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

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 6.0.2
- Storage Foundation and High Availability Solutions 6.0.3
- Storage Foundation and High Availability Solutions 6.0.4

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 RPMs

Table 1-1 lists the Red Hat Enterprise Linux 5 RPMs that are updated in this release.

Table 1-1 RPMs for Red Hat Enterprise Linux 5

Name	Version	Arch	Size in Bytes
VRTSamf	6.0.500.000	x86_64	1513348
VRTSaslapm	6.0.500.000	x86_64	251758
VRTScavf	6.0.500.000	i386	199167
VRTSdbac	6.0.500.000	x86_64	911204
VRTSdbed	6.0.500.000	x86_64	36607587
VRTSfsadv	6.0.500.000	x86_64	4784301
VRTSgab	6.0.500.000	x86_64	1005902
VRTSglm	6.0.500.000	x86_64	107952
VRTSllt	6.0.500.000	x86_64	951717
VRTSlvmconv	6.0.500.000	i686	73602
VRTSodm	6.0.500.000	x86_64	242846
VRTSperl	5.14.2.20	x86_64	17926880
VRTSsfcpi601	6.0.500.000	noarch	847792
VRTSspt	6.0.500.000	noarch	24337985
VRTSvcs	6.0.500.000	i686	63951917
VRTSvcsag	6.0.500.000	i686	16298256
VRTSvcsdr	6.0.500.000	x86_64	200286
VRTSvcsea	6.0.500.000	i686	327615
VRTSvcsvmw	6.0.500.000	i686	10069945
VRTSvxfen	6.0.500.000	x86_64	492756
VRTSvxfs	6.0.500.000	x86_64	11033719
VRTSvxvm	6.0.500.000	x86_64	31153223

Table 1-2 lists the Red Hat Enterprise Linux 6 RPMs that are updated in this release.

Table 1-2 RPMs for Red Hat Enterprise Linux 6

Name	Version	Arch	Size in Bytes
VRTSamf	6.0.500.000	x86_64	1609408
VRTSaslapm	6.0.500.000	x86_64	173528
VRTScavf	6.0.500.000	i386	180048
VRTSdbac	6.0.500.000	x86_64	779496
VRTSdbed	6.0.500.000	x86_64	36607587
VRTSfsadv	6.0.500.000	x86_64	3907800
VRTSgab	6.0.500.000	x86_64	1562636
VRTSglm	6.0.500.000	x86_64	119196
VRTSllt	6.0.500.000	x86_64	1407008
VRTSlvmconv	6.0.500.000	i686	66872
VRTSodm	6.0.500.000	x86_64	285428
VRTSperl	5.14.2.20	x86_64	14282320
VRTSsfcpi601	6.0.500.000	noarch	847792
VRTSspt	6.0.500.000	noarch	24337983
VRTSvcs	6.0.500.000	i686	48327332
VRTSvcsag	6.0.500.000	i686	12408720
VRTSvcsdr	6.0.500.000	x86_64	436364
VRTSvcsea	6.0.500.000	i686	238684
VRTSvcsvmw	6.0.500.000	i686	8945524
VRTSvxfen	6.0.500.000	x86_64	1692064
VRTSvxfs	6.0.500.000	x86_64	10470812
VRTSvxvm	6.0.500.000	x86_64	22251860

Table 1-3 lists the SUSE Linux Enterprise Server 10 RPMs that are updated in this release.

Table 1-3 RPMs for SUSE Linux Enterprise Server 10

Name	Version	Arch	Size in Bytes
VRTSamf	6.0.500.000	x86_64	5834901
VRTSaslapm	6.0.500.000	x86_64	201168
VRTScavf	6.0.500.000	i386	167428
VRTSdbac	6.0.500.000	x86_64	1743222
VRTSdbed	6.0.500.000	x86_64	36433491
VRTSfsadv	6.0.500.000	x86_64	4963999
VRTSgab	6.0.500.000	x86_64	2327983
VRTSglm	6.0.500.000	x86_64	141640
VRTSllt	6.0.500.000	x86_64	1646110
VRTSlvmconv	6.0.500.000	i586	64116
VRTSodm	6.0.500.000	x86_64	321717
VRTSperl	5.14.2.20	x86_64	15460166
VRTSsfcpi601	6.0.500.000	noarch	847792
VRTSspt	6.0.500.000	noarch	24337995
VRTSvcs	6.0.500.000	i586	88583929
VRTSvcsag	6.0.500.000	i586	15432536
VRTSvcsdr	6.0.500.000	x86_64	410261
VRTSvcsea	6.0.500.000	i586	272077
VRTSvcsvmw	6.0.500.000	i586	9917651
VRTSvxfen	6.0.500.000	x86_64	2410991
VRTSvxfs	6.0.500.000	x86_64	11199213
VRTSvxvm	6.0.500.000	x86_64	31148767

Table 1-4 lists the SUSE Linux Enterprise Server 11 RPMs that are updated in this release.

Table 1-4 RPMs for SUSE Linux Enterprise Server 11

Name	Version	Arch	Size in Bytes
VRTSamf	6.0.500.000	x86_64	2057440
VRTSaslapm	6.0.500.000	x86_64	193877
VRTScavf	6.0.500.000	i386	172198
VRTSdbac	6.0.500.000	x86_64	783895
VRTSdbed	6.0.500.000	x86_64	36433491
VRTSfsadv	6.0.500.000	x86_64	3523509
VRTSgab	6.0.500.000	x86_64	922831
VRTSglm	6.0.500.000	x86_64	130587
VRTSgms	6.0.400.000	x86_64	39450
VRTSllt	6.0.500.000	x86_64	810101
VRTSlvmconv	6.0.500.000	i586	65422
VRTSodm	6.0.500.000	x86_64	293459
VRTSperl	5.14.2.20	x86_64	13978050
VRTSsfcpi601	6.0.500.000	noarch	847792
VRTSspt	6.0.500.000	noarch	24338000
VRTSvcs	6.0.500.000	i686	76864215
VRTSvcsag	6.0.500.000	i686	12596264
VRTSvcsdr	6.0.500.000	x86_64	204168
VRTSvcsea	6.0.500.000	i686	234027
VRTSvcsvmw	6.0.500.000	i686	8972606
VRTSvxfen	6.0.500.000	x86_64	1036462
VRTSvxfs	6.0.500.000	x86_64	9162191
VRTSvxvm	6.0.500.000	x86_64	21467090

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, 6.0.2 or 6.0.4 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 can be 6.0.1, 6.0.2, or 6.0.4.

