicating shared disk groups in a Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) or Veritas Storage Foundation for Oracle RAC (SFRAC) environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg` commands:

```
vradmind not reachable on cluster peer
```

In addition, all other `vradmin` commands (except `vradmin printvol`) fail with the error:

```
"VxVM VVR vradmin ERROR V-5-52-488 RDS has configuration error related
to the master and logowner."
```

This is due to a defect in the internal communication sub-system, which will be resolved in a later release.

Workaround:

Restart `vradmind` on all the cluster nodes using the following commands:

```
# /etc/init.d/vras-vradmind.sh stop
# /etc/init.d/vras-vradmind.sh start
```

While vradmin commands are running, vradmind may temporarily lose heart beats (2071568, 2275444)

This issue may occasionally occur when you use `vradmin` commands to administer VVR. While the `vradmin` commands run, `vradmind` may temporarily lose heartbeats, and the commands terminate with the following error message:

```
VxVM VVR vradmin ERROR V-5-52-803 Lost connection to host host;
terminating command execution.
```

Workaround:

To resolve this issue

1. Depending on the application I/O workload and network environment, uncomment and increase the value of the `IPM_HEARTBEAT_TIMEOUT` variable in the `/etc/vx/vras/vras_env` on all the hosts of the RDS to a higher value. The following example increases the timeout value to 120 seconds.

   ```
   export IPM_HEARTBEAT_TIMEOUT
   IPM_HEARTBEAT_TIMEOUT=120
   ```

2. Restart `vradmind` on all the hosts of the RDS to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following on all the hosts of the RDS:

   ```
   # /etc/init.d/vras-vradmind.sh stop
   # /etc/init.d/vras-vradmind.sh start
   ```

vxassist relayout removes the DCM (145413)

If you perform a relayout that adds a column to a striped volume that has a DCM, the DCM is removed. There is no message indicating that this has happened. To replace the DCM, enter the following:

```
# vxassist -g diskgroup addlog vol logtype=dcm
```

Cannot relayout data volumes in an RVG from concat to striped-mirror (2129601)

This issue occurs when you try a relayout operation on a data volume which is associated to an RVG, and the target layout is a striped-mirror.

Workaround:

To relayout a data volume in an RVG from concat to striped-mirror

1. Pause or stop the applications.
2. Wait for the RLINKs to be up to date. Enter the following:

   ```
   # vxrlink -g diskgroup status rlink
   ```

3. Stop the affected RVG. Enter the following:

   ```
   # vxrvg -g diskgroup stop rvg
   ```

4. Disassociate the volumes from the RVG. Enter the following:

   ```
   # vxvol -g diskgroup dis vol
   ```

5. Relayout the volumes to striped-mirror. Enter the following:

   ```
   # vxassist -g diskgroup relayout vol layout=stripe-mirror
   ```

6. Associate the data volumes to the RVG. Enter the following:

   ```
   # vxvol -g diskgroup assoc rvg vol
   ```

7. Start the RVG. Enter the following:

   ```
   # vxrvg -g diskgroup start rvg
   ```

8. Resume or start the applications.

vradmin verifydata operation fails when replicating between versions 5.1 and 6.0 (2360713)

When replicating in a cross-version VVR environment consisting of hosts running Storage Foundation 5.1 and hosts running Storage Foundation 6.0, the `vradmin verifydata` command fails with the following error:

```
VxVM VVR vxrsync ERROR V-5-52-2222 [from host]: VxVM in.vxrsyncd
ERROR V-5-36-2125 Server volume access error during [assign volids]
volume path: [/dev/vx/dsk/dg/snapshot_volume] reason: [this could be
```

```
because a target volume is disabled or an rlink associated with a
target volume is not detached during sync operation].
```

Workaround:

There are two workarounds for this issue.

- Upgrade the hosts running Storage Foundation 5.1 to Storage Foundation 5.1SP1 or later and re-run the `vradmin verifydata` command.
- Follow the offline verification procedure in the "Verifying the data on the Secondary" section of the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*. This process requires ensuring that the secondary is up-to-date, pausing replication, and running the `vradmin syncrvg` command with the `-verify` option.

Replication hang when VVR logowner is on CVM slave node (2405943)

When VVR is used for asynchronous replication in shared disk group environment, one of the nodes of the cluster at the primary site is chosen as the logowner. When the logowner node is on a node which is a slave node for the underlying CVM cluster, in the presence of heavy I/O from a node that is not the logowner, it is possible to get into a replication hang. This is due to an internal defect which will be fixed in later releases.

