

Veritas Storage Foundation™ and High Availability Solutions 6.0.3 Release Notes - Linux

6.0.3 Maintenance Release

Veritas Storage Foundation and High Availability Solutions Release Notes 6.0.3

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

Product version: 6.0.3

Document version: 6.0.3 Rev 12

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 us regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apj@symantec.com

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

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

Documentation

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

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

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Contents

Technical Support	3
Chapter 1	
About Veritas Storage Foundation and High Availability Solutions	9
Introduction	10
About the installmr script	10
The installmr script options	10
Overview of the installation and upgrade process	15
Changes introduced in 6.0.3	16
Changes introduced in Veritas Cluster Server 6.0.3	16
Changes introduced in SFDB 6.0.3	17
System requirements	17
Supported Linux operating systems	17
VMware Environment	20
Database requirements	20
Veritas Storage Foundation memory requirements	21
Disk space requirements	21
Number of nodes supported	21
List of products	21
Fixed issues	22
Installation and upgrades fixed issues	22
Veritas Dynamic Multi-pathing fixed issues	23
Veritas Storage Foundation fixed issues	23
Veritas Cluster Server fixed issues	34
Veritas Storage Foundation and High Availability fixed issues	38
Veritas Storage Foundation for Oracle RAC fixed issues	38
Veritas Storage Foundation Cluster File System High Availability fixed issues	40
Known issues	41
Issues related to installation and upgrade	42
Veritas Dynamic Multi-pathing known issues	52
Veritas Storage Foundation known issues	55
Veritas Cluster Server known issues	88

	Veritas Storage Foundation and High Availability known issues	122
	Veritas Storage Foundation Cluster File System High Availability known issues	122
	Veritas Storage Foundation for Oracle RAC known issues	130
	Veritas Storage Foundation for Sybase ASE CE known issues	132
	Symantec VirtualStore known issues	134
	Software limitations	136
	Limitations related to rolling upgrade	136
	Veritas Dynamic Multi-pathing software limitations	137
	Veritas Storage Foundation software limitations	138
	Veritas Cluster Server software limitations	142
	Veritas Storage Foundation and High Availability software limitations	151
	Veritas Storage Foundation Cluster File System High Availability software limitations	152
	Veritas Storage Foundation for Oracle RAC software limitations	153
	Veritas Storage Foundation for Sybase ASE CE software limitations	154
	Symantec VirtualStore software limitations	155
	Documentation errata	156
	List of RPMs	156
	Downloading the 6.0.3 archive	158
Chapter 2	Installing the products for the first time	160
	Installing the Veritas software using the script-based installer	160
	Installing Veritas software using the Web-based installer	161
	Starting the Veritas Web-based installer	162
	Obtaining a security exception on Mozilla Firefox	162
	Installing 6.0.3 with the Veritas Web-based installer	162
Chapter 3	Upgrading to 6.0.3	164
	Prerequisites for upgrading to 6.0.3	164
	Downloading required software to upgrade to 6.0.3	164
	Supported upgrade paths	165
	About using the installer to upgrade when the root disk is encapsulated	165
	Upgrading to 6.0.3	166
	Performing a full upgrade to 6.0.3 on a cluster	167
	Upgrading to 6.0.3 on a standalone system	175

	Upgrading to 6.0.3 on a system that has encapsulated boot disk	178
	Performing a rolling upgrade using the <code>installmr</code> script	178
	Performing a phased upgrade to SFCFSHA and SFRAC	184
	Upgrading the operating system	206
	Verifying software versions	206
Chapter 4	Uninstalling version 6.0.3	207
	About removing Veritas Storage Foundation and High Availability Solutions 6.0.3	207
	Uninstalling Veritas Storage Foundation Cluster File System 6.0.3	207
	Preparing to uninstall Veritas Storage Foundation Cluster File System High Availability 6.0.3	208
	Uninstalling Veritas Storage Foundation Cluster File System High Availability	209
	Uninstalling Veritas Storage Foundation for Oracle RAC 6.0.3	210
	Uninstalling Veritas Storage Foundation for Sybase CE 6.0.3	212

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [Introduction](#)
- [About the installmr script](#)
- [Overview of the installation and upgrade process](#)
- [Changes introduced in 6.0.3](#)
- [System requirements](#)
- [List of products](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation errata](#)
- [List of RPMs](#)
- [Downloading the 6.0.3 archive](#)

Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 6.0.3 release.

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 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.

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

This Maintenance Release is available as 6.0.3.

About the installmr script

Veritas Storage Foundation and High Availability Solutions 6.0.3 provides an upgrade script.