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>
```

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 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`.

KVMGuest agent supports Red Hat Enterprise Virtualization 3.1 and 3.2

KVMGuest agent is enhanced to support Red Hat Enterprise Virtualization 3.1 and 3.2. A new key `UseRHEVMManualFencing` is added to `RHEVMInfo` attribute of KVMGuest.

`RHEVMInfo` attribute specifies the information about the RHEV environment The keys associated with this attribute are as follows:

Enabled	Specifies the virtualization environment. The default value is 0. If the value is 0, it means that KVM environment is enabled. If the Value is 1, it means that RHEV environment is enabled.
URL	Specifies the RHEV-M URL that can be used for REST API communication. For example:https://rhevm-server.domain.com:443
User	Specifies the RHEV-M User that can be used for REST API communication. For example: admin@internal , rhevadmin@example.com
Password	Specifies the encrypted password of RHEV-M User. The password should be encrypted using `vcsencrypt` command.
Cluster	Specifies the name of the RHEV-M cluster, to which the VCS host belongs.
UseManualRHEVMFencing	Enables or disables the use of manual RHEV-M fencing if physical host on which virtual machine is running crashes. The default value is 0. The value 0 signifies that the use of manual RHEV-M fencing is disabled. The value 1 signifies that the use of manual RHEV-M fencing is enabled.

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:

In 6.0.5 release, Symantec has introduced support for Kove Xpress Disk (XPD) devices over InfiniBand. Please refer to the following link for more information regarding Kove Xpress Disk (XPD) devices:

http://kove.com/xpress

placeholder for technotelink

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 DB2 veriosn 10.5.

Note: When updating DB2 database parameters, ensure that any deferred values are updated before using SFDB tools. The DB2 database needs to be reactivated for the parameter values to take effect.

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 Linux operating systems

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

Table 1-5 shows the supported operating systems for this release.

Note: SF Oracle RAC will soon announce support for RHEL 6.

Refer to the following TechNote for the latest information on supported operating systems and Oracle database versions:

http://www.symantec.com/docs/TECH44807

The TechNote will be updated once Oracle supports RHEL6 and Symantec qualifies the same.

Table 1-5 Supported operating systems

Operating systems	Levels	Kernel version
Red Hat Enterprise Linux 6	Update 1	2.6.32-131.0.15.el6
	Update 2	2.6.32-220.el6
	Update 3	2.6.32-279.el6
	Update 4	2.6.32-358.el6
	Update 5	2.6.32-431.el6
Red Hat Enterprise Linux 5	Update 5	2.6.18-194.el5
	Update 6	2.6.18-238.el5
	Update 7	2.6.18-274.el5
	Update 8	2.6.18-308.el5
	Update 9	2.6.18-348.el5
	Update 10	2.6.18-371.el5
SUSE Linux Enterprise 11	SP1	2.6.32.12-0.7
	SP2	3.0.13-0.27
	SP3	3.0.76-0.11
SUSE Linux Enterprise 10	SP4	2.6.16.60-0.85.1
Oracle Linux 6	Update 1	2.6.32-131.0.15.el6
	Update 2	2.6.32-220.el6
	Update 3	2.6.32-279.el6
	Update 4	2.6.32-358.el6
	**Update 5	2.6.32-431.el6
Oracle Linux 5	Update 5	2.6.18-194.el5
	Update 6	2.6.18-238.el5
	Update 7	2.6.18-274.el5
	Update 8	2.6.18-308.el5
	Update 9	2.6.18-348.el5
	Update 10	2.6.18-371.el5

** Also supported on UEK R2 2.6.39-400.17.1.el6uek kernel

Note: Unbreakable Enterprise Kernel (UEK) is supported by VCS only. Storage Foundation does not support UEK.

Note: Only 64-bit operating systems are supported on the AMD Opteron or the Intel Xeon EM64T (x86_64) Processor line.

If your system is running an older version of either Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Oracle Linux, upgrade it before attempting to install the Veritas software. Consult the Red Hat, SUSE, or Oracle documentation for more information on upgrading or reinstalling your operating system.

Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain Kernel Application Binary Interface (KABI) compatibility.

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

Mandatory patch required for Oracle Bug 4130116

If you are running Oracle versions 9.2.0.6 or 9.2.0.7, you must apply the Oracle patch for Oracle Bug 4130116. Contact Oracle to obtain this patch, and for details on how to apply it.

For more information, refer to the following TechNote:

http://www.symantec.com/docs/HOWTO19718

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

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

VMware Environment

For information about the use of this product in a VMware Environment, refer to http://entsupport.symantec.com/docs/289033

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 1-6 SFDB features database support for 6.0.5

SFDB feature	DB2	Oracle	Sybase
Oracle Disk Manager, Cached Oracle Disk Manager	No	Yes	No
Concurrent I/O	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes
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 1-7 Installation and upgrades fixed issues in 6.0.5

Fixed issues	Description
3125494	Reporting $padv->QEMU not defined.
3131744	installmp CPI should not enable DMP native support ,it should ask customer if they want to enable it.
3243089	During a live upgrade the installation process takes more time than expected.
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.
3356724	I/O fencing cannot be configured in Oracle Enterprise Linux (OEL) 6.4 withUnbreakable Enterprise Kernel (UEK) kernel for 6.0.4 release.
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".

Table 1-7 Installation and upgrades fixed issues in 6.0.5 *(continued)*

Fixed issues	Description
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.
3347126	On UEK2/SLES11SP3, no support is provided for VRTSperl.
3355917	Support Linux Update.
3384775	Installing patch 6.0.3.200 on RHEL 6.4 or earlier RHEL 6.* versions fails with ERROR: No appropriate modules found.

Installation and upgrades: issues fixed in 6.0.4

In this release, there were no fixed issues related to installation and upgrades.

Installation and upgrades: issues fixed in 6.0.3

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

Table 1-8 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.

Installation related fixed issues in 6.0.2

Table 1-9 Installation related fixed issues

Incident	Description
2622987	If a host is not reporting to any management server but sfmh discovery is running before you upgrade to 6.0, sfmh-discovery may fail to start after the 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 1-10 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.
2585899	On RHEL, unable to create storage for OCR and Vote disk when using FQDN instead of using only the node name.
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.

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 1-11 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues

Fixed issues	Description
3263336	Internal noise test on cluster file system hits the "f:vx_cwfrz_wait:2" and "f:vx_osdep_msgprint:panic" debug asserts.
3462694	The fsdedupadm(1M) command fails with error code 9 when it tries to mount checkpoints on 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.

Table 1-11 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
3444771	Internal noise test on cluster file system hits debug assert while creating a file.
3317116	Internal command conformance text for mount command on RHEL6 Update4 hit debug assert inside the vx_get_sb_impl()function.
3463464	Internal kernel functionality conformance test hits a kernel panic due to null pointer dereference.
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.
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.
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.

Table 1-11 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
3274048	VxFS hangs when it requests a cluster-wide grant on an inode while holding a lock on the inode.
3364309	Internal stress test on cluster file system hit debug assert in Group Lock Manager (GLM).
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 6.0.4

Table 1-12 describes the incidents that are fixed in Veritas Storage Foundation Cluster File System High Availability in 6.0.4.

Table 1-12 Veritas Storage Foundation Cluster File System High Availability 6.0.4 fixed issues

Incident	Description
3259634	A Cluster File System having more than 4G blocks gets corrupted because the blocks containing some file system metadata get eliminated.
3248954	The system fails to build modules on SUSE Linux Enterprise Server 11 (SLES 11) SP3 for Veritas Group Lock Manager (GLM), because it cannot find the `utsrelease.h` file needed to build the GLM module for SP3 on SLES 11.
3248953	The system fails to build modules on SLES11 SP3 for Veritas Group Messaging Sevice (GMS), because it cannot find the `utsrelease.h` file needed to build the GMS module for SP3 on SLES 11.
3214328	There was a mismatch between the states for the glm grant level and the glm data in a cfs inode.

Table 1-12 Veritas Storage Foundation Cluster File System High Availability 6.0.4 fixed issues *(continued)*

Incident	Description
3192985	Checkpoints quota usage on Cluster File System (CFS) can be negative.
3189562	Oracle daemons get stuck on the GLM lock.
3047134	The system panics during the internal testing due to a Group Atomic Broadcast (GAB) callback routine in an interrupt context with the following message: `Kernel panic - not syncing: GLM assert GLM_SPIN_LOCK:1`
3011959	The system may panic because of the file system locking or unlocking using the `fsadm`(1M) or the `vxumount`(1M) command.
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.
2956195	`mmap` in the CFS environment takes a long time to complete.
2495673	Mismatch of concurrent I/O-related data in an inode is observed during communication between the nodes in a cluster.
2107152	The system panics when you `umount` a mntlock protected VxFS file system, if that device is duplicately mounted on different directories.

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

Table 1-13 describes the Veritas Storage Foundation Cluster File System fixed issues in 6.0.3.

Table 1-13 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.
2906018	The vx_iread errors are displayed after successful log replay and mount of the file system.

Table 1-13 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues *(continued)*

Incident	Description
2857731	Internal testing hits an assert "f:vx_mapdeinit:1" .
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.
2590918	Delay in freeing unshared extents upon primary switch over.

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 1-14 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 1-15 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 1-15 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.
2223258	The vxdisksetup(1M) command initializes the disk which already has Logical Volume Manager (LVM) or File System (FS) on it.
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.
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.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
2857360	The vxconfigd(1M) command hangs when the vol_use_rq tunable ofVxVM is changed from 1 to 0.
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.
2874810	When you install DMP only solutions using the installdmp command, the root support is not enabled.
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.
2882412	The 'vxdisk destroy' command uninitializes a VxVM disk which belongs to a deported disk group.
2893530	With no VVR configuration, when system is rebooted, it panicked.
2898324	UMR errors reported by Purify tool in "vradmind migrate" command.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
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.
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.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2957555	The vxconfigd(1M) daemon on the CVM master node hangs in the userland during the vxsnap(1M) restore operation.
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.
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.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
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.
3010830	During a script-based or web-based installation, the post check verification forVRTSvxvm and VRTSaslapm may fail due to changed user and group permissions of some files after installation.
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.
3038684	The restore daemon enables the paths of Business Continuance Volumes-Not Ready (BCV-NR) devices.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
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.
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.
3060697	The vxrootmir(1M) utility fails with the following error message:VxVM vxdisk ERROR V-5-1-5433 Device sdb: init failed.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
3088059	On Red Hat Enterprise Linux 6.x (RHEL6.x), the type of hostbus adapter (HBA) is reported as SCSI instead of FC.
3088907	A node in a Cluster Volume Manager can panic while destroying a shared disk group.
3091916	The Small Computer System Interface (SCSI) I/O errors overflow the syslog.
3091978	The lvm.conf variable preferred_names is set to use DMP even if the dmp_native_support tunable is 'off'.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
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.
3136272	The disk group import operation with the "-o noreonline" option takes additional import time.
3140835	When the reclaim operation is in progress using the TRIM interface, the path ofone of the trim capable disks is disabled and this causes a panic in the system.
3142315	Disk is misidentified as clone disk with udid_mismatch flag.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3162418	The vxconfigd(1M) command dumps core due to wrong check in ddl_find_cdevno() function.
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.
3186149	On Linux System with LVM version 2.02.85, on enablingdmp_native_support LVM volume Groups will disappear.
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.
3220929	The Veritas Volume Manager (VxVM) fails to convert a logicalvolume manager (LVM) volume to Veritas Volume Manager (VxVM) volume.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
3244217	Cannot reset the clone_disk flag during vxdg import.
3247040	vxdisk scandisks enables the PowerPath (PP) enclosure which was disabled previously.
3254311	System panics when reattaching site to a site-consistent diskgroup having volumelarger than 1.05 TB
3258531	The vxcdsconvert(1M) command fails with error "Plex column/offset is not 0/0 for new vol volume".
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.
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.
3273435	VxVM disk group creation or import with SCSI-3 Persistent Reservation (PR) fails.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
3287940	Logical unit number (LUN) from EMC CLARiiON array and having NR (Not Ready) state are shown in the state of online invalid by Veritas Volume Manager (VxVM).
3289202	Handle KMSG_EPURGE error in CVM disk connectivity protocols.
3321269	vxunroot may hang during un-encapsulation of root disk.
3323548	In the Cluster Volume Replicator (CVR) environment, a cluster-wide vxconfigd hang occurs on primary when you start the cache object.
3325022	The VirtIO-disk interface exported disks from an SLES11 SP2 or SLES11 SP3 host are not visible.
3326900	Warnings are observed during execution of the vxunroot command.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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).
3416098	The vxvmconvert utility throws error during execution.
3416622	The hot-relocation feature fails for a corrupted disk in the CVM environment.
3417044	System becomes unresponsive while creating a VVR TCP connection.
3421236	The vxautoconvert/vxvmconvert utility command fails to convert LVM (Logical Volume Manager) to VxVM (Veritas Volume Manager) causing data loss.
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.
3433931	The AvxvmconvertA utility fails to get the correct LVM version.
3435225	In a given CVR setup, rebooting the master node causes one ofthe slaves to panic.
3277258	BAD TRAP panic in vxio:vol_mv_pldet_callback.
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.
3250096	/dev/raw/raw# devices vanish on reboot when selinux enabled.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
3236772	Heavy I/O loads on primary sites result intransaction/session timeouts between the primary and secondary sites.
3235350	I/O on grown region of a volume leads to system panic if the volume has instant snapshot.
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.
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.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
3107741	The vxrvg snapdestroy command fails with the "Transaction aborted waiting for io drain" error message.
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.
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.
3056311	For release < 5.1 SP1, allow disk initialization with CDS format using raw geometry.
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.
2982834	The /etc/vx/bin/vxdmpraw command fails to create a raw device for a full device when enabling Dynamic Multi-Pathing (DMP) support for Automatic Storage Management (ASM) and also does not delete all the raw devices when the DMP support is disabled for ASM.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2945658	If the Disk label is modified for an Active/Passive LUN, then the current passive paths don't reflect this modification after a failover.
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.
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.
2750782	The Veritas Volume Manager (VxVM) upgrade process failsbecause it incorrectly assumes that the root disk is encapsulated.
2743870	Continuous I/O errors are seen when retry I/O.
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.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
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.
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.

Table 1-15 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3403172	The EMC Symmetrix Not-Ready (NR) device paths gets disabled or enabled unexpectedly.
3182350	If there are more than 8192 paths in the system, the vxassist(1M) command hangs when you create a new VxVM volume or increase the existing volume's size.

Veritas Volume Manager: issues fixed in 6.0.4

Table 1-16 describes the incidents that are fixed in Veritas Volume Manager in 6.0.4. (Only for SLES11)

Table 1-16 Veritas Volume Manager 6.0.4 fixed issues

Incident	Description
3321269	The command `vxunroot` may hang during un-encapsulation of root disk.
3240858	The `/etc/vx/vxesd/.udev_lock` file may have different permissions at different instances.
3199056	Veritas Volume Replicator (VVR) primary system panics in the `vol_cmn_err` function due to the VVR corrupted queue.
3186971	The logical volume manager (LVM) configuration file is not correctly set after turning on DMP native support. As a result, the system is unbootable.
3186149	On Linux System with LVM version 2.02.85, on enabling `dmp_native_support` LVM volume Groups will disappear
3182350	If there are more than 8192 paths in the system, the `vxassist(1M)` command hangs when you create a new VxVM volume or increase the existing volume's size.
3182175	The `vxdisk -o thin,fssize list` command can report incorrect File System usage data.
3177758	Performance degradation is seen after upgrade from SF 5.1SP1RP3 to SF 6.0.1 on Linux.
3162418	The `vxconfigd(1M)` command dumps core due to wrong check in `ddl_find_cdevno()` function.
3146715	Rlinks do not connect with Network Address Translation (NAT) configurations on Little Endian Architecture.

Table 1-16 Veritas Volume Manager 6.0.4 fixed issues *(continued)*

Incident	Description
3130353	Continuous disable or enable path messages are seen on the console for EMC Not Ready (NR) devices.
3121380	I/O of Replicated Volume Group (RVG) hangs after one data volume is disabled.
3091916	The Small Computer System Interface (SCSI) I/O errors overflow the syslog.
3063378	Some VxVM commands run slowly when EMC PowerPath presents and manages "read only" devices such as EMC SRDF-WD or BCV-NR.
3015181	I/O can hang on all the nodes of a cluster when the complete non-A or A class of storage is disconnected.
3012929	The `vxconfigbackup(1M)` command gives errors when disk names are changed.
3010191	Previously excluded paths are not excluded after upgrade to VxVM 5.1SP1RP3.
2986596	The disk groups imported with mix of standard and clone Logical Unit Numbers (LUNs) may lead to data corruption.
2972513	Cluster Volume Manager (CVM) and PGR keys from shared data disks are not removed after stopping VCS.
2959733	Handling the device path reconfiguration in case the device paths are moved across LUNs or enclosures to prevent the `vxconfigd(1M)` daemon coredump.
2905579	VxVM RPM installation on SLES11 SP2 fails for kernel version 3.0.26-0.7.6 and above.
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.
2857360	The `vxconfigd(1M)` command hangs when the `vol_use_rq` tunable of VxVM is changed from 1 to 0.
2857044	System crashes while resizing a volume with data change object (DCO) version 30.
3052770	The `vradmin syncrvg` operation with a volume set fails to synchronize the secondary RVG with the primary RVG.

Veritas Volume Manager: issues fixed in 6.0.3

Table 1-17 describes the incidents that are fixed in Veritas Volume Manager in 6.0.3.

Table 1-17 Veritas Volume Manager 6.0.3 fixed issues

Incident	Description
3002770	Accessing NULL pointer in dmp_aa_recv_inquiry() caused system panic.
2971746	For single-path device, bdget() function is being called for each I/O, which cause high cpu usage and leads to I/O performance degradation.
2970368	Enhancing handling of SRDF-R2 WD devices in DMP.
2965910	vxassist dump core with the `-o ordered` option.
2964169	In multiple CPUs environment, I/O performance degradation is seen when I/O is done through VxFS and VxVM specific private interface.
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.
2943637	DMP IO statistic thread may cause out of memory issue so that OOM(Out Of Memory) killer is invoked and causes system 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 the master is switched.
2933138	panic in voldco_update_itemq_chunk() due to accessing invalid buffer
2930569	The LUNs in 'error' state in output of 'vxdisk list' cannot be removed through DR(Dynamic Reconfiguration) Tool.
2928764	SCSI3 PGR registrations fail when dmp_fast_recovery is disabled.
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.

Table 1-17 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
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.
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.
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.
2878876	vxconfigd dumps core in vol_cbr_dolog() due to race between two threads processing requests from the same client.
2869594	Master node panics due to corruption if space optimized snapshots are refreshed and "vxclustadm setmaster" is used to select master.
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.

Table 1-17 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2839059	vxconfigd logged warning `cannot open /dev/vx/rdmp/cciss/c0d` device to check for ASM disk format.
2837717	The `vxdisk(1M) resize` command fails if `da name` is specified.
2836798	Prevent DLE on simple/sliced disk with EFI label
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