Workaround:

Enable the PreOnline trigger of the RVGLogOwner agent so that the VVR logowner will always reside on the CVM master node. For the detailed procedure, refer to the RVGLogowner agent notes section in the *Veritas Cluster Server Bundled Agents Reference Guide*.

vradmin verifydata may report differences in a cross-endian environment (2834424)

When replicating between two nodes in a cross-platform environment, and performing an autosync or replication, the vradmin verifydata command may report differences. This is due to different endianness between the platforms. However, the file system on the secondary node will be consistent and up to date.

vradmin verifydata operation fails if the RVG contains a volume set (2808902)

In a VVR environment, the `vradmin verifydata` command fails with the following error if the replicated volume group (RVG) contains any volume set:

```
Message from Primary:
VxVM VVR vxrsync ERROR V-5-52-2009 Could not open device
```

```
/dev/vx/dsk/vvrdg/<volname> due to: stat of raw character volume path
failed
```

I/O hangs on the primary node when running vxrvg snaprestore operation (2762147)

In a CVR environment, if a secondary node is set as the logowner for an RVG, issuing the `vxrvg snaprestore` command on the primary node may result in an I/O hang.

The vxrecover command does not automatically recover layered volumes in an RVG (2866299)

The `vxrecover` command calls the recovery process for the top-level volume, which internally takes care of recovering its subvolumes. The `vxrecover` command does not handle layered volumes correctly. The recovery process fails to recover the subvolumes, which remain in the NEEDSYNC state.

Workaround:

Manually recover the layered volumes using the `vxvol` utility, as follows:

```
# vxvol -g diskgroup resync volume
```

vxassist and vxresize operations do not work with layered volumes that are associated to an RVG (2162579)

This issue occurs when you try a resize operation on a volume that is associated to an RVG and has a striped-mirror layout.

Workaround:

To resize layered volumes that are associated to an RVG

1. Pause or stop the applications.
2. Wait for the RLINKs to be up to date. Enter the following:

   ```
   # vxrlink -g diskgroup status rlink
   ```

3. Stop the affected RVG. Enter the following:

   ```
   # vxrvg -g diskgroup stop rvg
   ```

4. Disassociate the volumes from the RVG. Enter the following:

   ```
   # vxvol -g diskgroup dis vol
   ```

5 Resize the volumes. In this example, the volume is increased to 10 GB. Enter the following:

```
# vxassist -g diskgroup growto vol 10G
```

6 Associate the data volumes to the RVG. Enter the following:

```
# vxvol -g diskgroup assoc rvg vol
```

7 Start the RVG. Enter the following:

```
# vxrvg -g diskgroup start rvg
```

8 Resume or start the applications.

Veritas Storage Foundation for Databases (SFDB) tools known issues

The following are known issues in this release of Veritas Storage Foundation for Databases (SFDB) tools.

Some dbed SmartTier commands do not work correctly in non-POSIX locales (2138030)

Some dbed SmartTier commands do not work correctly in non-POSIX locale settings.

Workaround:

Set the environment variable `LANG=C` systemwide in the `/etc/profile` file.

Some dbed operations may fail in system configurations where the hostname “localhost” cannot be resolved [3436609]

With hostname “localhost” that fails to get resolved, many dbed operations may fail. For example, the “vxsfadm –o valid” operation fails with the following error messages:

```
bash-4.1$ /opt/VRTSdbed/bin/vxsfadm -s sos -a oracle -o valid -c \
/tmp/sn7130
Use of uninitialized value in concatenation (.) or string
at /opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 2119.
Use of uninitialized value in string at \
/opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 2120.

SFDB vxsfadm ERROR V-81-0728 The directory \
/etc/vx/vxdba/oracle/local/.sfae could not be created.
```

```
Reason: Operating system returned error: No such file or directory
```

Workaround: Ensure that the name "localhost" resolves to the local loopback interface address (e.g. 127.0.0.1). You can verify whether "localhost" name can be resolved on your host by using the `ping` command.

Example output on a system where "localhost" cannot be resolved:

```
bash-4.1# ping localhost
ping: unknown host localhost
```

Flashsnap clone fails under some unusual archivelog configuration on RAC (2846399)

In a RAC environment, when using FlashSnap, the archive log destination to snapshot must be a shared path, and must be the same across all the nodes. Additionally, all nodes must use the same archive log configuration parameter to specify the archive log destination. Configurations similar to the following are not supported:

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_2='location=/tpcc_arch'
tpcc3.log_archive_dest_3='location=/tpcc_arch'
```

Where tpcc1, tpcc2, and tpcc3 are the names of the RAC instances and /tpcc_arch is the shared archive log destination.

Workaround:

To use FlashSnap, modify the above configuration to *.log_archive_dest_1='location=/tpcc_arch'. For example,

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_1='location=/tpcc_arch'
tpcc3.log_archive_dest_1='location=/tpcc_arch'
```

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SFHA. There is no workaround at this point of time.

Database Storage Checkpoint unmount may fail with device busy (2591463)

In some cases, when a database that is cloned using a Database Storage Checkpoint is shut down, an error similar to the following may occur:

```
SFAE Error:0457: Failed to unmount device
/dev/vx/dsk/datadg/datavol:Ckpt_1317707593_rw_1317708154.
Reason: VxFS returned error : umount: /tmp/clonedb/data: device is
busy
```

Workaround:

As an Oracle user, force shut down the clone database if it is up and then retry the unmount operation.