See “[Supported upgrade paths](#)” on page 165.

Symantec recommends that you use the upgrade script. The `installmr` script allows you to upgrade all the patches associated with the packages installed, after which you can reboot to start all of the processes.

The installmr script options

Table 1-1 The command line options for the product upgrade script

Command Line Option	Function
[<i>system1 system2...</i>]	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.

Table 1-1 The command line options for the product upgrade script (*continued*)

Command Line Option	Function
[<code>-precheck</code>]	Use the <code>-precheck</code> option to confirm that systems meet the products' installation requirements before the installation.
[<code>-postcheck</code>]	Use the <code>-postcheck</code> option after an installation or upgrade to help you determine installation-related problems and provide troubleshooting information.
[<code>-responsefile response_file</code>]	Use the <code>-responsefile</code> option to perform automated installations or uninstalls using information stored in a file rather than prompting for information. <i>response_file</i> is the full path of the file that contains configuration definitions.
[<code>-logpath log_path</code>]	Use the <code>-logpath</code> option to select a directory other than <code>/opt/VRTS/install/logs</code> as the location where the <code>installmr</code> log files, summary file, and response file are saved.
[<code>-tmppath tmp_path</code>]	Use the <code>-tmppath</code> option to select a directory other than <code>/var/tmp</code> as the working directory for <code>installmr</code> . This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
[<code>-timeout timeout_value</code>]	Use the <code>-timeout</code> option to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the <code>-timeout</code> option overrides the default value of 1200 seconds. Setting the <code>-timeout</code> option to 0 will prevent the script from timing out. The <code>-timeout</code> option does not work with the <code>-serial</code> option.
[<code>-keyfile ssh_key_file</code>]	Use the <code>-keyfile</code> option to specify a key file for SSH. When you use this option the <code>-i ssh_key_file</code> is passed to every SSH invocation.

Table 1-1 The command line options for the product upgrade script (*continued*)

Command Line Option	Function
[<code>-hostfile <i>hostfile_path</i></code>]	Use the <code>-hostfile</code> option to specify the location of a file containing the system names for installer.
[<code>-patchpath <i>patch_path</i></code>]	Use the <code>-patchpath</code> option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by <code>installmr</code> .
[<code>-kickstart <i>kickstart_path</i></code>]	Use the <code>-kickstart</code> option to generate kickstart scripts which can be used by Redhat Linux Kickstart for automated installation of all rpms for every product, an available location to store the kickstart scripts should be specified as a complete path. The <code>-kickstart</code> option is supported on Redhat Linux only.
[<code>-yumgroupxml <i>yum_group_xml_path</i></code>]	Use the <code>-yumgroupxml</code> option to generate a yum group definition XML file which can be used by <code>createrepo</code> command on Redhat Linux to create yum group for automated installation of all rpms for a product. An available location to store the XML file should be specified as a complete path. The <code>-yumgroupxml</code> option is supported on Redhat Linux only.

Table 1-1 The command line options for the product upgrade script (*continued*)

Command Line Option	Function
<pre>[-serial -rsh -redirect -pkgset -pkgtable -pkginfo -listpatches]</pre>	<p>Use the <code>-serial</code> option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion.</p> <p>Use the <code>-rsh</code> option when <code>rsh</code> and <code>rcp</code> are to be forced for communication though <code>ssh</code> and <code>scp</code> is also setup between the systems.</p> <p>Use the <code>-redirect</code> option to display progress details without showing the progress bar.</p> <p>Use the <code>-pkgset</code> option to discover the package set installed on the systems specified.</p> <p>Use the <code>-pkgtable</code> option to display product rpms in correct installation order.</p> <p>Use the <code>-pkginfo</code> option is used to display the correct installation order of packages and patches. This option is available with or without one of following options: <code>-allpkgs</code>, <code>-minpkgs</code>, and <code>-recpkgs</code>.</p> <p>Use the <code>-listpatches</code> option to display product patches in the correct installation order.</p>

Table 1-1 The command line options for the product upgrade script (*continued*)

Command Line Option	Function
<p>[-makeresponsefile -comcleanup -version -nolic]</p>	<p>Use the <code>-makeresponsefile</code> option to generate a response file without doing an actual installation. Text displaying install, uninstall, start, and stop actions are simulations. These actions are not being performed on the system.</p> <p>Use the <code>-comcleanup</code> option to remove the ssh or rsh configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of ssh or rsh are abruptly terminated.</p> <p>Use the <code>-version</code> option to check the status of installed products on the system.</p> <p>Use the <code>-nolic</code> option to install product rpms on systems without entering product licenses. Configuration, startup, or installation of license based features are not performed when using this option.</p>