2801962	Grow of a volume takes significantly large time when the volume has version 20 DCO (Data Change Object) attached to it
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.
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

Table 1-17 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2273190	Incorrect setting of the 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 DG 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 the Nolicense flag can lead to dmp database inconsistency.
1859018	The `link detached from volume` warnings are displayed when a linked-breakoff snapshot is created.
1765916	VxVM socket files don't have proper write protection
1725593	The `vxdmpadm listctlr` command has to be enhanced to print the count of device paths seen through the controller.

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 and Cluster Volume Manager fixed issues.

Table 1-18 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.

Table 1-18 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2812355	CVS rolling upgrade : vxconfigd hung in join when tried to join upgraded slave node to cluster during upgrade from 5.1sp1rp2 to 6.0sp1 on "sles11sp1-Issue 2".
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.
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.
2751102	Random panics seen in vx_worklist_thr on SLES11 and VxFS.
2747032	Write is taking long time to complete when read/write happen simultaneously.
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.
2726148	System unbootable after `/usr/lib/vxvm/bin/vxupdatelvm` script updates filter in `lvm.conf` file.
2721807	Root disk encapsulation: On SLES11 SP2, machine went to maintenance mode during final reboot after encap.
2711312	Missing symbolic link is created after pulling FC cable on RHEL6.
2710579	Do not write backup labels for CDS disk - irrespective of disk size.

Table 1-18 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2710147	Node panics in dmp_pr_do_reg during key registration with fencing enabled.
2709743	Inplace upgrade is not working from 6.0.
2703858	Site failure (storage and all nodes including master node) led to 'configuration daemon not accessible' error on all the sites.
2701654	Phantom DMP disk partition causes panic.
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.
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.
2664825	Disk group import fails when disk contains no valid UDID tag on config copy and config copy is disabled.
2660151	The `vxconfigd` daemon is generating a series of LVM header messages for devices (CLONES/replicated devices). Secondary EMC MirrorView LUNS in an error state.
2656803	Race between `vxnetd start` and `stop` operations causes panic.
2652485	Inactive snapshot LUNs cause trespassing.
2648176	Performance difference on Master versus Slave during recovery with Data Change Object (DCO).

Table 1-18 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
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.
2644248	The `vxunroot` command fails as root partition "logvol" mounted on `/var/log`.
2643634	Message enhancement for a mixed (non-cloned and cloned) disk group import.
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.
2623182	vxvm-boot not cleaning up `/tmp/vx.*` directories whenever system reboot is done for Linux environment.
2620555	I0 hang due to SRL overflow & CVM reconfig.
2612301	Upgrading kernel on encapsulated boot disk does not work on Red Hat Enterprise Linux (RHEL) 5, 6, and SUSE Linux Enterprise Server (SLES) 10.
2607706	Encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path names are not the same.
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()`.
2564092	Automate the LUN provisioning (addition) / removal steps using `vxdiskadm`.
2553729	Status of the EMC Clariion disk changed to "online clone_disk" after upgrade.
2486301	"VXFS" RPM installation failed.
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.
1675482	The `vxdg list` *`dgname`* command gives error 'state=new failed'.

Table 1-18 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
1533134	DMP: depreciated SCSI `ioctl` use sg_io type of error.
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.
2682491	The vxvmconvert command displays an error message while converting an LVM volume.

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 1-19 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.
2794625	Unable to configure ASM to use DMP native block device path.
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.
2700086	EMC BCV (NR) established devices are resulting in multiple dmp events messages (paths being disabled/enabled).
2652485	Inactive snapshot luns cause trespassing.
2626199	vxdmpadm list dmpnode printing incorrect path-type.
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 1-20 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.
2414266	The fallocate(2) system call fails on Veritas File System (VxFS) file systems in the Linux environment.
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.
2982157	During internal testing, the f:vx_trancommit:4 debug asset was hit when the available transaction space is lesser than required.
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.
3071622	On SLES10, bcopy(3) with overlapping address does not work.
3340286	After a file system is resized, the tunable setting of dalloc_enable gets reset to a default value.
3096834	Intermittent vx_disable messages are displayed in the system log.
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
3233284	FSCK binary hangs while checking Reference Count Table (RCT).
3249958	When the /usr file system is mounted as a separate file system, Veritas File system (VxFS) fails to load.

Table 1-20 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.
3271892	Veritas File Replicator (VFR) jobs fail if the same ProcessID (PID) is associated with the multiple jobs working on different target filesystems.
3369020	The Veritas File System (VxFS) module fails to load in the RHEL 6 Update 4 environment.
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.
3313756	The file replication daemon exits unexpectedly and dumps coreon the target side.
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.
3331047	Memory leak occurs in the vx_followlink() function in error condition.
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.

Table 1-20 Veritas File System 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3335272	The mkfs (make file system) command dumps core when the log size provided is not aligned.
3337806	The find(1) command may panic the systems with Linux kernels with versions greater than 3.0.
3349649	The Oracle Disk Manager (ODM) module fails to load on RHEL6.5
3349651	VxFS modules fail to load on RHEL6.5
3350804	System panic on RHEL6 due to kernel stack overflow corruption.
3351937	Command vfradmin(1M) may fail while promoting job on locally mounted VxFS file system due to "relatime" mount option.
3352059	High memory usage occurs when VxFS uses Veritas File Replicator (VFR) on the target even when no jobs are running.
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.
3364303	Internal stress test on a locally mounted file system hits a debug assert in VxFS File Device Driver (FDD).
3364306	Stack overflow seen in extent allocation code path.
3383147	The C operator precedence error may occur while turning offdelayed allocation.
3394803	A panic is observed in VxFS routine vx_upgrade7() functionwhile running the vxupgrade command(1M).
3415639	The type of the fsdedupadm(1M) command always shows as MANUAL even it is launched by the fsdedupschd daemon.
3433786	The vxedquota(1M) command fails to set quota limits for some users.
3436699	An assert failure occurs because of a race condition between clone mount thread and directory removal thread while pushing data on clone.
3465035	The VRTSvxfs and VRTSfsadv packages displays incorrect "Provides" list.

Veritas File System: issues fixed in 6.0.4

Table 1-21 describes the incidents that are fixed in Veritas File System (VxFS) in 6.0.4.

Table 1-21 Veritas File System 6.0.4 fixed issues

Incident	Description
3312030	The default quota support on Veritas File System version 6.0.4 is changed to 32 bit.
3270210	On SUSE Linux Enterprise Server 11 (SLES 11) SP3, no support is provided for `VRTSfsadv`, `VRTSfssdk`, `VRTSvxfs` and `VRTSvxvs`.
3270200	Support for SLES 11 SP3 is added.
3268931	Support for SLES 11 SP3 is added.
3268916	Support for SLES 11 SP3 is added.
3268908	Support for SLES 11 SP3 is added.
3257467	Support for SLES 11 SP3 is added.
3248955	The system fails to build modules on SLES 11 SP3 for Veritas Oracle Disk Manager (ODM), because it cannot find the `utsrelease.h` file needed to build the ODM module for SP3 on SLES 11.
3247752	The system panics with the following message during the internal stress tests: `Unable to handle kernel paging request at <addr>`
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.
3149174	Veritas Oracle Disk Manager (ODM) clone shutdown fails with the `ORA-03113: end-of-file on communication channel` error.
3142045	With Oracle 12c version, the Veritas Oracle Disk Manager (ODM) library causes a version mismatch on the RHEL6 platform.
3140990	Requirement for the ability to turn off VxFS's invalidation of pages for some Network File System (NFS) workloads.
3121933	The `pwrite(2)` function fails with the EOPNOTSUPP error.
3101418	The current time returned by the operating system (Oracle error code `ORA-01513`) during Oracle startup is invalid.
3089210	The message `V-2-17: vx_iread_1 filesystem file system inode inode number marked bad incore` is displayed in the system log.

Table 1-21 Veritas File System 6.0.4 fixed issues *(continued)*

Incident	Description
3079215	Oracle RAC Database creation fails with the `Ora-00600 [ksfd_odmio1]` error when Veritas ODM links.
3068902	In case of stale NFS mounts, the `statfs()` function calls on non-VxFS file systems may cause df commands to hang.
3042485	During internal stress testing, the `f:vx_purge_nattr:1` assert fails.
2999493	The file system check validation fails after a successful full fsck during the internal testing with the message: `run_fsck : First full fsck pass failed, exiting`
2991880	In low memory conditions on a VxFS, certain file system activities may seem to be non-responsive.
2983248	The `vxrepquota`(1M) command dumps core.
2977828	The file system is marked bad after an inode table overflow error occurs.
2966277	Systems with high file system activity like read/write/open/lookup may panic the system.
2963763	When the `thin_friendly_alloc()` and `deliache_enable()` functionality is enabled, VxFS may enter a deadlock.
2926684	In rare cases, the system may panic while performing a logged write.
2908391	It takes a long time to remove checkpoints from the VxFS file system, when there are a large number of files present.
2907930	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.
2907923	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.
2907919	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.
2907908	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.
2885592	The `vxdump` operation is aborted on a file system which is compressed using the `vxcompress` command.
2839871	On a system with DELICACHE enabled, several file system operations may hang.

Table 1-21 Veritas File System 6.0.4 fixed issues *(continued)*

Incident	Description
2834192	You are unable to mount the file system after the full `fsck`(1M) utility is run.

Veritas File System: issues fixed in 6.0.3

Table 1-22 describes the incidents that are fixed in Veritas File System in 6.0.3.

Table 1-22 Veritas File System 6.0.3 fixed issues

Incident	Description
3004466	Installation of 5.1SP1RP3 fails on RHEL 6.3.
2895743	Accessing named attributes for some files seems to be slow.
2893551	When nfs connections are under load , file attribute information is replaced by question marks.
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.
2821152	Internal Stress test hit an assert “f:vx_dio_physio:4,1” on locally mounter file system.
2806466	fsadm -R resulting in panic at LVM layer due to vx_ts.ts_length set to 2GB.
2773383	Read/Write operation on a memory mapped files seems to be hung.
2641438	Modifications to user name space extended attributes are lost after a system reboot.
2624262	fsdedup.bin hit oops at vx_bc_do_brelse.
2611279	Filesystem with shared extents may panic.
2417858	VxFS quotas do not support 64 bit limits.

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 1-23 Veritas File System fixed issues

Incident	Description
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.
2682550	Access a VxFS file system via NFS could cause system panic on Linux while unmount is in progress.
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.
2597347	fsck should not coredump when only one of the device record has been corrupted and the replica is intact.
2584531	vxfs hangs on ls, du and find.
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.
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 1-24 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.
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.
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.
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.
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.

Table 1-24 Veritas Cluster Server 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3140359	Global Atomic Broadcast (GAB) fails to start when the gabconfig -c and gabconfig -cx commands are executed simultaneously on the same system.
3140961	The Asynchronous Monitoring Framework (AMF) driver panics the system when a Veritas File System (VxFS) file system is mounted or unmounted.
3140967	Veritas File System (VxFS) fails to unload.
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.
3177476	The AMF kernel driver causes the system to panic when an event is unregistered.
3207663	Incorrect user privileges are set in case of incorrect use of the '-group' option in command "hauser -addprive".
3211054	VCS support for Oracle Linux Unbreakable Enterprise KernelRelease 2 (UEK2) on Oracle Linux 6.
3211683	Mount agent does not detect "bind" file system mounts.
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.
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.

Table 1-24 Veritas Cluster Server 6.0.5 fixed issues *(continued)*

Fixed issues	Description
3302589	Veritas Cluster Server (VCS) support for SuSE LinuxEnterprise Server 11 Service Pack 3
3318335	VMwareDisks resource faults during data store migration of VM or during creation, deletion, or reversal of the VMsnapshot.
3318764	Unexpected deletion of temporary files causes the VCS agents to report an incorrect state.
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).
3354227	The clean entry point operation of the IPMultiNIC agent fails in the absence of a state file that contains current active device information.
3362108	The system panics if LLT receives a corrupt packet from the network.
3409593	The ASMDG agent shows offline before volumes are released and service group fail-over will be delayed because volumes won't stop.
3410926	The KVMGuest agent does not support virtual machine monitoring with the latest Red Hat Enterprise Virtualization (RHEV).

Veritas Cluster Server: Issues fixed in 6.0.4

This section describes Veritas Cluster Server fixed issues in 6.0.4.

LLT, GAB, and I/O fencing fixed issues in 6.0.4

Table 1-30 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-25 LLT, GAB, and I/O fencing fixed issues

Incident	Description
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.
3137520	Low Latency Transport (LLT) detects a duplicate node ID incorrectly, even if the nodes are using different ethernet SAP values.

Table 1-25 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
3302589	Veritas Cluster Server (VCS) support for SuSE Linux Enterprise Server 11 Service Pack 3.

Bundled agents fixed issues in 6.0.4

Table 1-31 lists the fixed issues for bundled agents.

Table 1-26 Bundled agents fixed issues

Incident	Description
3065929	VCS agent for VMwareDisks fails to detach or attach the VMware disks if the ESX host and the virtual machine are not in the same network.
3098114	Installation of VRTSvcsag package changes permission of `/opt/VRTSvcs/bin` directory to 744 from 755.
3125918	AMF driver panics the node when vxconfigd is unresponsive.
3028760	NFSRestart resource does not start NFS processes, such as statd and lockd, during online or offline operations.

VCS engine fixed issues in 6.0.4

Table 1-32 lists the fixed issues for VCS engine.

Table 1-27 VCS engine fixed issues

Incident	Description
3090710	The High Availability Daemon (HAD) starts and stops before the VxFEN driver configuration completes.
3112608	Resource fails to become online after switch operations fails for a service group.
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.
3042450	A parent service group which is frozen and configured with online local hard dependency is brought offline when its child service group faults.

AMF fixed issues

This section describes the issues fixed in AMF since the previous major release.

Fixed issues related AMF in 6.0.4

Table 1-28 lists the AMF fixed issues.

Table 1-28 AMF fixed issues

Incident	Description
3090229	AMF driver panics the node when vxconfigd is unresponsive.
3097342	AMF driver panics the node when AMF is being stopped.
3259682	If vxconfigd hangs, then registration thread of imfd trying to get diskgroup status from vxconfigd also hangs. Therefore, the `amfregister` command waiting for IMFD gets stuck.

Veritas Cluster Server: issues fixed in 6.0.3

Table 1-29describes the incidents that are fixed in Veritas Cluster Server in 6.0.3.

Table 1-29 Veritas Cluster Server 6.0.3 fixed issues

Incident	Description
3035052	Add support for special characters `.` and `/` in queue manager name, in the WebSphereMQ wizard.
3028989	Add support for the custom log path and data path for WebSphereMQ versions 7.0, 7.0.1 and 7.1 in a unified way.
3013962	Monitor fails to detect DB2 resource online for DB2 version 10.1.
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.
2964772	If you take an NFSRestart resource offline, the NFSRestart agent may unexpectedly stop NFS processes in a local container Zones.
2951467	AMF driver panics the machine when a VXFS filesystem is unmounted.
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

Table 1-29 Veritas Cluster Server 6.0.3 fixed issues *(continued)*

Incident	Description
2813773	AMF driver panics the box with the message "AMF ASSERT panic: FAILED(dev_name)"
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 1-30 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-30 LLT, GAB, and I/O fencing fixed issues

Incident	Description
2708619	If you set the scsi3_disk_policy attribute to dmp, you cannot enable the Veritas fencing module (VxFEN). The VxFEN source code is updated to pick up the dmp device path that contains the full disk name instead of a partition or slice.
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/lltab` file floods the system log file with large number of log messages.

Bundled agents fixed issues in 6.0.1

Table 1-31 lists the fixed issues for bundled agents.

Table 1-31 Bundled agents fixed issues

Incident	Description
2850904	Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and PanicSystemOnDGLoss is set to 0.

Table 1-31 Bundled agents fixed issues *(continued)*

Incident	Description
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.
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.
2679251	After stopping the VCS forcefully (hastop -local -force), if disk reservation resource is configured then user may observe a system panic.
2800021	Clean entry point fails due to discrepancy in command path on RHEL 5.
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 1-32 lists the fixed issues for VCS engine.

Table 1-32 VCS engine fixed issues

Incident	Description
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.

Table 1-32 VCS engine fixed issues *(continued)*

Incident	Description
2647049	VCS logs do not update the log time lines once the time-zone is changed. Thus, VCS keeps logging messages with old timings even after the time- zone is updated on the system.
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.
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 1-33 lists the fixed issues for enterprise agents.

Table 1-33 Enterprise agents fixed issues

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

Agent framework fixed issues in 6.0.1

Table 1-34 lists the fixed issues for agent framework.

Table 1-34 Agent framework fixed issues

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

Veritas 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 1-35 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 1-36 lists the issues fixed in 6.0.1.

Table 1-36 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.`

Table 1-36 Issues fixed in 6.0.1 *(continued)*

Incident	Description
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.
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.

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 1-37 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.

Storage Foundation for Databases (SFDB) tools: issues fixed in 6.0.4

Table 1-38 describes the incidents that are fixed in SFDB tools in 6.0.4.

Table 1-38 Storage Foundation for Databases (SFDB) tools 6.0.4 fixed issues

Incident	Description
3211391 (3211388)	During a Veritas Database Edition (DBED) instant checkpoint cloning, if you enable the Block Change Tracking feature, it results in an `ORA-00600` error.
3277940 (3290416)	Some DBED operations may fail with the following error message: `ORA-01406: fetched column value was truncated`
2937529	Support for the newer DB2 10.1 version of the database.

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

Table 1-39describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 6.0.3.

Table 1-39 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 1-40 describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-40 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 1-40 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 6.0.4

There are no Symantec Virtual Store fixed issues in 6.0.4.

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

Issues related to installation and upgrade

This section describes the known issues during installation and upgrade.

On RHEL 6 x86_64, system may panic during SFCFSHA rolling upgrade to 6.0.5 from 5.1SP1RPx, if frozen service groups exist [3471289]

During SFCFSHA rolling upgrade to 6.0.5 on RHEL 6 x86_64 from 5.1SP1RPx, if frozen online service groups depending on CVM exist, CVM may fail to be brought offline. Thereby system may crash when the installer stops vxglm due to a known issue in vxglm in 5.1SP1RPx on RHEL 6 x86_64.

Workaround:

Ensure all the service groups can be offline before the rolling upgrade.

If the system crashes, do the following:

1 Restart the system.

2 Ensure service groups are not in the **Frozen** state.

3 Perform the rolling upgrade again.

Rolling upgrade to 6.0.5 may fail if you set LC_ALL to POSIX on SLES (3437811)

Sybase is unable to start if `LC_ALL` is set to `POSIX`. The reason that the Sybase resources work fine before any upgrade is due to the different methods of starting VCS engine (HAD). VCS engine's (HAD) environment is inherited by the Agent and thus is passed to the ONLINE entry point of the Sybase resource. When you start HAD using `/etc/init.d/vcs start` and `LC_ALL` is set to `POSIX`, the Sybase resource doesn't start.

Workaround:

For the Sybase resources, add the following lines in the `SYBASE.sh` file. This file is present in directory specified in the attribute `$Home`.

File: `$Home/SYBASE.sh`