Attempt to use SmartTier commands fails (2332973)

The attempts to run SmartTier commands such as `dbdst_preset_policy` or`dbdst_file_move` fail with the following error:

```
fsppadm: ERROR: V-3-26551: VxFS failure on low level mechanism
with message - Device or resource busy
```

This error occurs if a sub-file SmartTier command such as `dbdst_obj_move` has been previously run on the file system.

There is no workaround for this issue. You cannot use file-based SmartTier and sub-file SmartTier simultaneously.

Attempt to use certain names for tiers results in error (2581390)

If you attempt to use certain names for tiers, the following error message is displayed:

```
SFORA dbdst_classify ERROR V-81-6107 Invalid Classname BALANCE
```

This error occurs because the following names are reserved and are not permitted as tier names for SmartTier:

- BALANCE
- CHECKPOINT
- METADATA

Workaround:

Use a name for SmartTier classes that is not a reserved name.

Clone operation failure might leave clone database in unexpected state (2512664)

If the clone operation fails, it may leave the clone database in an unexpected state. Retrying the clone operation might not work.

Workaround:

If retrying does not work, perform one the following actions depending on the point-in-time copy method you are using:

- For FlashSnap, resync the snapshot and try the clone operation again.
- For FileSnap and Database Storage Checkpoints, destroy the clone and create the clone again.
- For space-optimized snapshots, destroy the snapshot and create a new snapshot.

Contact Symantec support if retrying using the workaround does not succeed.

FlashSnap resync fails if there is an existing space-optimized snapshot (2479901)

If you try a FlashSnap resync operation when there is an existing space-optimized snapshot, the resync operation fails with the following error:

```
Error: VxVM vxdg ERROR V-5-1-4597 vxdg join FS_oradg oradg failed
datavol_snp : Record already exists in disk group
archvol_snp : Record already exists in disk group
```

Workaround:

Destroy the space-optimized snapshot first and then perform the FlashSnap resync operation.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0.5 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.0.5.

When upgrading from SFHA version 5.0 to SFHA 6.0.5 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File:  is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround

Before running `sfua_rept_migrate`, rename the startup script NO_S*vxdbms3 to S*vxdbms3.

Clone command fails if PFILE entries have their values spread across multiple lines (2844247)

If you have a `log_archive_dest_1` in single line in the init.ora file, then `dbed_vmclonedb` will work but `dbed_vmcloneb` will fail if you put in multiple lines for `log_archive_dest_1`.

Workaround

There is no workaround for this issue.

Frequent occurrence of SFDB remote or privileged command error (2869262)

If you installed a single instance database and try to run SFDB-related commands, then an error similar to the following might occur:

```
$ /opt/VRTSdbed/bin/dbed_update

No repository found for database faildb, creating new one.

SFDB vxsfadm ERROR V-81-0450 A remote or privileged command could not
be executed on host1

Reason: This can be caused by the host being unreachable or the vxdbd
daemon not running on that host.

Action: Verify that the host swpa04 is reachable.  If it is, verify
that the vxdbd daemon is running using the /opt/VRTS/bin/vxdbdctrl
status command, and start it using the /opt/VRTS/bin/vxdbdctrl start
command if it is not running.
```

There is no workaround at this point of time.

Data population fails after datafile corruption, rollback, and restore of offline Storage Checkpoint (2869259)

Sometimes when a datafile gets corrupted below its reservation size, the rollback may not pass and the file may not be rolled back correctly.

There is no workround at this point of time.

Storage Checkpoint clone fails if the `archive log` destination is same as the datafiles destination (2869266)

Storage Checkpoint cloning fails if the `archive log` destination is the same as the datafiles destination. The error is similar to:

```
Use of uninitialized value $path in hash element
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 121.
Use of uninitialized value $path in concatenation (.) or string
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 124.
Use of uninitialized value $path in pattern match (m//)
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 126.

SFDB vxsfadm ERROR V-81-0564 Oracle returned error.

Reason: ORA-02236: invalid file name (DBD ERROR: error possibly near
<*> indicator at char 172 in 'CREATE CONTROLFILE REUSE SET DATABASE
'TClone03' RESETLOGS NOARCHIVELOG
```

Workaround: For the 6.0.5 release, create distinct archive and datafile mounts for the Storage Checkpoint service.

FileSnap detail listing does not display the details of a particular snap (2846382)

FileSnap does not support displaying a detailed listing of a snapshot or clone. FileSnap only supports displaying a summary of all the snapshots or clones. For example, for the CLI `vxsfadm -s filesnap -a oracle --name=snap1 -o list`, a summary listing all the snapshots is displayed, instead of a detailed listing of a particular snapshot.

Workaround:

There is no workaround for this issue.

'vxdbd' process is online after Flash archive installation (2869269)

After a Flash archive installation of the SF stack, the `vxdbd` process is up, even if the stack is not configured.

Workaround: You can ignore, or stop the `vxdbd` process using the `/opt/VRTSdbed/common/bin/vxdbdctrl stop`command.