Table 1-1 The command line options for the product upgrade script (*continued*)

Command Line Option	Function
<pre>[-upgrade_kernelpkgs -upgrade_nonkernelpkgs -rolling_upgrade -rollingupgrade_phase1 -rollingupgrade_phase2]</pre>	<p>The <code>-upgrade_kernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase1</code>.</p> <p>The <code>-upgrade_nonkernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase2</code>.</p> <p>Use the <code>-rolling_upgrade</code> option to perform rolling upgrade. Using this option, installer will detect the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.</p> <p>Use the <code>-rollingupgrade_phase1</code> option to perform rolling upgrade phase 1. During this phase, the product kernel rpms will be upgraded to the latest version.</p> <p>Use the <code>-rollingupgrade_phase2</code> option to perform rolling upgrade phase 2. During this phase, VCS and other agent rpms will be upgraded to the latest version. During this phase, product kernel drivers will be rolling-upgraded to the latest protocol version.</p>

Overview of the installation and upgrade process

Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

To install or upgrade

- 1 If you are upgrading to 6.0.3, skip to step 2.

If you are installing 6.0.3 for the first time:

- Download Storage Foundation and High Availability Solutions 6.0.1 from <http://fileConnect.symantec.com>.
- Extract the tar ball into a directory called `/tmp/sfha601`.
- Check <http://sort.symantec.com/patches> to see if there are any patches available for the 6.0.1 Installer. Download applicable P-patches and extract them to the `/tmp` directory.

- Change the directory to `/tmp/sfha601`:

```
# cd /tmp/sfha601
```
 - Install the 6.0.1 software. Follow the instructions in the Installation Guide.

```
# ./installer -require complete_path_to_601_installer_patch
```
- 2 Download SFHA 6.0.3 from <http://sort.symantec.com/patches> and extract it to a directory called `/tmp/sfha603`.
 - 3 Check <http://sort.symantec.com/patches> to see if there are patches available for the 6.0.3 installer. Download applicable P-patches and extract them to the `/tmp` directory.
 - 4 Change the directory to `/tmp/sfha603`:

```
# cd /tmp/sfha603
```
 - 5 Install 6.0.3:

```
# ./installmr -require complete_path_to_603_installer_patch
```

Changes introduced in 6.0.3

This section lists the changes in 6.0.3.

Changes introduced in Veritas Cluster Server 6.0.3

This section lists the Veritas Cluster Server changes in 6.0.3.

Db2udb Agent support extended to DB2 10.1

The DB2udb Agent for VCS 6.0.3 now supports DB2 10.1.

Virtual to virtual clustering support in Red Hat Enterprise Virtualization (RHEV) 3.1 and 3.2

In release 6.0.3, VCS provides support for virtual to virtual clustering in RHEV 3.1 and 3.2 environment. The following table lists the software required to run Veritas Cluster Server (VCS) in-guest in the RHEV environment.

Table 1-2 VCS system requirements for KVM-supported RHEV configurations

Supported OS versions in physical hosts	Supported OS versions in virtual machines	Hardware requirement
Red Hat Enterprise Linux 6.3, 6.4 Red Hat Enterprise Virtualization – Hypervisor 3.1, 3.2	RHEL 5 Update 5, 6, 7, 8, 9 RHEL 6 Update 1, 2, 3, 4	Full virtualization-enabled CPU

Note: Before configuring VCS inside virtual machine for virtual to virtual clustering, apply the LLT P-patch: <https://sort.symantec.com/patch/detail/7695>. It solves the issue of LLT port open failure with the error “device or resource busy”.

For information on the implementation details, see *Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide* for Linux.

Changes introduced in SFDB 6.0.3

This section lists the SFDB changes in 6.0.3.

DB2 support is now extended to DB2 10.1

Storage Foundation and High Availability now supports DB2 10.1, 9.7, and 9.5 with fixpack 2 or greater.

System requirements

This section describes the system requirements for this release

Supported Linux operating systems

[Table 1-3](#) lists the supported operating systems for this release of SFHA Solutions.