```
...
LC_ALL="C"
export LC_ALL
```

During VCS upgrades to 6.0.5, when you configure WebSphereMQ application with VCS wizards, error messages appear (3395997)

During VCS upgrades to 6.0.5, when you configure WebSphereMQ application with VCS wizards, you may see the following message:

```
Appropriate version of Package [VRTSmq6] is not installed. Install version
[5.1.14.0] or higher, and then retry.
```

The error message is displayed because the version of VRTSmq6 RPM installed is lower than 5.1.14.0. VRTSmq6 RPM is the WebSphereMQ agent RPM.

Workaround:

Install the latest VRTSmq6 RPM available at:

https://sort.symantec.com/agents

Note: The version of VRTSmq6 RPM should be 5.1.14.0 or later.

Error messages in syslog (1630188)

If you install or uninstall a product on a node, you may see the following warnings in syslog: /var/log/message. These warnings are harmless and can be ignored.

```
Jul  6 10:58:50 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:54 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:59 swlx62 setroubleshoot: SELinux is preventing the
restorecon from using potentially mislabeled files
```

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.

The uninstaller does not remove all scripts (2696033)

After removing SFHA, some of the RC scripts remain in the `/etc/rc*.d/` folder. This is due to an issue with the chkconfig rpm in RHEL6 and updates. You can manually remove the scripts from the `/etc/rc*.d/` folder after removing the VxVM packages.

Workaround:

Install the chkconfig-1.3.49.3-1 chkconfig rpm from the RedHat portal. Refer to the following links:

http://grokbase.com/t/centos/centos/117pfhe4zz/centos-6-0-chkconfig-strange-behavior

http://rhn.redhat.com/errata/RHBA-2012-0415.html

NetBackup 6.5 or older version is installed on a VxFS file system (2056282)

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to Veritas Storage Foundation (SF) 6.0.1, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/openv`), then while upgrading to SF 6.0.1, the installer fails to check if NetBackup is installed

on the same machine and uninstalls the shared infrastructure RPMs `VRTSpbx`, `VRTSat`, and `VRTSicsco`. This causes NetBackup to stop working.

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the `/usr/openv/netbackup/bin/version` file and `/usr/openv/netbackup/version` file to the `/tmp` directory. If you have clustered NetBackup installed, you must also copy the `/usr/openv/netbackup/bin/cluster/NBU_RSP` file to the `/tmp` directory. After you unmount the NetBackup file system, manually copy these two version files from `/tmp` to their original directories. If you have clustered NetBackup installed, you must also copy the `/usr/openv/netbackup/bin/cluster/NBU_RSP` file from `/tmp` to its original directory.

If the `version` files' directories do not exist, create the directories:

```
# mkdir -p /usr/openv/netbackup/bin
```

Run the installer to finish the upgrade process. After upgrade process completes, remove the two version files and their directories.

If your system is already affected by this issue, then you must manually install the `VRTSpbx`, `VRTSat`, and `VRTSicsco` RPMs after the upgrade process completes.

Ignore certain errors after an operating system upgrade—after a product upgrade with encapsulated boot disks (2030970)

Ignore certain errors after an operating system upgrade after a product upgrade with encapsulated boot disks.

You can ignore the following errors after you upgrade the operating system after a product upgrade that occurred with an encapsulated boot disk. Examples of the errors follow:

```
The partioning on disk /dev/sda is not readable by
The partioning tool parted, which is used to change the
partition table.
You can use the partitions on disk /dev/sda as they are.
You can format them and assign mount points to them, but you
cannot add, edit, resize, or remove partitions from that
disk with this tool.
```

Or

```
Root device: /dev/vx/dsk/bootdg/rootvol (mounted on / as reiserfs)
Module list: pilix mptspi qla2xxx silmage processor thermal fan
reiserfs aedd (xennet xenblk)
```

```
Kernel image; /boot/vmlinuz-2.6.16.60-0.54.5-smp
Initrd image: /boot/initrd-2.6.16.60-0.54.5-smp
```

The operating system upgrade is not failing. The error messages are harmless.

Workaround: Remove the /boot/vmlinuz.b4vxvm and /boot/initrd.b4vxvm files (from an un-encapsulated system) before the operating system upgrade.

Upgrading from Veritas Storage Foundation 5.1 Service Pack 1 Rolling Patch 2 to 6.0.1 with rootability enabled fails (2581313)

Upgrading from Veritas Storage Foundation (SF) 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 to 6.0.1 while using an encapsulated root disk fails because the post installation scripts of Veritas Volume Manager (VxVM) are unable to start the `initrd` daemon.

Workaround: To upgrade from 5.1 SP1 RP2 to 6.0.1 while using an encapsulated root disk, you must reinstall the nash utility on the system prior to the upgrade.

To upgrade from 5.1 SP1 RP2 to 6.0.1 while using an encapsulated root disk

1. Encapsulate the root disk.
2. Reinstall the nash utility.
3. Upgrade to the SF 6.0.1 release.

After finishing a kernel upgrade on a master node the cvm group on a slave node does not come online (2439439)

After successfully finishing a kernel upgrade on one node, the cvm group does not come online on the second node.

Workaround: Check that your cluster is not in a jeopardy state before you perform a rolling upgrade.

After upgrade from VxVM version 6.0, 6.0.1, or 6.0.3 with an encapsulated boot disk, the system fails to boot (2750782)

On Red Hat Enterprise Linux 6 (RHEL6), during the Veritas Volume Manager (VxVM) upgrade from 6.0, 6.0.1, or 6.0.3 to higher versions, the RPM runs the installation scripts of the VxVM higher version first.Then, the RPM runs the uninstallation scripts of the VxVM 6.0, 6.0.1, or 6.0.3 version. Due to a defect in the 6.0, 6.0.1, and 6.0.3 uninstallation scripts, it corrupts the file installed by the higher version. This leads to the boot failure.

Workaround:

1 Unroot the encapsulated root disk.

2 Uninstall `VRTSvxvm` (6.0, 6.0.1, or 6.0.3) RPM.

3 Install `VRTSvxvm` of higher version (above 6.0.3).

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
```

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.

Adding a node to a cluster fails if you did not set up passwordless ssh or rsh

Adding a node to a cluster fails if you did not set up passwordless `ssh` or `rsh` prior to running the `./installsfcfsha<version> -addnode` command.

Workaround: Set up passwordless `ssh` or `rsh`, and then run the `./installsfcfsha<version> -addnode` command.

Where <*version*> is the current release version.

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
```

Issue with soft links getting deleted in a manual upgrade

While performing a manual upgrade (from 5.1 to 6.0) of the VRTSvlic RPM, some of the soft links created during your previous installation are deleted. As a result, `vxkeyless` binary is not found in its specified path.

To prevent this, use the `--nopreun` option.

For example: `rpm -Uvh --nopreun VRTSvlic-3.02.61.003-0.x86_64.rpm`

While upgrading the VCS stack from a version prior to VCS 5.1, reconfiguration of MultiNICA IPv4RouteOptions attribute is required (2003864)

The 5.1SP1 MultiNICA agent now uses `ip` command by default. Due to behavioral differences in `ip` and `ifconfig` commands in regards to route configuration, MultiNICA flushes routes and sets them back for the new active device. If the MultiNICA resource configuration is not intended to make use of `ifconfig` command (see table below), you must configure IPv4RouteOptions attribute in MultiNICA resource definition.

Note: RouteOptions values are used by the `route` command where as the IPv4RouteOptions value is used by the `ip route` command. The values to be configured for these two attribute are very specific to their respective commands.

Table 1-41 Whether attributes are configured and required actions that you need to perform during upgrade

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
Configured	May or may not be configured	May or may not be configured	In this case the `ifconfig` command is used. If RouteOptions is set, attribute value is used to add/delete routes using command `route`. As the Options attribute is configured, IPv4RouteOptions values are ignored.	No need to configure IPv4RouteOptions .
Not configured	May or may not be configured	Must be configured	In this case the `ip` command is used. IPv4RouteOptions must be configured and are used to add/delete routes using the `ip route` command. As Options attribute is not configured, RouteOptions value is ignored.	Configure IPv4RouteOptions and set the IP of default gateway. The value of this attribute typically resembles: IPv4RouteOptions = “default via *gateway_ip*” For example: IPv4RouteOptions = “default via 192.168.1.1”

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]
   ```

SELinux error during installation of VRTSvcsag RPM

During the installation of VRTSvcsag RPM on RHEL 5 SELinux enabled machine, you may observe following SELinux error:

```
/usr/sbin/semodule: Failed on /opt/VRTSvcs/bin/selinux/vcsag.pp!
```

This error occurs due to improper installation of the SELinux package. As a result, SELinux commands may not function properly.

Workaround: Reinstall the SELinux package and relabel filesystem by either init or fixfiles method.

Web installer has no option to remove node from a cluster

Web Installer does not provide the option to remove node from a cluster.

Workaround: Manually remove nodes from a cluster. There is no option to remove nodes available from Web Installer or CPI.

CPI installer throws CPI WARNING V-9-40-4493 if enabling DMP root support when upgrading SF to 6.0.5 (3064919)

If you upgrade SF to 6.0.5 with the dmp_native_support option enabled, CPI installer may report the following message:

```
CPI WARNING V-9-40-4493 Failed to turn on
dmp_native_support tunable on <SERVER_NAME>. Refer to
Dynamic Multi-Pathing Administrator's guide to determine
the reason for the failure and take corrective action.
```

This message is inappropriate and should not be displayed. It does not affect the upgrading functions.

Workaround:

There is no workaround for this issue. You can safely ignore the message.

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.

Erroneous resstatechange trigger warning

You may encounter the following warning when you restart resources:

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

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

The Web installer hangs at the end of the rolling upgrade process (2792835)

At the end of a rolling upgrade, the Web installer completes all the processes successfully but does not show the completion page.

Workaround:

Even though you don't see a completion page, the upgrade process executes successfully. Refresh the browser to begin using it for other purposes.

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
   ```

Enabling or installing DMP for native support may not migrate LVM volumes to DMP (2737452)

On Linux System with LVM version 2.02.85, installing DMP or enabling dmp_native_support for DMP may not migrate LVM volumes to DMP. LVM Volume Groups may disappear.

From LVM version 2.02.85 onwards, device list is obtained from udev by default if LVM2 is compiled with UDEV support. This setting is managed using obtain_device_list_from_udev variable in /etc/lvm/lvm.conf. As DMP devices are not managed by UDEV, they will not be used by LVM. Thus LVM volumes are not migrated.

Workaround:

For LVM version 2.02.85 onwards, for DMP native support, always disable UDEV support for LVM by adding following line to /etc/lvm/lvm.conf in “devices” section:

```
obtain_device_list_from_udev = 0
```

Then install the package or enable dmp_native_support tunable. If dmp_native_support is already enabled then run the following command for applying changes:

```
# vxdmpadm settune dmp_native_support=on
```

Veritas Volume Manager known issues

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

Asymmetric Logical Unit Access (ALUA) LUNs assigned as VirtIO-SCSI devices not visible inside KVM guests (3341432)

Inside the KVM guest, during device discovery, Veritas Volume Manager executes RTPG IOCTL on disks to fetch the disk properties. The RTPG IOCTL fails on virtual devices backed by ALUA DMP nodes which is exported from the host using the VirtIO-SCSI interface. Therefore, the ASL fails to claim the disks inside the guest and the disks are not visible to Volume Manager. For example, if sdb is the DMPNODE-backed disk where DMPNODE belongs to ALUA enclosure and is exported to guest using virtio-scsi interface, the corresponding ALUA vendor ASL fails to claim the disk and the DDL-STATUS displays the following error:

```
# vxddladm list devices
DEVICE TARGET-ID STATE DDL-STATUS (ASL)
============================================================
```

```
sdb - Online ERROR (libvxxiv.so)
sda - Online CLAIMED (OTHER_DISKS)
```

Therefore the disk is not visible to volume manager.

```
# vxdisk list
DEVICE TYPE DISK GROUP STATUS
sda auto:none - - online invalid
```

Workaround: Export the underlying subpaths of DMPNODE to the guest using VirtIO-SCSI interface.

Enclosure attributes reset to default after root disk encapsulation (3188158)

After you encapsulate the root disk, the Enclosure attributes specified by the `vxtune setattr` command reset to their default values after restart. And after unroot, the Enclosure attributes show updated values.

Workaround:

If the changed attributes are requisite, unroot the disk encapsulation. If the changed attributes are not requisite, no action is required.

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.

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.

device.map must be up to date before doing root disk encapsulation (2202047)

If you perform root disk encapsulation while the `device.map` file is not up to date, the `vxdiskadm` command displays the following error:

```
VxVM vxencap INFO V-5-2-6098 Missing entry for root disk <rootdisk name>
in /boot/grub/device.map
```

Workaround:

Before you perform root disk encapsulation, run the the following command to regenerate the device.map file:

```
grub-install --recheck /dev/sdb
```

Required attributes for Veritas Volume Manager (VxVM) devices to avoid boot failures (1411526)

To support iSCSI devices, Veritas Volume Manager (VxVM) does not start non-root devices until runlevel2. The boot process expects all local (non-NFS) mount points in the `/etc/fstab` file to be present at boot time. To avoid boot failures, all VxVM entries in the `/etc/fstab` file must have the _netdev attribute, and must not have

the fsck required flag set. These attributes enable VxVM to defer mounting of VxVM devices until after VxVM has started.

vxrestored daemon fails to restore disabled paths (1663167)

The `vxrestored` daemon fails to restore disabled paths on RHEL 5 with direct attached disks.

Workaround:

Enable the `mpt_disable_hotplug_remove` tunable so that path level failover and failback function properly on RHEL 5 machines with direct attached disks.

To enable the mpt_disable_hotplug_remove tunable

1 Edit the /etc/modprobe.conffile and add the following line to the end of the file:

```
options mptsas mpt_disable_hotplug_remove=0
```

2 Rebuild the initrd image:

```
# mkinitrd -f /boot/initrd-`uname -r`.img `uname -r`
```

3 Reboot the system.

System hangs or panics after disabling 3 of 4 arrayside ports (1724260)

The system hangs or panics after you disable 3 of 4 arrayside ports.

Workaround:

This issue is fixed with a Novell patch for SLES 11 as indicated in Bugzilla ID 524347:

https://bugzilla.novell.com/show_bug.cgi?id=524347

Machine fails to boot after root disk encapsulation on servers with UEFI firmware (1842096)

Certain new servers in the market such as IBM x3650 M2, Dell PowerEdge T610, come with support for the UEFI firmware. UEFI supports booting from legacy MBR type disks with certain restrictions on the disk partitions. One of the restrictions is that each partition must not overlap with other partitions. During root disk encapsulation, it creates an overlapping partition that spans the public region of the root disk. If the check for overlapping partitions is not disabled from the UEFI firmware, then the machine fails to come up following the reboot initiated after running the commands to encapsulate the root disk.

Workaround:

The following workarounds have been tested and are recommended in a single-node environment.

For the IBM x3650 series servers, the UEFI firmware settings should be set to boot with the "Legacy Only" option.

For the Dell PowerEdge T610 system, set "Boot Mode" to "BIOS" from the "Boot Settings" menu.

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
   ```

Root disk encapsulation is not supported if the root disk (DMP node) has customized name or user-defined name (1603309)

Encapsulation of the root disk fails if it has been assigned a customized name with vxdmpadm(1M) command. If you want to encapsulate the root disk, make sure that you have not assigned a customized name to its corresponding DMP node.

See the vxdmpadm(1M) manual page.

See the "Setting customized names for DMP nodes" section of the *Veritas Storage Foundation Administrator's Guide*.

VxVM starts before OS device scan is done (1635274)

While working with some arrays, VxVM may start before all devices are scanned by the OS. This slow OS device discovery may result in malfunctioning of VM, fencing and VCS due to partial disks seen by VxVM.

Workaround:

After the fabric discovery is finished, issue the `vxdisk scandisks` command to bring newly discovered devices into the VxVM configuration.

Converting LVM volumes to VxVM volumes fails when multipathed storage devices are present (1471781, 1931727)

The `vxvmconvert` utility cannot convert LVM volumes to VxVM volumes when multipathed storage devices are present. This issue occurs with SLES 11 and RHEL5, due to changes in the LVM utilities. If multipathed devices are detected, the `vxvmconvert` utility exits with the following error:

```
vxvmconvert cannot convert multipathed devices on SLES11 systems.
... Exiting.
```

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.

During system boot, some VxVM volumes fail to mount (2622979)

During system boot, some VxVM volumes that exist in the `/etc/fstab` file fail to mount with the following error messages:

```
# fsck
Checking all file systems.
  error on stat() /dev/vx/dsk//volume: No such
file or directory
```

The load order of kernel modules in Linux results in the VxFS file system driver loading late in the boot process. Since the driver is not loaded when the `/etc/fstab` file is read by the operating system, file systems of the type vxfs will not mount.

Workaround:

To resolve the failure to mount VxFS file systems at boot, specify additional options in the `/etc/fstab` file. These options allow the filesystems to mount later in the boot process. An example of an entry for a VxFS file system:

```
/dev/vx/dsk/testdg/testvolume /mountpoint  vxfs  _netdev,hotplug  1 1
```

To resolve the issue, the fstab entry for VxVM data volumes should be as per following template:

```
/dev/vx/dsk/testdg/testvol    /testmnt    vxfs   _netdev     0 0
```

Removing an array node from an IBM Storwize V7000 storage system also removes the controller (2816589)

When using an IBM Storwize V7000 storage system, after removing one array node, the corresponding controller is also removed.

Workaround: The following procedure resolves this issue.