Storage Checkpoint clone fails in CFS environment if cloned using same Storage Checkpoint and same clone name on both nodes (2869268)

The Storage Checkpoint clone of an oracle database fails in a CFS environment, if you create a clone with a clone name and Storage Checkpoint name same as another clone up on a different CFS node.

Workaround:

There is no workaround. Create a clone with a different clone name.

Very long off-host cloning times for large number of datafiles (2849540)

When cloning off-host in certain Oracle database configurations, particularly with several hundred datafiles, the cloning can take a very long time, upto an hour or more. This problem does not cause the cloning to fail. The problem applies to all services such as FlashSnap, Space-optimized snapshots, FileSnap, and Storage Checkpoint.

Workaround:

There is no workaround at this point of time.

`sfua_rept_migrate` fails after phased SFRAC upgrade from 5.0MP3RP5 to 6.0.1 (2874322)

Command `sfua_rept_migrate` sometimes gives an error when upgrading to 6.0.1, and fails to unmount the repository volume. The error message is similar to:

```
# ./sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
Unmounting SFUA Sybase ASA repository.
UX:vxfs umount: ERROR: V-3-26388: file system /rep has been mount
locked
SFORA sfua_rept_migrate ERROR V-81-5550 umount /dev/vx/dsk/repdg/repvol
```

```
failed.
SFORA sfua_rept_migrate ERROR V-81-9162 Failed to umount repository.
```

Workaround:

The error does not hamper the upgrade. The repository migration works fine, but the old repository volume does not get unmounted. Unmount the mount using the manual option.

For example, use `/opt/VRTS/bin/umount -o mntunlock=VCS /rep`.

For more information, see TECH64812.

Veritas Storage Foundation and High Availability known issues

For known issues of Veritas Storage Foundation and High Availability, refer toVeritas Storage Foundation known issues and Veritas Cluster Server known issues.

Veritas Storage Foundation Cluster File System High Availability known issues

This section describes the known issues in this release of Veritas Storage Foundation Cluster File System High Availability.

ClearClone attribute not added automatically to `CVMTypes.cf` after upgrade [3483394]

The ClearClone attribute is not added automatically to the CVMVolDg agent configuration file after upgrade.

Workaround:

Manually add ClearClone attribute to the `/etc/VRTSvcs/conf/config/CVMTypes.cf` file as follows:

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/haattr -add CVMVolDg ClearClone -integer 0
# /opt/VRTSvcs/bin/hatype -modify CVMVolDg ArgList -add ClearClone
# /opt/VRTSvcs/bin/haconf -dump -makero
```

If you roll back the software from 6.0.5 Maintenance Release, use the following steps:

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/hatype -modify CVMVolDg ArgList -delete ClearClone
# /opt/VRTSvcs/bin//haattr -delete CVMVolDg ClearClone
# /opt/VRTSvcs/bin/haconf -dump -makero
```

The svsdatastore(1M) command may set the return value to ‘zero’ even when an error is reported (3313498)

The svsdatastore(1M) command may set the return value to ‘zero’ even when an error is reported.

For example,

```
#svsdatastore add <invalid disk name> Error: V-35-585: Disk
invaliddisk does not exists # echo $? 0
```

Workaround: There is no workaround for this issue.

The internal test is unresponsive due to high memory consumption due to fork failure (3402643)

The high swap usage issue is observed with Solaris 11 Update 1 bits and predominantly with internal stress or noise testing. The high usage causes the file system to become unresponsive.

You may not observe this issue during the stress load, which is induced using vxbench or TPCC. Thus, this issue is considered a corner case for production.

Workaround: There is no workaround for this issue.

Internal stress test on cluster file system hits debug assert in GLM [3364309]

In Group Lock Manager (GLM), the code to handle last revoke for a lock may cause deadlock, which is caught upfront by debug assert.

Workaround: There’s no workaround for the issue.

NFS issues with VxFS Storage Checkpoint (2027492)

NFS clients mounting VxFS Storage Checkpoints that are NFS-exported by SFCFSHA or SFHA cluster nodes using a virtual IP may receive the following error message upon virtual IP failover:

```
Stale NFS file handle
```

This is a result of major numbers of VxFS Storage Checkpoints not necessarily being the same on all SFCFSHA or SFHA cluster nodes.

There is no workaround at this time.

The mount command may hang when there are large number of inodes with extops and a small vxfs_ninode, or a full fsck cannot fix the link count table corruptions (2689326)

You might encounter one of the following issues:

- If there are large number of inodes having extended operations (extops), then the number of inodes used by the `mount` command reaches the maximum number of inodes that can be created in core. As a result, the `mount` command will not get any new inodes, which causes the `mount` command to run slowly and sometimes hang.
  Workaround: Increase the value of vxfs_ninode.
- The link count table (LCT) file can get damaged such that the flag is set, but the attribute inode is already freed. In this case, the mount command tries to free an inode that has been already freed thereby marking the file system for a full structural file system check.
  Workaround: There is no workaround for this issue.

CFS commands might hang when run by non-root (3038283)

The CFS commands might hang when run by non-root.

Workaround

To resolve this issue

- Use `halogin` command to save the authentication information before running any CFS commands on a non-root sessions.

  When you run the `halogin` command, VCS stores encrypted authentication information in the user's home directory.

Miscalculated file set usage (2123429)

When file set quotas are enabled, it may be possible for VxFS to get into a state where it thinks a very large number of blocks are allocated to Storage Checkpoints. This issue can be seen using the `fsckptadm` command:

```
# fsckptadm getquotalimit /mnt1
Filesystem   hardlimit   softlimit   usage   action_flag
/mnt1        10000       10000       18446744073709551614
```

This could cause writes to Storage Checkpoints to fail. It could also trigger the removal of removable Storage Checkpoints.

Workaround

If this occurs, disabling and re-enabling file set quotas causes VxFS to recalculate the number of blocks used by Storage Checkpoints:

```
# fsckptadm quotaoff /mnt1
# fsckptadm quotaon /mnt1
# fsckptadm getquotalimit /mnt1
Filesystem   hardlimit   softlimit   usage   action_flag
/mnt1        10000       10000       99
```

Multiple CFSmount resources are in a single service group they may not all come online after a reboot (2164670)

In some cases when multiple CFSmount resources are in a single service group, they all may not come online after a reboot. You will need to manually bring them online after a reboot.

Workaround

Create a resource dependency between the various CFSmount resources.

Panic due to null pointer de-reference in vx_bmap_lookup() (3038285)

A null pointer dereference in the `vx_bmap_lookup`() call can cause a panic.

Workaround: Resize the file system with the `fsadm` command from the primary node of the cluster.

fsdedupadm status shows FAILED on already deduped filesystem [2715413]

If a file system is already deduped, then after dryrun fsdedupadm status shows FAILED all time.

Workaround: Only do the dryrun on a fresh file system. Or If you do dryrun on an already deduped file system, then start dedup again:

fsdedupadm start *mountpoint*

Veritas Storage Foundation for Oracle RAC known issues

This section describes the known issues in this release of Veritas Storage Foundation for Oracle RAC.

Oracle RAC issues

This section lists the known issues in Oracle RAC.

Oracle Grid Infrastructure installation may fail with internal driver error [3473408]

The Oracle Grid Infrastructure installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

Workaround:

Perform one of the following steps depending on the type of installer you use for the installation:

- Script-based installer
  Export the `OUI_ARGS` environment variable, before you run the SFHA installation program:

  ```
  export OUI_ARGS=-ignoreInternalDriverError
  ```

  For more information, see the Oracle Metalink document: 970166.1
- Web-based installer
  When you run the Web-based installer, in the **Enter the arguments to be passed to the Oracle installer** text box, enter the value `-ignoreInternalDriverError`.
  For more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

During installation or system startup, Oracle Grid Infrastructure may fail to start [1933542]

After successful installation of Oracle RAC 11g Release 2 Grid Infrastructure, while executing the `root.sh` script, `ohasd` may fail to start. Similarly, during system startup, Oracle Grid Infrastructure may fail to start though the VCS engine logs may indicate that the cssd resource started Oracle Grid Infrastructure successfully.

The following message may be displayed on running the `strace` command:

```
# /usr/bin/strace -ftt -p pid_of_ohasd.bin
14:05:33.527288 open("/var/tmp/.oracle/npohasd",
O_WRONLY <unfinished ...>
```

For possible causes and workarounds, see the Oracle Metalink document: 1069182.1

Enabling ODM in Oracle RAC 11 Release 2 installations causes errors (1913013)

Enabling ODM in Oracle RAC 11 Release 2 installations causes the following error:

```
'ODM ERROR V-41-4-1-253-12 Not enough space'
Oracle instance may also crash with same error.
```

The error is observed if the DISM (Dynamic Intimate Shared memory) feature is enabled. In Solaris, the Oracle database uses DISM if it is available on the system, and if the value of the `sga_max_size` initialization parameter is larger than the size required for all SGA components combined.

Workaround: Make sure that the file `ORACLE_HOME/bin/oradism` is owned by the root user with "execute" and "setuid" permissions. If the problem persists after correcting the permissions, uncomment the `sga_max_size` and `memory_target` `init.ora` parameters.

Oracle VIP Configuration Assistant fails with an error message (1182220)

During Oracle RAC 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

```
The given interface(s), "" is not public.
Public interfaces should be used to configure virtual IPs.
```

This message appears only when the VIP is not from the regular public IP range (for example, 200.).

Workaround: Invoke the `vipca` utility manually as the superuser.

```
# export DISPLAY=nebula:0.0
# $CRS_HOME/bin/vipca
```

Oracle Cluster Verification utility displays a warning message

During the final stage of Oracle RAC 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification utility.