To resolve this issue

1 Set the `iotimeout` tunable to 600:

```
# vxdmpadm setattr enclosure encl1 recoveryoption=throttle \
 iotimeout=600
```

2 After you re-add the SAN VC node, run the `vxdctl enable` command for Dynamic Multi-Pathing (DMP) to detect the added paths:

```
# vxdctl enable
```

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.

System panic occurs on RHEL6.0 when VxVM is installed on EMC PowerPath managed devices (2573229)

System panic occurs on RHEL6.0 due to page cache issue when VxVM is installed on top of EMC PowerPath managed devices.

Workaround :

Install VxVM before installing and configuring EMC PowerPath

Hardware paths for operating system paths have changed in DMP 6.0 (2410716)

In DMP 6.0, the hardware paths for operating system paths have changed. After upgrading to DMP 6.0, path attributes are reset to the default values. You must reconfigure any path-level attributes that were defined in the /etc/vx/dmppolicy.info file.

Workaround:

To configure path-level attributes

1 Remove the path entries from the `/etc/vx/dmppolicy.info` file.

2 Reset the path attributes.

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.

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
```

cvm_clus resource goes into faulted state after the resource is manually panicked and rebooted in a 32 node cluster (2278894)

The `cvm_clus` resource goes into faulted state after the resource is manually panicked and rebooted in a 32 node cluster.

Workaround: There is no workaround for this issue.

DMP disables subpaths and initiates failover when an iSCSI link is failed and recovered within 5 seconds. (2100039)

When using iSCSI S/W initiator with an EMC CLARiiON array, iSCSI connection errors may cause DMP to disable subpaths and initiate failover. This situation occurs when an iSCSI link is failed and recovered within 5 seconds.

Workaround:

When using iSCSI S/W initiator with an EMC CLARiiON array, set the node.session.timeo.replacement_timeout iSCSI tunable value to 40 secs or higher.

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.

The SCSI registration keys are not removed even if you stop VCS engine for the second time (3037620)

If you stop VCS engine for the first time, the SCSI registration keys can be removed. But if you stop VCS engine for the second time, the keys are not removed.

Workaround:

There is no workaround for this issue.

Using `vxsnap` to addmir for a volume is very slow and the plex is detached after the operation (3037712)

If you use vxsnap to addmir for a volume on RHEL6.3, the process is very slow and the plex becomes detached after the operation.

Workaround:

There is no workaround for this issue.

Upgrading VxVM package (in-place upgrade) on an encapsulated boot disk does not work on SLES11 (3061521)

On SLES11, after the in-place upgrade on root encapsulated system, the following message is displayed on reboot:

```
Waiting for device /dev/vx/dsk/bootdg/rootvol to appear:
..............could not find /dev/vx/dsk/bootdg/rootvol.
Want me to fall back to <resume device>?(Y/n)
```

Workaround:

For an SLES11 machine with encapsulated root disk, perform the following steps to upgrade VxVM.

To upgrade VxVM:

1. Unroot the encapsulated root disk:

   ```
   # /etc/vx/bin/vxunroot
   ```

2. Upgrade the VxVM.

   You can perform this using the CPI installer or using the `rpm` command.

   ```
   # rpm -Uvh VRTSvxvm-version.rpm
   ```

3 Reboot the system.

4 Re-encapsulate the root disk:

```
# /etc/vx/bin/vxencap -c -g root_diskgroup rootdisk=root_disk
```

Issue with Veritas Volume Manager and Red Hat Enterprise Linux 6 Update 4 (3137543)

Starting with Red Hat Enterprise Linux (RHEL) 6 Update 2 and later updates, the operating system includes the trusted boot software. When selecting a full installation and not the standard installation, the trusted boot package gets installed on the system. This modifies the `grub/menu.list` file, which breaks the root disk encapsulation feature from Veritas Volume Manager (VxVM). This change impacts the root disk encapsulated systems, which might prevent booting from the encapsulated root disk.

Workaround:

You can resolve this issue by using either of the following workarounds:

- Ensure that the trusted boot software is removed from the system before rebooting the system after upgrading to the RHEL 6 Update 4 operating system.
- Unencapsulate the system before upgrading to the RHEL 6 Update 4 operating system.

For more information, see note 3 of the following TechNote:

http://www.symantec.com/docs/TECH203099

After encapsulation of the root disk, the system may fail to boot up with the "Unrecognized device string" error. (3326269)

If the `/boot/device.map` file contains multiple `hd` number entries pointing to the root disk, it causes the `vxencap` script to add incorrect `hd` number expressions in the `vxvm_root` entry and therefore, the GRUB (Grand Unified Bootloader) fails to identify the root device.

Workaround:

Remove duplicate entries pointing to the root disk from the `/boot/device.map` file before encapsulation, and make sure the correct entry is pointing to the root disk.

OR

Manually correct the root(hdx,y) in the `vxvm_root` entry in the `menu.lst` file immediately after the encapsulation and before a reboot.

Nodes with encapsulated root disks fail to reboot after you upgrade VxVM from version 6.0.3 to 6.0.4. (3300580)

On an SUSE Linux Enterprise Server (SLES) 11 system with encapsulated root disks, a reboot after upgrading Veritas Volume Manager (VxVM) from 6.0.3 to 6.0.4 prompts you to fall back to the resume device. The prompt message is similar to following:

```
could not find /dev/vx/dsk/bootdg/rootvol, want me to fall back to
/dev/sda2? (Y/N)
```

Workaround:

To resolve this issue

1. Select **Y** on the prompt to continue booting up.
2. Backup the existing `/boot/VxVM_initrd.img` file.
3. Regenerate the `VxVM_initrd.img` file using the following command:

```
# /etc/vx/bin/vxinitrd  /boot/VxVM_initrd.img `uname -r`
```

Nodes with encapsulated root disks fail to restart after you perform an operating system upgrade for SLES11SPx on 6.0.1 systems. (2612301)

Upgrading the kernel or operating system on an encapsulated boot disk does not work on SUSE Linux Enterprise Server (SLES) 11.

Note: This issue though fixed in release 6.0.1 on Red Hat Enterprise Linux (RHEL) 5, 6 and SLES 10. It occurs on SLES 11 because of a program issue.

Workaround:

Perform the following procedure on the system with the encapsulated root disk to upgrade the kernel or operating system.

To upgrade the kernel or operating system on a system with an encapsulated root disk

1. Unroot the encapsulated root disk:

   ```
   # /etc/vx/bin/vxunroot
   ```

2. Upgrade the kernel or the operating system.

   You may upgrade the kernel using the following rpm command:

   ```
   # rpm -Uvh Kernel-upgrade_version
   ```

   OR

   You may upgrade the operating system using tools like YaST or Zypper.

3. Reboot the system.
4. Re-encapsulate the root disk:

   ```
   # /etc/vx/bin/vxencap -c -g root_diskgroup rootdisk=root_disk
   ```

Veritas Cluster Server known issues

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

Stale NFS file handle on the client across failover of a VCS service group containing LVMLogicalVolume resource (2016627)

A VCS service group for a LVM volume group will be online automatically after a failover. However, the client applications may fail or be interrupted by stale NFS file handle error.

Workaround: To avoid the stale NFS file handle on the client across service group failover, specify "fsid=" in the Options attribute for Share resources.

Operational issues for VCS

LVMLogicalVolume online entry point stops responding and times out for mirrored volumes on SLES10 [2077294]

`LVMLogicalVolume` uses `lvchange` command to activate the logical volumes. In case of mirrored volumes, the `lvchange` command itself stops responding when it is invoked through a script. This causes online entry point to time out and the online entry point of `LVMLogicalVolume` resource stops responding. This is an issue with the SLES10.

LVM SG transition fails in all paths disabled status [2081430]

If you have disabled all the paths to the disks, the `LVM2 vg` commands stop responding and wait until at least one path to the disks is restored. As `LVMVolumeGroup` agent uses `LVM2` commands, this behavior causes online and offline entry points of `LVMVolumeGroup` agent to time out and `clean EP` stops responding for an indefinite time. Because of this, the service group cannot fail over to another node.

Workaround: You need to restore at least one path.

SG goes into Partial state if Native LVMVG is imported and activated outside VCS control

If you import and activate LVM volume group before starting VCS, the `LVMVolumeGroup` remains offline though the `LVMLogicalVolume` resource comes online. This causes the service group to be in a partial state.

Workaround: You must bring the VCS `LVMVolumeGroup` resource offline manually, or deactivate it and export the volume group before starting 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.

Switching service group with DiskGroup resource causes reservation conflict with UseFence set to SCSI3 and powerpath environment set [2749136]

If UseFence is set to SCSI3 and powerpath environment is set, then switching the service group with DiskGroup resource may cause following messages to appear in syslog:

```
reservation conflict
```

This is not a SFHA issue. In case UseFence is set to SCSI3, the diskgroups are imported with the reservation. This message gets logged while releasing and reserving the disk.

Workaround: See the tech note available at http://www.symantec.com/business/support/index?page=content&id=TECH171786.

Issues related to the VCS engine

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.

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

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

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

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

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

Workaround: No workaround.

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

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

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

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

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

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>
```

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.

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

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

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

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

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

Kernel may panic when you stop the HAD process with the force flag on a node and then reboot the node (3387018)

When you execute the following steps on any of the nodes in a cluster, kernel may panic on any of the nodes in the cluster:

1 Stop the HAD process with the force flag with the following command:

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

2 Reboot or shut down the node where HAD ran.

Workaround:

No workaround available.

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.

Issues related to the bundled agents

KVMGuest resources may fault due to RHEV-M takes longer time to detect the node fault and initiate manual fencing [3473541]

In case of node crash, RHEV-M takes longer time to detect it and change the host state as `Non-Responsive`. Even after it detects the crash and changes the host state, sometimes, initiation of RHEV-M manual fencing may be stuck. As a result, VCS KVMGuest resources are timed out and then fault. In such situations, VCS KVMGuest agents are not able to fence out the node. A Red Hat Enterprise Virtualization limitation causes this issue.

Workaround:

Manually initiate fencing through RHEV-M by clicking **Confirm Host has been rebooted**.

KVMGuest agent fails to communicate with “Internal” RHEV-M domains in RHEV 3.0 environments (3421010)

KVMGuest agent uses REST APIs to communicate with Red Hat Enterprise Virtualization Manager. The default domain is set to internal, which is a local domain, while configuring RHEV-M. In RHEV 3.0 environment, the REST APIs fail to

communicate with RHEV-M using internal domain. Red Hat has fixed this issue in RHEV 3.1.

Workaround: Symantec recommends use of RHEV 3.1 and later versions. For more information, see the Red Hat Enterprise Virtualization documentation.

Manual configuration of RHEVMInfo attribute of KVMGuest agent requires all its keys to be configured (3277994)

The RHEVMInfo attribute of KVMGuest agent has 6 keys associated with it. When you edit main.cf to configure RHEVMInfo attribute manually, you must make sure that all the keys of this attribute are configured in main.cf. If any of its keys is left unconfigured, the key gets deleted from the attribute and the agent does not receive the complete attribute. Hence, it logs a Perl error `Use of uninitialized value` in the engine log. This is due to the VCS engine behavior of handling the attribute with key-value pair.

Workaround: Use the high availability (ha) commands to add or modify the RHEVMInfo attribute of KVMGuest resource.

System panics after starting KVM virtualized guest or initiating KVMGuest resource online [2337626]

System panics when the KVM guest is started or when the KVMGuest resource online is initiated. This issue is rarely observed.

The issue is observed due to the file descriptor leak in the libvirtd process. The maximum file open limit of file descriptor for libvirtd process is 1024. You may sometimes observe that more than 1024 file descriptors are opened when the KVM guest is started. Therefore, if the maximum file open limit is crossed, any attempt to start the KVM guest or to open a new file causes the system to panic. VCS cannot control this behavior as it suspects a file descriptor leak in the libvirtd process.

Workaround: There is no definite resolution for this issue; however, you can check the number of files opened by the libvirtd process in `/proc/<pid of libvirtd>/fd/`. If the file count exceeds 1000, restart libvirtd with the following command:

```
/etc/init.d/libvirtd restart
```

KVMGuest monitor entry point reports resource ONLINE even for corrupted guest or with no OS installed inside guest [2394235]

The VCS KVMGuest monitor entry point reports resource state as ONLINE in spite of the operating system inside the guest being corrupted or even if no operating system is installed inside the guest. The VCS KVMGuest agent uses `virsh` utility to determine the state of the guest. When the guest is started, the `virsh` utility

reports the state of the running guest as running. Based on this running state, VCS KVMGuest agent monitor entry point reports the resource state as ONLINE.

In case the operating system is not installed inside the guest or the installed operating system is corrupted, `virsh` utility still reports the guest state as running. Thus, VCS also reports the resource state as ONLINE. Since RedHat KVM does not provide the state of the operating system inside guest, VCS cannot detect the guest state based on the state of the operating system.

Workaround: No workaround for this known issue.

LVM logical volume may get stuck with reiserfs file system on SLES11 [2120133]

LVM logical volume may get stuck with reiserfs file system on SLES11 if the service group containing the logical volume is switched continuously between the cluster node.

This issue may be observed:

- During the continuous switching of the service group having the LVM logical volume with reiserfs file system.
- On SLES11 and with reiserfs file system only.
- Due to the behavior of device-mapper on SLES11.

However, the issue is not consistent. Sometimes, the device-mapper gets stuck while handling the logical volumes and causes the logical volume to hang. In such a case, LVM2 commands also fail to clear the logical volume. VCS cannot handle this situation as the LVM2 commands are unable to deactivate the hung logical volume.

Resolution: You must restart the system on which the logical volumes are stuck in this situation.

KVMGuest resource comes online on failover target node when started manually [2394048]

The VCS KVMGuest resource comes online on failover target node when VM guest started manually, even though the resource is online on the primary node.

Red Hat kernel-based virtual machine (KVM) allows you to start the guest using same guest image on multiple nodes. The guest image is residing on the cluster file system. If the guest image is stored on the cluster file system, then it becomes available on all the cluster nodes simultaneously.

If the KVMGuest resource of VCS has made the guest online on one node by starting it using the guest image on cluster file system and if you manually start the same guest on the other node, Red Hat KVM does not prevent you from doing so. However, as this particular guest is under VCS control, VCS does not allow the

resource to be ONLINE on multiple nodes simultaneously (unless it is in parallel service group configuration). VCS detects this concurrency violation and brings down the guest on the second node.

Note: This issue is also observed with CVM raw volume.

Workaround: No workaround required in VCS. VCS concurrent violation mechanism handles this scenario appropriately.

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.

DiskReservation agent may call clean if you configure large number of its resources in a single service group [2336391]

DiskReservation agent may call clean if you configure large number of DiskReservation resource (more than 400 resources) in a single service group and try to offline the service group.

In a single service group configuration with more than 400 DiskReservation resources and equal number of Mount resources, the service group offline may cause the DiskReservation agent to call clean entry point. This issue is not observed if you configure about 150 resources.

Workaround: No workaround.

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.

DiskGroup agent is unable to offline the resource if volume is unmounted outside VCS

DiskGroup agent is unable to offline the resource if volume is unmounted using the `umount -l` command outside VCS.

A service group contains DiskGroup, Volume and Mount resources and this service group is online. Volume is mounted by Mount resource with VxFSMountLock enabled. An attempt to manually unmount the volume using `umount -l` system command causes the mount point to go away; however, the file system lock remains as it is. The volume cannot be stopped as it is mount locked and hence the disk group cannot be imported. This causes the disk group resource to go into UNABLE to OFFLINE state. Also, any attempt to again mount the file system fails, because it is already mount locked. This issue is due to file system behavior on Linux.