For example:

```
Utility
============================================================
OUI-25031: Some of the configuration assistants failed. It is
strongly recommended that you retry the configuration
assistants at this time. Not successfully running any "
Recommended" assistants means your system will not be correctly
configured.
1. Check the Details panel on the Configuration Assistant Screen
```

```
to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button
to retry them.
============================================================
```

Workaround: You may safely ignore this message if the cluster is operating satisfactorily.

Oracle Database Configuration Assistant displays an error

The Database Configuration Assistant utility displays the following error:

```
SGA size cannot be greater than maximum shared memory
segment size (0).
```

Workaround: Ignore this message and manually configure the database memory parameters for Oracle. In the "Memory" tab of the Oracle Database Creation Assistant (DBCA), select a Custom and Manual shared memory management configuration and enter the appropriate values.

Oracle Universal Installer (OUI) abruptly closes without responding

During the installation of Oracle Grid Infrastructure 11.2.0.4 on SF Oracle RAC nodes running Solaris 10/11, the Oracle Universal Installer (OUI) abruptly closes without responding if the “Shared File System” option is selected on the “Oracle Cluster Registry” page of the installer.

Workaround:

- For “Typical Installation” of Oracle Grid Infrastructure 11.2.0.4, enter:

  ```
  # ./runInstaller -J-Doracle.install.cvu.getSharedPartitionList=false \
  -J-Doracle.install.cvu.checkSharedStorageFileSystemType=false \
  -J-Doracle.install.grid.validate.QuickInstallUI=false
  ```

  Note: The option "-J-Doracle.install.grid.validate.QuickInstallUI=false" skips validations on the "Specify Install Locations" page.

  Ensure that you provide correct inputs for all the fields on this page, that is, manually verify the inputs before proceeding with the installation.

- For “Advanced Installation” of Oracle Grid Infrastructure 11.2.0.4, enter:

  ```
  # ./runInstaller -J-Doracle.install.cvu.getSharedPartitionList=false \
  -J-Doracle.install.cvu.checkSharedStorageFileSystemType=false \
  ```

```
-J-Doracle.install.grid.validate.OCRStorageUI=false \
-J-Doracle.install.grid.validate.VDSKStorageUI=false
```

Note: The options skip all the checks for Oracle Cluster Registry and voting disk storage. Ensure that you manually verify the location that is provided before proceeding with the installation.

See http://www.symantec.com/docs/TECH213369 for more details.

SFHA issues

This section lists the known issues in SFHA for this release.

PrivNIC and MultiPrivNIC agents not supported with Oracle RAC 11.2.0.2 and later versions

The PrivNIC and MultiPrivNIC agents are not supported with Oracle RAC 11.2.0.2 and later versions.

For more information, see the following Technote:

http://www.symantec.com/business/support/index?page=content&id=TECH145261

Issue with format of the last 8-bit number in private IP addresses (1164506)

The PrivNIC/MultiPrivNIC resources fault if the private IP addresses have a leading 0 in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1. or X.0X.X.1 or 0X.X.X.1, where X is an octet of the IP address.

When you configure private IP addresses for Oracle Clusterware, ensure that the IP addresses have a format as displayed in the following two-node example:

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

Confirm the correct format by viewing the PrivNIC or MultiPrivNIC resource in the `/etc/VRTSvcs/conf/config/main.cf` file.

Warning message displayed on taking cssd resource offline if LANG attribute is set to "eucJP" (2123122)

When you take the cssd resource offline using the `hares -offline cssd` command and the LANG attribute is set to “eucJP”, the following message may be observed in the `hamsg engine_A` command output:

```
VCS INFO V-16-2-13716 Could not find message V-16-2-13716
```

You may ignore the message.

PrivNIC resource faults in IPMP environments on Solaris 11 systems (2838745)

The PrivNIC resource faults on Solaris 11 systems when private interfaces used by IPMP are configured under PrivNIC resource.

Workaround: Avoid using PrivNIC or MultiPrivNIC agents in IPMP environments.

Error displayed on removal of VRTSjadba language package (2569224)

Removal of the VRTSjadba language package displays the following error on the screen:

```
Executing postremove script.
Generating BMC map file...
bmcmap ERROR V-33-1000-10001 Unable to create BMC map
```

You may ignore the error.

Oracle Universal Installer fails to start on Solaris 11 systems (2784560)

The Oracle Universal Installer (OUI) fails to start when the SF Oracle RAC installer invokes the OUI for the installation of Oracle Clusterware/Grid Infrastructure software.

Workaround: Install the following packages before installing Oracle Clusterware/Grid Infrastructure.

```
SUNWxwplt
SUNWmfrun
```

For instructions, see the Oracle documentation.

On nodes with heavy load, the CSSD resource may fault [3404403]

The CSSD agent checks the status of Oracle Clusterware using the Oracle Clusterware command `crsctl check crs`. On nodes with heavy load, the command does not complete within the period that the MonitorTimeout defines. After the 4 (default value of the FaultOnMonitorTimeout attribute) successive monitor timeouts, the CSSD resource goes to the FAULT state.

Workaround: Set the value of the FaultOnMonitorTimeouts attribute to 0 and use the AlertOnMonitorTimeouts attribute.