Workaround: Do not use `umount -l` command to unmount the VxFS file system when the mount lock is enabled. Instead, first unlock the mount point using the`/opt/VRTS/bin/fsadm` command and then unmount the file system.

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.

VVR setup with FireDrill in CVM environment may fail with CFSMount Errors [2564411]

When you try to bring the FireDrill service group online through Java Console or `hagrp -online` command, the CFSMount resource goes into faulted state.

Workaround: Run the `fsck` command.. You can find these commands in the engine logs.

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.

RVGsnapshot agent does not work with volume sets created using vxvset [2553505]

RVGsnapshot agent does not work with volume sets created using `vxvset`. This happens during FireDrill in a VVR einviroment.

Workaround: No workaround.

No log messages in engine_A.log if VCS does not find the Monitor program [2563080]

No message is logged in the engine_A.log, when VCS cannot find the Monitor program with KVM guest with service group online.

Workaround: In case resource state is unknown , also refer to agent log files for messages.

No IPv6 support for NFS [2022174]

IPv6 is not supported for NFS.

Workaround: No workaround.

Guest virtual machine may fail on RHEL 6.1 if KVM guest image resides on CVM-CFS [2659944]

If a KVM guest image file resides on CVM-CFS, the migration of that guest virtual machine may fail with "Permission Denied" error on RHEL 6.1. This causes guest virtual machine to go in "shut-off" state on both source and destination node, and the associated VCS KVMGuest.

Workaround: Make sure that the guest image file is having 777 permission.

KVMGuest agent fails to communicate with RHEV-M domain as “Internal” (2738491)

KVMGuest agent uses REST APIs to communicate with Red Hat Enterprise Virtualization Manager. The default domain that is set while configuring the RHEV-M is **internal**, which is a local domain. The REST APIs fail to communicate with RHEV-M using internal domain.

Workaround: Add an additional valid domain using `rhevm -manage -domains` command. Red Hat has provided steps to use this command to add a valid domain for REST API communication. Please refer Red Hat Enterprise Virtualization documentation for more details.

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.

Concurrency violation observed when host is moved to maintenance mode [2735283]

When a Red Hat Enterprise Virtualization host running a virtual machine is moved to maintenance state, the virtual machine migration is initiated by RHEV. SFHA detects the migration according to virtual machine state, such as "migrating". Due to timing issue RHEV Manager occasionally sends the virtual machine state as "up" even if the migration is in progress. Due to this state, the resource is marked ONLINE on the node to which it migrates and may cause concurrency violation.

Workaround: No workaround.

Concurrency violation observed during migration of monitored virtual machine [2755936]

If a VCS service group has more than one KVMGuest resource monitoring virtual machine and one of the virtual machines is migrated to another host, a service group level concurrency violation occurs as the service group state goes into PARTIAL state on multiple nodes.

Workaround: Configure only one KVMGuest resource in a Service group.

KVMGuest agent fails to recognize paused state of the VM causing KVMGuest resource to fault [2796538]

In a SUSE KVM environment, when a virtual machine is saved, its state is changed to paused and then shut-off. The paused state remains for a very short period of time, due to timing in case that the KVMGuest agent misses this state. Then the resource state will be returned as OFFLINE instead of INTENTIONAL OFFLINE, which causes the KVMGuest resource to fault and failover.

This is due to the limitation of SUSE KVM as it does not provide a separate state for such events.

Workaround: No workaround.

Resource does not appear ONLINE immediately after VM appears online after a restart [2735917]

During a VM restart the resource does not come ONLINE immediately after the VM starts running. As the VM state is 'Reboot in Progress' it reports INTENTIONAL OFFLINE and after VM is UP the resource cannot immediately detect it as the next monitor is scheduled after 300 seconds.

Workaround: Reduce the OfflineMonitorInterval and set it to suitable value.

Logical volume resources fail to detect connectivity loss with storage when all paths are disabled in KVM guest [2871891]

In a KVM environment if all storage paths are disabled, then LVMLogicalVolume and LVMVolumeGroup resources fails to detect the loss of connectivity with storage. This is because even if the storage paths are disabled, native LVM commands return success, which causes VCS to report the resource state ONLINE and hence no SG failover is initiated. If any application monitored by VCS is doing a read/write I/O to this volumes, then it can detect the fault and initiate service group failover.

Workaround: No workaround.

LVM Logical Volume will be auto activated during I/O path failure [2140342]

LVM Logical Volume gets auto activated during the I/O path failure. This causes the VCS agent to report "Concurrency Violation" errors, and make the resource groups offline/online temporarily. This is due to the behavior of Native LVM.

Workaround: Enable the LVM Tagging option to avoid this issue.

DiskGroup agent does not fail over the service group if storage connectivity is lost and I/O fencing is configured (3408712)

If storage connectivity is lost, DiskGroup agent initiates a service group failover, provided I/O fencing is configured. If DiskGroup resource Reservation attribute is set to "ClusterDefault", then the DiskGroup resource clean entry point returns a failure, as it fails to deport the disk group. As a result, it keeps on attempting to deport the diskgroup and service group failover is not initiated.

Workaround: Set Reservation attribute to "SCSI3". If the Reservation attribute is set to "SCSI3", clean entry point returns success even if it fails to deport the diskgroup and service groups fails over to another node.

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
```

Independent Persistent disk setting is not preserved during failover of virtual disks in VMware environment [3338702]

VMwareDisks agent supports Persistent disks only. Hence, Independent disk settings are not preserved during failover of virtual disk.

Workaround: No workaround.

System having DiskReservation resource online panics with the loss of storage connectivity (3321322)

If the storage connectivity is lost on the system where DiskReservation resource is online, that system panics.

Workaround: No workaround.

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

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

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

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.

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.

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 global clusters

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

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

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

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

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

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

LLT known issues

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

LLT connections are not formed when a vlan is configured on a NIC (2484856)

LLT connections are not formed when a vlan is configured on a NIC that is already used to configure an LLT link.

Workaround: Do not specify the MAC address of a NIC in the `lltab` file while configuring LLT if you want to configure a vlan later. If you have already specified the MAC address of a NIC, then delete the MAC address from the `lltab` file, and update the file before you restart LLT.

LLT may fail to detect when bonded NICs come up (2604437)

When LLT is configured over a bonded NIC and that bonded NIC is DOWN with the `ifconfig` command, LLT marks the corresponding link down. When the bonded NIC is UP again using the `ifconfig` command, LLT fails to detect this change and it doesn't mark the link UP.

Workaround: Close all the ports and restart LLT, then open the ports again.

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.

LLT may incorrectly declare port-level connection for nodes in large cluster configurations [1810217]

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This occurs in some corner cases even though a port is not even registered on the peer node.

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.

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.

I/O fencing known issues

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

VCS fails to take virtual machines offline while restarting a physical host in RHEV and KVM environments (3320988)

In RHEV and KVM environments, the virtualization daemons `vdsmd` and `libvirtd`required to operate virtual machines are stopped before VCS is stopped during a reboot of the physical host. In this scenario, VCS cannot take the virtual machine resource offline and therefore the resource fails to stop. As a result , LLT,

GAB and fencing fail to stop. However, the virtual network bridge is removed leading to the loss of cluster interconnects and causing a split-brain situation.

Workaround: If the virtual network bridge is not assigned to any virtual machine, remove the virtual bridge and configure LLT to use the physical interface. Alternatively, before initiating a reboot of the physical host, stop VCS by issuing the `hastop -local` command. The `-evacuate` option can be used to evacuate the virtual machines to another physical host.

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.

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 RPM 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.

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.

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 RPM is not installed on the system (2858190)

If VRTSvxfen RPM 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 RPM installed on the system you cannot run the utility from the install media.

Workaround: Install VRTSvxfen RPM, 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.

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.

Veritas Cluster Server agents for Veritas Volume Replicator known issues

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in 6.0.5 release.

fdsetup cannot correctly parse disk names containing characters such as "-" (1949294)

The fdsetup cannot correctly parse disk names containing characters such as "-".

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

Stale entries are found in sample main.cf file for RVGLogowner agent and RVGPrimary agent.

The stale entries are present in the main.cf.seattle and main.cf.london files on the RVGLogowner agent which includes CFSQlogckd resource. However, CFSQlogckd is not supported since VCS 5.0.

On RVGPrimary agent, the stale entries are present in file `main.cf.seattle` and `main.cf.london` and the stale entry includes the DetailMonitor attribute.

Workaround

1 For main.cf.seattle for RVGLogowner agent in the cvm group:

- Remove the following lines.

```
CFSQlogckd qlogckd (
             Critical = 0
              )

cvm_clus requires cvm_vxconfigd
        qlogckd requires cvm_clus
        vxfsckd requires qlogckd


        // resource dependency tree
        //
        //      group cvm
        //      {
        //      CFSfsckd vxfsckd
        //          {
        //          CFSQlogckd qlogckd
        //              {
        //              CVMCluster cvm_clus
        //                  {
```

```
//                  CVMVxconfigd cvm_vxconfigd
//                  }
//              }
//          }
//      }
```

- Replace the above lines with the following:

```
cvm_clus requires cvm_vxconfigd
        vxfsckd requires cvm_clus


        // resource dependency tree
        //
        //      group cvm
        //      {
        //      CFSfsckd vxfsckd
        //          {
        //           CVMCluster cvm_clus
        //              {
        //               CVMVxconfigd cvm_vxconfigd
        //              }
        //          }
        //      }
```

2 For main.cf.london for RVGLogowner in the cvm group:

- Remove the following lines

```
CFSQlogckd qlogckd (
             Critical = 0
              )

cvm_clus requires cvm_vxconfigd
        qlogckd requires cvm_clus
        vxfsckd requires qlogckd


        // resource dependency tree
        //
        //      group cvm
        //      {
        //      CFSfsckd vxfsckd
        //          {
```

```
//          CFSQlogckd qlogckd
//              {
//              CVMCluster cvm_clus
//                  {
//                  CVMVxconfigd cvm_vxconfigd
//                  }
//              }
//          }
//      }
```

- Replace the above lines with the following:

```
cvm_clus requires cvm_vxconfigd
        vxfsckd requires cvm_clus


        // resource dependency tree
        //
        //      group cvm
        //      {
        //      CFSfsckd vxfsckd
        //          {
        //           CVMCluster cvm_clus
        //              {
        //               CVMVxconfigd cvm_vxconfigd
        //              }
        //          }
        //      }
```

3 For main.cf.seattle for RVGPrimary agent in the cvm group:

- In the group ORAGrp and for the Oracle resource database, remove the line: `DetailMonitor = 1`

4 For main.cf.london for RVGPrimary agent in the cvm group:

- In the group ORAGrp and for the Oracle resource database, remove the line: `DetailMonitor = 1`

Issues related to Intelligent Monitoring Framework (IMF)

Registration error while creating a Firedrill setup [2564350]

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

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

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

Workaround: No workaround.

Perl errors seen while using haimfconfig command

Perl errors seen while using `haimfconfig` command:

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

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

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.

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

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

Workaround: Use the correct commands:

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

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).

Cluster Manager (Java Console) may display an error while loading templates (1433844)

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

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

Workaround: Ignore the error.

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.

Issues related to virtualization

Cluster communication breaks when you revert a snapshot in VMware environment (3409586)

If VCS is running on the guest operating system when a VMware virtual machine snapshot is taken, the virtual machine snapshot contains the run-time state of the cluster. When you restore the snapshot, the state of the cluster which is restored can be inconsistent with other nodes of the cluster. Due to the inconsistent state, VCS is unable to communicate with other nodes of the cluster.

Workaround: Before you take a snapshot of the virtual machine, stop the VCS services running inside the virtual machine.

Host fails to reboot when the resource gets stuck in ONLINE|STATE UNKNOWN state [2738864]

In a Red Hat Enterprise Virtualization environment, if a host reboot is performed on which the KVMGuest resource monitoring the virtual machine is ONLINE, then the host reboot fails. This is because the VDSM is getting stopped before VCS could shutdown the virtual machine. In this case, the virtual machine state remains ONLINE|STATE UNKNOWN, and hence VCS stop fails eventually failing the host reboot as well.

Workaround: Switch the service group to other node before initiating a host reboot.

VM state is in PAUSED state when storage domain is inactive [2747163]

If the storage domain associated with the running virtual machine becomes inactive, the virtual machine may go to **paused** state.

Workaround: Make sure that the storage domain is always active when running virtual machine.

Switching KVMGuest resource fails due to inadequate swap space on the other host [2753936]

Virtual machine fails to start on a host if the host does not have the sufficient swap space available.

Workaround: Please make sure that each host has sufficient swap space available for starting the virtual machine.

Policies introduced in SLES 11SP2 may block graceful shutdown if a VM in SUSE KVM environment [2792889]

In a SUSE KVM environment, virtual machine running SLES11 SP2 inside may block the virtual machine graceful shutdown request due to some policies introduced in SLES 11SP2. SUSE recommends turning off the policy with `polkit-gnome-authorization` for a virtual machine.

Workaround: Make sure that all the policies blocking any such request are turned off.

Load on `libvirtd` may terminate it in SUSE KVM environment [2824952]

In a SUSE KVM environment, occasionally `libvirtd` process may get terminated and `/etc/init.d/libvirtd status` command displays:

```
#/etc/init.d/libvirtd  status
Checking status of libvirtd                    dead
```

This may be due to heavy load on `libvirtd` process.

Workaound: Restart the `libvirtd` process and run:

```
# service libvirtd stop
# service libvirtd start
```

SFHA may detect the migration event during the regular monitor cycle due to the timing issue [2827227]

In a virtualization environment, SFHA detects the virtual machine migration initiated outside SFHA and changes the state accordingly. However, occasionally, SFHA may miss the migration event due to timing issue and detect the migration during the regular monitor cycle. For example, if you set OfflineMonitorInterval as 300sec,

it takes up to 5 minutes for SFHA to report ONLINE on the node where the virtual machine got migrated.

Workaround: No workaround available.

Increased memory usage observed even with no VM running [2734970]

Increased memory usage was observed on the hosts even when VMs were either not running or had stopped. This is due to the RHEV behavior.

Workaround: No workaround.

Resource faults when it fails to ONLINE VM beacuse of insufficient swap percentage [2827214]

In virtualization environment, if VCS fails to start the virtual machine due to unavailability of required virtualization resources such as CPU, memory, or disks, the resource goes into FAULTED state.

Workaround: Make sure that the required virtualization resources are always available in a virtualization environment.

Offline or switch of KVMGuest resource fails if the VM it is monitoring is undefined [2796817]

In a SUSE KVM environment, if a running virtual machine is undefined using `virsh undefine` command, an attempt to offline or switch the KVM guest resource monitoring that VM fails because the agent is not able to get the information from the KVM hypervisor.

Workaround: To undefine the VM on a particular node, first switch the service group containing the KVMGuest resource to another node and then undefine the VM on the first node.

Migration of guest VM on native LVM volume may cause libvirtd process to terminate abruptly (2582716)

When the guest VM image is on native LVM volume, then the migration of that guest initiated by the administrator may cause `libvirtd` process to terminate abruptly.

Workaround: Start the `libvird` process manually.

VMware virtual environment specific issues

vCenter Integrated Installer may fail to install VCS on OL 6.x guests [2905806]

In a VMware environment, if you create a virtual machine by setting **Guest Operating System Version** to **Oracle Linux 4/5/6 (64-bit)**, you cannot use the

vSphere Client to install Veritas Cluster Server (VCS) on the virtual machine. The installation wizard is unable to distinguish between OL 5.x and OL 6.x guests and therefore does not support installation on OL 6.x guests.

Workaround: When you create the virtual machine, set the **Guest Operating System Version** to **RedHat Enterprise Linux 6 (64-bit)** and then use vCenter Integrated Installer. Alternatively, install VCS from the command line, using one of the installer-based options. For more information, see the *Veritas Cluster Server Installation Guide*.

Oracle configuration wizard may not allow to add or remove listeners [2868770]

While configuring an Oracle application, the configuration wizard fails to accept new listeners or to remove the existing listeners if you navigate back from the virtual IPs page to the listener page.

Workaround: No workaround.

Unable to configure application when a new LSI Logic SCSI controller is added dynamically [2937623]

In case of SLES operating system, it is possible that hot plug drivers do not get loaded and hence the operating system is unable to recognize the newly added SCSI controllers.

Workaround: Load the drivers into the kernel before attempting to configure application or add a new SCSI controller. You can add the drivers by using the command:

```
# modprobe acpiphp
# modprobe pci_hotplug
```

During snapshot reversal, SFHA commands fail, and an unrelated error message may appear [2939177]

If a virtual machine is under SFHA control, and you revert to the virtual machine configuration snapshot, the SFHA configuration file moves to the read- only mode. During this time, if you run a SFHA command, the command fails. Also, the following unrelated error message appears in the log entries:

```
Failed to delete the attribute key <value of the key>
 of attribute DiskPaths in the VCS configuration with error 11335
(Configuration must be ReadWrite : Use haconf -makerw)
```

Workaround: No workaround. This message is safe to ignore.

Application instances are ignored if its dependent storage is configured for other instance [2966123]

Application instances are ignored if its dependent storage is already configure under VCS for other instances. The following message is logged in healthview_A.log when you try to configure a new application instance.

```
Skipping the <Application> instance: <instance_name>,
because the storage is already configured or an application
running under VCS control.
```

Workaround: Reconfigure all the instances that use common storage in a single run of the wizard.

Symantec High Availability tab does not report LVMVolumeGroup resources as online [2909417]

The Symantec High Availability tab does not automatically report the online status of activated LVMVolumeGroup resources in the following case:

- If you created the VCS cluster as part of the Symantec High Availability Configuration Wizard workflow.

Workaround: Start the LVMVolumeGroup resources from the Symantec High Availability tab. For more information, see the Symantec High Availability Solutions Guide for VMware.

vSphere UI may not show applications and their status [2938892]

The vSphere UI may show that it cannot connect to the virtual machine and may prompt for username and password of the virtual machine. Even after entering correct username and password, vSphere UI may not show the list of applications monitored by VCS. This can happen if the VOM MH configuration has authorization problems.

Workaround: Restart the xprtld daemon on the virtual machine with the following commands:

```
/etc/init.d/xprtld stop
/etc/init.d/xprtld start
```

SSO configuration fails if the system name contains non-English locale characters [2910613]

If you install the Symantec High Availability guest components on a system that has non-English locale characters in its name, then the SSO configuration between such a system and the Symantec High Availability Console host fails. (2910613)

Workaround: Ensure that the system name does not contain any non-English locale characters.

VMwareDisks agent may fail to start or storage discovery may fail if SELinux is running in enforcing mode [2854053]

The VMwareDisks agent and discFinder binaries refer to the libvmwarevcs.so shared library. SELinux security checks prevent discFinder from loading the libvmwarevcs.so library, which requires text relocation. If SELinux is running in enforcing mode, these two executables may report the "`Permission denied`" error and may fail to execute.

Workaround: Enter the following command and relax the security check enforcement on the Symantec libvmwarevcs.so library:

```
# chcon -t textrel_shlib_t '/opt/VRTSvcs/lib/libvmwarevcs.so'
```

During virtual machine migration, the migrating node might panic in a Red Hat Enterprise Virtualization Hypervisor environment (3106238)

During virtual machine migration, a node might take more than 16 seconds to send packets to the second node. If the `peerinact` tunable has the default value of 16 seconds, then this occurrence can cause either of the nodes to panic in a Red Hat Enterprise Virtualization Hypervisor (RHEV-H) environment because of a split brain scenario.

The LLT `peerinact` tunable specifies how long LLT will wait to receive a packet before marking the link of a peer node as "inactive." Once a link is marked as inactive, LLT will not send any data on that link.

Workaround:

Use the `lltconfig -T peerinact:`*`value`* command to set the `peerinact` tunable to a higher value.

To set the peerinact tunable

1. Determine how long the migrating node is unresponsive in your environment.
2. If that time is less than the default LLT `peerinact` value of 16 seconds, then this panic should not occur. Otherwise, set the value to greater than the time that the migrating node is unresponsive:

   ```
   # lltconfig -T peerinact:2000
   ```

   The `peerinact` tunable is in centiseconds. This example sets the value to 20 seconds.

3 Verify that `peerinact` has been set to the specified time:

```
# lltconfig -T query
Current LLT timer values (.01 sec units):
 heartbeat = 50
 heartbeatlo = 100
 peertrouble = 200
 peerinact = 2000
 oos = 10
 retrans = 10
 service = 100
 arp = 30000
 arpreq = 3000
 Current LLT flow control values (in packets):
 lowwater = 40
```

4 Repeat step 2 and step 3 on all nodes in the cluster, and then restart the migration.

5 After the migration completes, use the `lltconfig` command on all nodes in the cluster to reset the `peerinact` tunable to the default value:

```
# lltconfig -T peerinact:1600
```

Optionally, if you want to retain the new `peerinact` tunable value, you can make the value persistent across reboots by appending the following line at the end of `/etc/llttab` file:

```
set-timer peerinact:2000
```

This example sets the value to 20 seconds.

The vCenter Server tasks shown in the vSphere Client may fail to display the correct status if the Symantec High Availability Guest Components installation fails (2824900)

This issue occurs in case of VMware vCenter Server version 5.1.

If you choose to install the Symantec High Availability Guest Components using the vSphere Client menu and if the installation fails, then the vCenter Server task shown in the vSphere Client gets held-up at 10%. It however, does not display that the installation has failed.

Workaround:

Refer to the installer logs at the following locations on the Symantec High Availability Console server. Rectify the cause and re-run the wizard.

```
% ALLUSERSPROFILE %\Symantec\ApplicationHA\Logs\ApplicationHA.log

% ALLUSERSPROFILE %\Symantec\ApplicationHA\Logs\VII\
```

You can also refer to logs on the system at the following location:

```
(1) /opt/INSTALLER/vii_logs
```

Alternatively, choose to install the guest components using the product installer or the CLI. For details refer to the product-specific installation and upgrade guide.

Veritas Dynamic Multi-pathing known issues

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

DMP marks the subpaths as DISABLED while these subpaths are accessible from OS level (2037222)

For iSCSI devices on SLES 10 SP3, the DMP tunable parameter dmp_fast_recovery needs to be turned off.

```
# vxdmpadm settune dmp_fast_recovery=off
```

DMP panics if a DDL device discovery is initiated immediately after loss of connectivity to the storage (2040929)

When using EMC Powerpath with VxVM 5.1SP1 on SLES11, set the fast_io_fail_tmo on the HBA port to any non-zero value that is less than the dev_loss_tmo value so as to avoid a panic in case a DDL device discovery is initiated by the `vxdisk scandisks` command or the `vxdctl enable` command immediately after loss of connectivity to the storage.

Upgrading the Linux kernel when the root volume is under DMP control (2080909)

This section includes the procedures for upgrading the Linux kernel when the root volume is under DMP control.

Linux kernel can be upgraded on RHEL5 systems without turning off the DMP native support. Only one reboot is required to bring system LVM volume on DMP after kernel upgrade.

To update the kernel on a RHEL5 system

1 Update kernel with the rpm command.

```
# rpm -ivh kernel_rpm
```

2 Turn on the dmp_native_support tunable:

```
# vxdmpadm settune dmp_native_support=on
```

This enables booting with new kernel with LVM devices with DMP.

3 Reboot.

On SLES10 or SLES11

On SLES, the kernel can not be upgraded in a single reboot due to limitation in mkinitrd command.

To update the kernel on a SLES10 or SLES11 system

1 Turn off DMP native support

```
# vxdmpadm settune dmp_native_support=off
```

2 Reboot the system.

3 Upgrade kernel using the rpm command

```
# rpm -ivh kernel_rpm
```

4 Turn on DMP native support.

```
# vxdmpadm settune dmp_native_support=on
```

5 Reboot the system to bring the root LVM volume under DMP control.

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.

After rebooting the array controller for CX4-240-APF array, I/O errors occur on shared file systems (2616315)

For Linux hosts, rebooting the array controller for a CX4-240-APF array may result in I/O errors on shared file systems.

Workaround:

To work around this issue

- Set the tunable parameter dmp_lun_retry_timeout to 120 seconds before rebooting the array controller.

```
# vxdmpadm settune dmp_lun_retry_timeout=120
```

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
```

DMP path is disabled under I/O load on SuSE10 SP4 and 3PAR array (3060347)

When running I/O load, some DMP paths are disabled. Error messages like the following are displayed in `/var/log/message`:

```
sd 1:0:0:14: SCSI error: return code = 0x08070002
sdp: Current: sense key: Aborted Command
     ASC=0x4b <<vendor>> ASCQ=0xc4
end_request: I/O error, dev sdp, sector 10633904
qla2xxx 0000:0b:00.0: scsi(1:0:14)
Dropped frame(s) detected (0x80000 of 0x80000 bytes), \
firmware reported underrun.
qla2xxx 0000:0b:00.0: scsi(1:0:14) \
FCP command status: 0x15-0x302 (0x70002) \
portid=060100 oxid=0x29d ser=0x1048e \
cdb=2a0000 len=0x80000 rsp_info=0x8 resid=0x0 fw_resid=0x80000
VxVM vxdmp V-5-0-112 [Info] disabled path 8/0xf0 \
belonging to the dmpnode 201/0xf0 due to...
```

Workaround:

Use the newer firmware version of disk array and HBA card.

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 1-42 shows the Hitachi arrays that have new array names.

Table 1-42 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

Turning off the DMP native support does not reset the preferred_names field in lvm.conf to the original values (2421823)

When you turn off the native support, the preferred_names field in lvm.conf is not reset to the original value. LVM does not function correctly with Device Mapper Volumes.

Workaround: Manually edit the lvm.conf file, and then Run `vgscan` command

Dynamic multi-pathing (DMP) does not co-exist with `scsi_dh_emc` module on SLES11 (3466160)

On SLES11 SP2 and SP3 computer with EMC CLARiiON AP/F storage, dynamic multi-pathing (DMP) does not support coexistence with `scsi_dh_emc` module, as it disrupts the operations of DMP. You may check if `scsi_dh_emc` module is loaded with following command:

```
# lsmod | grep scsi_dh_emc
```

Workaround: Use the following procedure to remove the `scsi_dh_emc` module.

To remove `scsi_dh_emc` module

1 Stop `multipathd`.

```
#/etc/init.d/multipathd stop
```

```
# chkconfig multipathd off
```

2 Unload the `scsi_dh_emc` module.

```
# rmmod scsi_dh_emc
```

3 Regenerate the initrd file.

```
# mkinitrd
```

Automatic Storage Management (ASM) disks fail to display "asm" tag [3380535]

Veritas Volume Manager does not dislplay the "asm" tag for disks that are used by the ASM disk group if an entire disk or LUN is used as an ASM disk.

Workaround: Use a disk partition instead of an entire disk for disks that are used by the ASM disk group.

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

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

Installer exits upgrade to 5.1 RP1 with Rolling Upgrade error message (1951825, 1997914)

Installer exits upgrade to 5.1 RP1 with Rolling Upgrade error message, if protocol version entries are present in `/etc/gabtab` and `/etc/vxfenmode` files. Installer program may exit with either one of the following error messages during upgrade from 5.1 to 5.1 RP1:

```
SF51 is installed. Rolling upgrade is only supported from 5.1 to
higher version for the products
```

Or

```
To do rolling upgrade, VCS must be running on <node>.
```

Workaround: If the protocol version entries are present in `/etc/gabtab` and `/etc/vxfenmode` files, then installer detects it as Rolling Upgrade (RU). If you are not attempting RU, and doing full upgrade to 5.1 RP1, remove the protocol version entries from these two files for installer to proceed with regular upgrade.

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.

Process start-up may hang during configuration using the installer (1678116)

After you have installed a Storage Foundation product, some Veritas Volume Manager processes may hang during the configuration phase.

Workaround: Kill the installation program, and rerun the configuration.

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.

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.

Veritas File System known issues

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

Full fsck cannot repair the file system [3451617]

The `full fsck` command cannot repair the file system.

Workaround:

Run the following `full fsck` twice to repair the file system:

```
# fsck -o full -y device_name
```

Some inode operations on RHEL 6 systems may cause stack overflow error (3412667)

Some inode update operations on RHEL 6 systems may lead to stack overflow errors thereby causing the systems to panic.

Workaround:

No workaround is available.

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.

Task blocked messages display in the console for RHEL6 (2560357)

On RHEL6, the kernel occasionally displays messages in the console similar to the following example:

```
INFO: task seq:16957 blocked for more than 120 seconds.
```

These messages display because the task is blocked for a long time on sleep locks. However, the task is not hung and the messages can be safely ignored.

Workaround: You can disable these messages by using the following command:

```
# echo 0 > /proc/sys/kernel/hung_task_timeout_secs
```

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.

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.

System unable to select ext4 from the file system (2691654)

The system is unable to select ext4 from the file system.

Workaround: There is no workaround.

A low memory machine with RHEL 6.2 or higher may panic with General Protection Fault (3046544)

If a machine with low memory (less than 3GB) is installed with RHEL 6.2 or higher system, it may hit the General Protection Fault assert with the following stack trace:

```
machine_kexec at ffffffff8103284b
crash_kexec at ffffffff810ba972
oops_end at ffffffff81501860
die at ffffffff8100f26b
do_general_protection at ffffffff815013f2
general_protection at ffffffff81500bc5
shrink_page_list.clone.0 at ffffffff8112db97
shrink_inactive_list at ffffffff8112e10c
shrink_zone at ffffffff8112ee8f
balance_pgdat at ffffffff811303d9
kswapd at ffffffff81130606
kthread at ffffffff81091df6
kernel_thread at ffffffff8100c14a
```

Workaround:

There is no workaround for this issue.

The replication operation fails (3010202)

On SLES 10 SP3 or SLES 11 SP2, the replication operation fails with the following message:

```
btree.c     1827:       ASSERT(0) failed
```

This issue occurs if you use the `bcopy` function because it is deprecated in these releases.

Workaround:

You can turn off the shared extent optimization on these machines. Contact Symantec support to get the exact commands.

The de-duplication operation may seem to be hung (2753687)

After the de-duplication operation completes, some of the temporary files are not cleaned. Hence the shutdown script fails to kill all the processes and the de-duplication operation may hang.

Workaround:

Use the `kill` command to manually kill the de-duplication operation.

The `fsdedupadm` command does not show the full name of the machine (3015478)

The length of the name for a node is limited to 15 characters. So whenever the name of a machine exceeds this length, it shows only the first 15 characters of name.

Workaround:

There is no workaround for this issue.

Incorrect option for VxFS file system in /etc/fstab (3059189)

When you create a VxFS file system, VOM sets incorrect option 'suid' for it in the `/etc/fstab` file. This would make the operating system unable to startup normally after reboot:

```
# NOTE: When adding or modifying VxFS or VxVM entries, add '_netdev'
# to the mount options to ensure the filesystems are mounted after
# VxVM and VxFS have started.
/dev/vx/dsk/dg3/dg3vol2 /dg3vol2 vxfs suid 0 2
```

Workaround:

Before reboot, manually edit the `/etc/fstab` file and change 'suid' to '_netdev'.