1 Change the permission on the VCS configuration file to read-write mode. Enter:

```
# haconf -makerw
```

2 Set the AlertOnMonitorTimeouts attribute value to 4 for the CSSD resource. Enter:

```
# hatype -display Application | grep AlertOnMonitorTimeouts
Application AlertOnMonitorTimeouts 0

# hares -override cssd_resname AlertOnMonitorTimeouts
# hatype -modify Application AlertOnMonitorTimeouts 4
```

3 Set the FaultOnMonitorTimeouts attribute value to 0 for the CSSD resource. Enter:

```
# hatype -display Application | grep FaultOnMonitorTimeouts
Application FaultOnMonitorTimeouts 4
# hares -override cssd_resname FaultOnMonitorTimeouts
# hatype -modify Application FaultOnMonitorTimeouts 0
```

4 Verify the AlertOnMonitorTimeouts and FaultOnMonitorTimeouts settings. Enter:

```
# hatype -display Application |
egrep "AlertOnMonitorTimeouts|FaultOnMonitorTimeouts”
Application  AlertOnMonitorTimeouts 4
Application  FaultOnMonitorTimeouts 0
```

5 Change the permission on the VCS configuration file to read-only mode. Enter:

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

Veritas Storage Foundation for Sybase ASE CE known issues

This section describes the known issues in this release of Veritas Storage Foundation for Sybase ASE CE.

SFHA issues

This section lists the known issues in SFHA for this release.

Sybase Agent Monitor times out (1592996)

Problem: The Sybase Agent Monitor has issues of timing out, in cases where qrmutil reports delay.

The Sybase Agent monitor times out, if qrmutil fails to report the status to the agent within the defined MonitorTimeout for the agent.

Solution: If any of the following configuration parameters for Sybase Database is increased, it will require a change in its MonitorTimeout value:

- quorum heartbeat interval (in seconds)
- Number of retries

If the above two parameters are changed, Symantec recommends that the MonitorTimeout be set to a greater value than the following: ((number of retries + 1) * (quorum heartbeat interval)) + 5.

Installer warning (151550)

Problem: During configuration of Sybase instance under VCS control, if the quorum device is on CFS and is not mounted, the following warning message appears on the installer screen:

```
Error: CPI WARNING V-9-40-5460 The quorum file
/qrmmnt/qfile cannot be accessed now. This may be due to a
file system not being mounted.
```

The above warning may be safely ignored.

Unexpected node reboot while probing a Sybase resource in transition (1593605)

Problem: A node may reboot unexpectedly if the Sybase resource is probed while the resource is still in transition from an online to offline state.

Normally the monitor entry point for Sybase agent completes with 5-10 seconds. The monitor script for the Sybase agent uses the qrmutil binary provided by Sybase. During a monitor, if this utility takes longer time to respond, the monitor entry point will also execute for longer duration before returning status.

Resolution: During the transition time interval between online and offline, do not issue a probe for the Sybase resource, otherwise the node may reboot.

Unexpected node reboot when invalid attribute is given (2567507)

Problem: A node may reboot unexpectedly if the Home, Version, or Server attributes are modified to invalid values while the Sybase resources are online in VCS.

Resolution: Avoid setting invalid values for the Home, Version, or Server attributes while the Sybase resources are online in VCS, to avoid panic of the node.

AutoFailOver = 0 attribute absent in the sample files at /etc/VRTSagents/ha/conf/Sybase (2615341)

Problem: AutoFailOver = 0 attribute is not present in the sample files at /etc/VRTSagents/ha/conf/Sybase.

Resolution: If you copy the main.cf file from the /etc/VRTSagents/ha/conf/Sybase location, add the AutoFailOver = 0 attribute to the binmnt and sybasece service groups.

Symantec VirtualStore known issues

This section describes the known issues in this release of Symantec VirtualStore.

Symantec VirtualStore issues

The svsiscsiadm create lun command fails if you create a LUN greater than the available space on the file system (2567517)

The svsiscsiadm create lun command fails if you create a LUN of a size greater than the total amount of space available on the file system. The underlying iscsitadm command fails with the following error message:

```
iscsitadm: Error Requested size is too large for system
```

The report of this error is logged in the /var/VRTSvcs/log/engine_A.log file.

If you then try to create a LUN on the same target, the LUN creation call fails again with the following error message:

```
iscsitadm: Error Failed to create a symbolic link to the backing store
```

The report of this error is logged in the /var/VRTSvcs/log/engine_A.log file.

This makes the target unusable.

Workaround

To resolve this issue

1. Note the TargetID and LunID on which the `svsiscsiadm create lun` command failed. To find the failed LunID, note the last LunID for the target on which `svsiscsiadm create lun` command failed with the use of the `svsiscsiadm list` command. To calculate the failed LunID, add 1 to last LunID seen by the `svsiscsiadm list` command.
2. Go to the configuration directory for the TargetID:

   ```
   # cd /etc/iscsi/TargetID .
   ```

3. Delete the symlink pointing to path of LUN backing file which failed to get added. The below LunID is the failed LunID, which is a result of calculation in point 1:

   ```
   # rm -f /etc/iscsi/TargetID/lun.($lunid + 1)
   ```

After removal of the symlink you should be able to add LUNs on the unusable target.