Rolling upgrade from version 6.0.3 may cause corruption in the Veritas File System external quota file (2933571)

The 6.0.3 release supported 64-bit quota, which allowed quota limits higher than 1 terabyte. However, there is a possibility of corruption during rolling upgrade when two nodes may be on different patch levels—one node using 64-bit quota limit and the other node using 32-bit quota limit.

Workaround: Disable quotas before you start rolling upgrade. Back up the external quota file. This will help to restore the original file in the event that the quota file becomes corrupted during the upgrade.

For more information, see the technote:
http://www.symantec.com/docs/TECH211178

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.

RVG monitor script may display command not found messages (1709034)

On VCS hosts with VVR resources configured, the following error message displayed in engine_A.log indicates a script error:

```
/opt/VRTSvcs/bin/RVG/monitor: line 124: {print $6}: command not found
/opt/VRTSvcs/bin/RVG/monitor: line 124: {print $6}: command not found
/opt/VRTSvcs/bin/RVG/monitor: line 124: {print $6}: command not found
```

This may fail online/monitor the bunker RVG resources, when they are configured.

Workaround:

Manually edit the following files to update the script:

```
/opt/VRTSvcs/bin/RVG/monitor
/opt/VRTSvcs/bin/RVG/online
/opt/VRTSvcs/bin/RVG/offline
```

In each file, modify the following line:

```
sys=`LC_ALL=C; export LC_ALL; $hasys -nodeid | $awk '{print $6}'`
```

to

```
sys=`LC_ALL=C; export LC_ALL; $hasys -nodeid | awk '{print $6}'`
```

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 -t 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 -t vxfs /dev/vx/dsk/dg/data_volume
```

Running SUSE Linux and using Novell's YaST tool to configure an IPv6 address may result in an error (1679261)

When Novell's YaST tool is invoked to configure an IPv6 address on a different network interface and if:

- the host name, the DNS server name and domain name are specified to the YaST tool.
- IPv6 address is assigned by the Dynamic Host Configuration Protocol (DHCP).
- the "Write Hostname to /etc/hosts" option is selected (this is selected by default).

This results in the `vradmin` command returning the following error:

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

This happens because the YaST tool can replace the `/etc/hosts` entry containing `127.0.0.2` from the IPv4 host name to the specified new IPv6 host name. For example:

```
127.0.0.2 v6hostname.space.ipv6.com v6hostname
```

Workaround:

The following procedure resolves this issue.

To resolve this issue

1 Edit the `/etc/hosts` file to specify the correct IPv6 address.

2 Restart the `vradmind` daemon on all VVR hosts:

```
# /etc/init.d/vras-vradmind.sh restart
```

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 restart
```

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 restart
```

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.

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.

vradmin functionality may not work after a master switch operation (2163712)

In certain situations, if you switch the master role, `vradmin` functionality may not work. The following message displays:

```
VxVM VVR vxrlink ERROR V-5-1-15861  Command is not supported for
command shipping. Operation must be executed on master
```

Workaround:

To restore vradmin functionality after a master switch operation

1 Restart `vradmind` on all cluster nodes. Enter the following:

```
# /etc/init.d/vras-vradmind.sh restart
```

2 Re-enter the command that failed.

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 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.

Clone fails with error "ORA-01513: invalid current time returned by operating system" with Oracle 11.2.0.3 (2804452)

While creating a clone database using any of the point-in-time copy services such as Flashsnap, SOS, Storage Checkpoint, or Filesnap, the clone fails. This problem appears to affect Oracle versions 11.2.0.2 as well as 11.2.0.3.

You might encounter an Oracle error such as the following:

```
/opt/VRTSdbed/bin/vxsfadm -s flashsnap -o clone
-a oracle -r dblxx64-16-v1 --flashsnap_name TEST11 --clone_path
/tmp/testRecoverdb --clone_name clone1
USERNAME:  oragrid
STDOUT:
Retrieving snapshot information ...                              Done
Importing snapshot diskgroups ...                                Done
Mounting snapshot volumes ...                                    Done

ORA-01513: invalid current time returned by operating system
```

This is a known Oracle bug documented in the following Oracle bug IDs:

- Bug 14102418: DATABASE DOESNT START DUE TO ORA-1513
- Bug 14036835: SEEING ORA-01513 INTERMITTENTLY

Workaround:

Retry the cloning operation until it succeeds.

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.

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.

Clone command errors in a Data Guard environment using the MEMORY_TARGET feature for Oracle 11g (1824713)

The `dbed_vmclonedb` command displays errors when attempting to take a clone on a STANDBY database in a dataguard environment when you are using the MEMORY_TARGET feature for Oracle 11g.

When you attempt to take a clone of a STANDBY database, the `dbed_vmclonedb` displays the following error messages:

```
Retrieving snapshot information ...                          Done
Importing snapshot diskgroups ...                            Done
Mounting snapshot volumes ...                                Done
Preparing parameter file for clone database ...              Done
Mounting clone database ...
ORA-00845: MEMORY_TARGET not supported on this system


SFDB vxsfadm ERROR V-81-0612 Script
/opt/VRTSdbed/applications/oracle/flashsnap/pre_preclone.pl failed.
```

This is Oracle 11g-specific issue known regarding the MEMORY_TARGET feature, and the issue has existed since the Oracle 11gr1 release. The MEMORY_TARGET feature requires the `/dev/shm` file system to be mounted and to have at least 1,660,944,384 bytes of available space. The issue occurs if the `/dev/shm`file system is not mounted or if the file system is mounted but has available space that is less than the required minimum size.

Workaround:

To avoid the issue, remount the `/dev/shm` file system with sufficient available space.

To remount the /dev/shm file system with sufficient available space

1. Shut down the database.
2. Unmount the `/dev/shm` file system:

```
# umount /dev/shm
```

3. Mount the `/dev/shm` file system with the following options:

```
# mount -t tmpfs shmfs -o size=4096m /dev/shm
```

4. Start the database.

Offline mode Storage Checkpoint or FlashSnap does not confirm the offline status of the database in CFS environment, leading to clone failure (2869260)

In a cluster file system for Single Instance Oracle, if an offline snapshot or Storage Checkpoint, and clone is created on the node where the database is inactive, then the cloning would fail with an error similar to SFDB vxsfadm ERROR V-81-0564 Oracle returned error.

```
... Reason: ORA-01194: file 1 needs more recovery to be consistent
ORA-01110: data file 1: /var/tmp/ikWxDkQ1Fe/data/sfaedb/system01.dbf'
(DBD ERROR: OCIStmtExecute) ...
```

Workaround: There is no workaround for this. In case of a Single Instance database installed on a cluster file system, create the Storage Checkpoint or snapshot on the active node.

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.

vxdbd fails to start after upgrade from 6.0.1 to 6.0.3 MR on linux (2969173)

The `vxdbd` daemon fails to start after manual upgrade from 6.0.1 to 6.0.5 Maintenance Release (MR) on Linux platform.

Workaround:

Use the `--nopreun` option for the `rpm` upgrade command to start up `vxdbd` properly after manual upgrade.

For example:

```
$ rpm -Uvh --nopreun VRTSdbed
```

The dbdst_show_fs(1M) command may fail with a Perl warning message (3263177)

The `dbdst_show_fs`(1M) command may fail with the following Perl warning message:

```
oracle@testbox:~> dbdst_show_fs -S $ORACLE_SID -m /snap_data11r2 -o volume
perl: warning: Setting locale failed.
perl: warning: Please check that your locale settings:
```

```
        LANGUAGE = (unset),
        LC_ALL = "",
        LANG = "en_US.UTF-8"
    are supported and installed on your system.
perl: warning: Falling back to the standard locale ("C").
SFORA dbdst_show_fs ERROR V-81-6209 Repository Empty.
```

Note: This issue is observed only in Linux SLES 10.

Workaround: Use the following command for the default locale settings and retry:

```
oracle@testbox:~> export LC_ALL=C
```

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.

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.

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.

Workaround:

For SFCFS:

- You can specify the `fsid share` option during `cfsshare share` to force the `fsid` of an NFS-exported VxFS Storage Checkpoint to remain the same on all cluster nodes.

  For example:

  To NFS-export a VxFS Storage Checkpoint of a VxFS file system that has already been added to VCS configuration and mounted at `/ckpt1`, run the following command:

  ```
  # cfsshare share /ckpt1 "fsid=num"
  ```

  where *num* is any 32-bit number that is unique amongst all the exported file systems.

  See the `exports`(5) manual page for more information.

For SFHA:

- You can modify the Options attribute of the Share resource corresponding to the VxFS Storage Checkpoint and add the `fsid` share option to force the `fsid` of an NFS-exported VxFS Storage Checkpoint to remain the same on all cluster nodes.

  See the `exports`(5) manual page for more information.

  For example:

  If `share1` is the VCS Share resource corresponding to an NFS-exported VxFS Storage Checkpoint, then run the following commands:

  ```
  # haconf -makerw
  # hares -modify share1 Options "fsid=num"
  # haconf -dump -makero
  ```

  where *num* is any 32-bit number that is unique amongst all the exported filesystems.

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.

NFS resource might not come online while configuring CNFS share (2488685)

If SELinux is configured as `enforcing` or `permissive`, NFS resource might not come online and go into `FAULTED` state while configuring CNFS share `cfsnfssg` service group.

Sample output:

```
# hastatus -sum

-- SYSTEM STATE
-- System               State                Frozen

A  swlx14               RUNNING              0

-- GROUP STATE
-- Group           System   Probed     AutoDisabled    State

B  cfsnfssg        swlx14   Y          N               OFFLINE|FAULTED
B  cfsnfssg_dummy  swlx14   Y          N               OFFLINE
B  cvm             swlx14   Y          N               ONLINE
B  vip1            swlx14   Y          N               OFFLINE

-- RESOURCES FAILED
-- Group           Type                 Resource             System

D  cfsnfssg        NFS                  nfs                  swlx14
```

Workaround

To resolve this issue you need to add the Ethernet port into the trusted list for SELinux.

- In the System Setup->Firewall configuration, select customize.
- In the Trusted device, select the Ethernet port.

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.

tail -f run on a cluster file system file only works correctly on the local node (2613030)

When using the `tail -f` command to monitor a file on a cluster file system, changes to the file made on remote nodes are not detected. This is due to the `tail` command now utilizing `inotify`. Symantec is currently unable to support `inotify` with a cluster file system due to GPL restrictions.

Workaround: To revert to the old behavior, you can specify the `---disable-inotify` option with the `tail` command.

"Configuration must be ReadWrite : Use haconf -makerw" error message appears in VCS engine log when hastop -local is invoked (2609137)

A message similar to the following example appears in the `/var/VRTSvcs/log/engine_A.log` log file when you run the `hastop -local` command on any system in a SFHA cluster that has `CFSMount`resources:

```
2011/11/15 19:09:57 VCS ERROR V-16-1-11335 Configuration must be
ReadWrite : Use haconf -makerw
```

The `hastop -local` command successfully runs and you can ignore the error message.

Workaround: There is no workaround for this issue.

Issues observed with force unmounting a parent cluster file system mount before unmounting a nested child VxFS or cluster file system mount (2621803)

When you have nested mounts in which a secondary VxFS file system is mounted in the name space of the primary file system in the cluster, if the primary file system gets force umounted before unmounting the secondary, then unmounting the secondary at a later time can cause unpredictable issues.

Workaround: There is no workaround for this issue.

Full file system check takes over a week (2628207)

On a large file system with many Storage Checkpoints, a full file system check using the`fsck_vxfs`(1M) command might appear to be hung. The `fsck` command is not actually hung; the process can take an extremely long time to complete.

Workaround: There is no workaround for this issue.

Performance degradation seen on a CFS filesystem while reading from a large directory (2644485)

Performance degradation is seen on a CFS filesystem while reading from a large directory.

Workaround: There is no workaround.

Some ODM operations may fail with "ODM ERROR V-41-4-1-328-22 Invalid argument" (3323866)

On systems having heavy database activity using ODM some operations may fail with the following error:

```
ODM ERROR V-41-4-1-328-22 Invalid argument.
```

Workaround:

Retry the operation

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.

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

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

Node fails to join the SFHA cluster if the file system containing Oracle Clusterware is not mounted (2611055)

The sequence number of the startup script for Oracle High Availability Services daemon (ohasd) is lower than some of the SFHA components such as VXFEN and VCS. During system startup, if the file system containing Oracle Clusterware does not get mounted before the ohasd startup script is executed, the script continuously waits for the file system to become available. As a result, the other scripts (including those of SFHA components) are not executed and the node being started does not join the SFHA cluster.

Workaround: If the rebooted node does not join the SFHA cluster, the cluster can be started manually using the following command:

```
# installsfrac -start node1 node2
```

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.

Rolling upgrade not supported for upgrades from SF Oracle RAC 5.1 SP1 with fencing configured in `dmp`mode.

Rolling upgrade is not supported if you are upgrading from SF Oracle RAC 5.1 SP1 with fencing configured in `dmp`mode. This is because fencing fails to start after the system reboots during an operating system upgrade prior to upgrading SF Oracle RAC.

The following message is displayed:

```
VxVM V-0-0-0 Received message has a different protocol version
```

Workaround: Perform a full upgrade if you are upgrading from SF Oracle RAC 5.1 SP1 with fencing configured in `dmp`mode.

Veritas Volume Manager can not identify Oracle Automatic Storage Management (ASM) disks (2771637)

Veritas Volume Manager (VxVM) commands can not identify disks that are initialized by ASM. Administrators must use caution when using the VxVM commands to avoid accidental overwriting of the ASM disk data.

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 cluster node may panic (2524087)

On SLES 10 SP4, the cluster node may panic while iSCSI initiator access the LUN from the target.

Workaround

There is no workaround at this time.

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.

The Symantec Quick Clone Virtual Machine Wizard may not behave as expected when multiple instances are open (2309702)

The wizard may not behave as expected, if you invoke multiple parallel session of the wizard from a single vSphere Client at the same time.

For example, if you do the following:

- Right-click wingoldvm1 and invoke the wizard.
- Then soon after, right-click slesgoldvm1 and invoke the wizard.

This causes you to have two instances of the wizard running from the same vSphere Client and can cause unexpected behavior.

Workaround

To resolve this issue:

- Close both instances of the wizard.
- Reopen a new instance of the wizard.

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.

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.

Known issues common to VCS and SFCFSHA

This section describes the known issues in this release common to Veritas Cluster Server (VCS) and Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

Network File System (NFS) lock failover is not supported on RHEL6 Update 3 and Update 4, and SLES 11 SP1, SP2 and SP3 (3331646)

If a file is locked from a NFS client, the other client may also get the lock on the same file after failover of the NFS share service group.

Workaround:

No workaround available.

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 kickstart.
- 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).

Support for SmartSync with database mounted on raw volumes [3416016]

The SmartSync feature with the database configured on raw volumes depends on support from the database vendor. If supported by the database vendor, the SmartSync feature uses an extended interface between VxVM volumes and the database software to avoid unnecessary work during mirror resynchronization.

Verify with the database vendor that the database software supports the SmartSync feature.

The SmartSync feature is supported on all platforms when the database uses VxFS file systems mounted on Veritas Volume Manager volumes, through the Veritas Extension for Oracle Disk Manager (VRTSodm) interface.

Oracle RAC 10g Release 2 instances in the 6.0.5 release documents [3470952]

SF Oracle RAC supports Oracle RAC versions 11g Release 2 and later.

You may disregard references to earlier Oracle RAC versions in the installation guide and release notes documents of this release.