VirtualStore machine clones created while the VirtualStore cluster reboots will probably not start (2164664)

In some cases when you clone while rebooting the SVS nodes, you may receive several of the following error messages:

```
clone vms could not start X server
```

Workaround

Delete all the clones that got created while the node crashed and redo the cloning operation.

Cloning may not work (2348628)

If you cannot clone and you are using the VMware vAPP and OVF templates, then you must disable the vApp.

Workaround

To disable the vAPP

1. In VI Client, right-click on the virtual machine > **Edit Settings** > **Options** > **vApp Options**.
2. Click **Disable**.

Need intelligent NDMP/NBU backups for virtual machines (2378396)

When using NDMP or the NBU client to backup a virtual machine, the space consumed by the backup is equivalent to the size of the disks in the virtual machine, even though not all of the disk space in the virtual machine is used.

If a VMDK (Virtual Machine Disk) file is 10GB in size, but only consumes 1GB of disk space, an backup done by NDMP or the NBU client generates 10GB of backup data, even though the original VMDK file contains 9GB of unassigned disk space.

Workaround

Use VMware-specific backup applications (such as NetBackup for VMware) to create space-efficient backups.

Virtual machines created by the Symantec Quick Clone Virtual Machine Wizard might not boot correctly if during the process the FileStore cluster node, the ESX Server, or the vCenter Server reboots (2164664, 2374229)

In some cases when you clone using the wizard, and one of the following servers crashes or reboots while the clone process is in progress, the clones might not get created correctly:

- FileStore nodes
- ESX host on which the clones are being created
- vCenter Server

Even if the clones appear in the vCenter inventory as created, the clones GuestOS might not be able to boot.

Workaround

Delete all of the clones that were created when the servers crashed or were rebooted, and redo the wizard operation.

Error message does not always display when you select an incorrect cluster to clone (2372713)

In cases where multiple FileStore clusters are registered with the same Virtual Center, the Symantec Quick Clone Virtual Machine Wizard might not provide a warning that you selected an incorrect cluster to clone a golden image. This could happen if all of the FileStore clusters are exporting the same file system path, such as `/mnt`. Instead of an advanced warning that you selected the wrong cluster, you instead see an error on the final page of the wizard when the wizard attempts to clone the disks (vmdks) of the golden image. The error that displays is similar to the following example:

```
/mnt/goldvm/goldvm.vmdk no such file or directory...
```

Workaround

There is no workaround for this issue.

Cloning issue in a Japanese environment (2623471)

You might be unable to clone with Guest OS Customization or VMware View integration. While using the FileSnap wizard, options may be missing or in an error state.

Workaround

The workaround involves temporarily running the vCenter Server in the English locale.

To resolve this issue

1. From the vCenter Server, stop the following services using the Task Manager or services.msc:

   ```
   VMware VCMDS
   VMware VirtualCenter Server
   VMware VirtualCenter Management Webservices
   VMware vCenter Update Manager Services
   ```

2. Rename the following language directories `ja` to `ja-x`:

   ```
   C:\Program Files\VMware\Infrastructure\VirtualCenter Server\ja
   C:\Program Files\VMware\Infrastructure\VirtualCenter Server\locale\ja
   C:\Program Files\VMware\Infrastructure\VirtualCenter Server\imgres\ja
   ```

3. Restart the services from step 1.
4. Take FileSnap clones with customization and/or View integration using the FileSnap wizard.
5. To return the vCenter Server to the Japanese locale, reverse steps 1-3.

The Virtual machine's local Administrator password may be set to blank (2676078, 2676079)

When clones are made of a Windows 2008 Virtual Machine and Guest OS Customization is enabled, the Virtual Machine's local Administrator password is set to blank.

Workaround: There is no workaround for this issue.

The installer output states, "Registering SVS license," even if you enabled keyless licensing

When installing, if you enable keyless licensing, the installer's output includes the following message:

```
Registering SVS license
```

Workaround: This message is harmless and can be ignored. The product will successfully install without a license key.

Software limitations

This section covers the software limitations of this release.

Limitations related to installation

This is the limitations related to installation in the 6.0.5 release.

Limitations related to web-based installer for SFRAC

- Web-based installer on local disk is not supported.
- If SFRAC is not configured before upgrade, the web-based installer does not support to upgrade SFRAC to 6.0.5.

Limitations related to Install Bundles

- Web-based installer doesn't support the Install Bundles feature.
- The feature doesn't support native OS install or upgrade methods, such as JumpStart, flash archive, and so on.
- The Install Bundles feature for 6.0.5 does not support hot fix installation.

Documentation errata

The following sections cover additions or corrections for the product documentation. These additions or corrections may be included in later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).