Table 1-3 Supported operating systems

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 6	Update 1, 2, 3 Note: Update 4 and 5 are also supported if you install VxFS patch 6.0.300.300. For more information, See "Support for RHEL 6.4, RHEL 6.5, OL 6.4, and OL 6.5" on page 19.	2.6.32-131.0.15.el6 2.6.32-220.el6 2.6.32-279.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only
Red Hat Enterprise Linux 5	Update 5, 6, 7, 8, 9,10	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5 2.6.18-308.el5 2.6.18-348.el5 2.6.18-371.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 11	SP1, SP2	2.6.32.12-0.7.1 3.0.42-0.7-default	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 10	SP4	2.6.16.60-0.85.1	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Linux 6	Update 1, 2, 3 Note: Update 4 and Update 5 are also supported if you install VxFS patch 6.0.300.300. For more information, See "Support for RHEL 6.4, RHEL 6.5, OL 6.4, and OL 6.5" on page 19.	2.6.32-131.0.15.el6 2.6.32-220.el6 2.6.32-279.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only

Table 1-3 Supported operating systems (*continued*)

Operating systems	Levels	Kernel version	Chipsets
Oracle Linux 5	Update 5, 6, 7, 8, 9,10	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5 2.6.18-308.el5 2.6.18-348.el5 2.6.18-371.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only

For important updates regarding this release, including the latest supported Red Hat erratas and updates and SUSE service packs, review the Late-Breaking News TechNote on the Symantec Technical Support website:

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

For Veritas Storage Foundation and High Availability Solutions Oracle Support Matrix, refer to <http://www.symantec.com/docs/DOC5081>.

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

Symantec does not support Oracle Linux Unbreakable Enterprise kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Support for RHEL 6.4, RHEL 6.5, OL 6.4, and OL 6.5

Veritas Storage Foundation and High Availability Solutions (SFHA) 6.0.3 by default does not support RHEL 6.4, RHEL 6.5, OL 6.4, or OL 6.5 due to breakages in the kernel interface in the Veritas File System (VxFS) component of SFHA.

To support RHEL 6.4, RHEL 6.5, OL 6.4, or OL 6.5, you must install SFHA 6.0.3, and then install VxFS patch 6.0.300.300 (public hotfix 3 for installation on VxFS 6.0.3) to resolve the kernel incompatibility.

You can download VxFS patch 6.0.300.300 from the following location:

<https://sort.symantec.com/patch/detail/8227>

For more information on support for RHEL 6.4, see the following TechNote:

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

For more information on support for RHEL 6.5, see the following TechNote:

<http://www.symantec.com/docs/TECH212892>

With the combination of SFHA 6.0.3 and VxFS patch 6.0.300.300, the supported RHEL 6 and OL 6 operating systems for this release are as follows:

Table 1-4 Supported Red Hat Enterprise Linux 6 and Oracle Linux 6 operating systems

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 6	Update 1, 2, 3, 4, 5	2.6.32-131.0.15.el6 2.6.32-220.el6 2.6.32-279.el6 2.6.32-358.el6 2.6.32-431.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Linux 6	Update 1, 2, 3, 4, 5	2.6.32-131.0.15.el6 2.6.32-220.el6 2.6.32-279.el6 2.6.32-358.el6 2.6.32-431.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only

VMware Environment

For information about the use of this product in a VMware Environment, refer to <http://www.symantec.com/docs/TECH51941>

Note: This TechNote includes information specific to all 5.1 releases. Please check this technote for the latest information.

Database requirements

The following TechNote identifies the most current information on supported databases (Oracle Single Instance, DB2, and Sybase) and operating system combinations:

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

Note: Veritas Storage Foundation (SF) and Veritas Storage Foundation Cluster File System (SFCFS) do not support running SFDB tools with Sybase, but they support running Oracle and Sybase on VxFS and VxVM.

SF Oracle RAC will announce support for RHEL6 once Oracle supports it. For latest information on support for Oracle database versions with SF Oracle RAC, see the Technical Support TechNote:

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

Veritas Storage Foundation memory requirements

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

Disk space requirements

Before installing your products, confirm that your system has enough free disk space.

Use the **Perform a Pre-installation Check (P)** menu for the Web-based installer to determine whether there is sufficient space.

Or, go to the installation directory and run the installer with the `-precheck` option.

```
# ./installmr -precheck
```

See [“About the installmr script”](#) on page 10.

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

List of products

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

- Veritas Dynamic Multi-Pathing (DMP)
- Veritas Volume Manager (VM)
- Veritas File System (FS)
- 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)

Fixed issues

This section describes the issues fixed in 6.0.3.

See the `README_SYMC.xxxxx-xx` files in the `<architecture>/rpms` directory on the installation media for the symptom, description, and resolution of the fixed issue.

- [Installation and upgrades fixed issues](#)
- [Veritas Dynamic Multi-pathing fixed issues](#)
- [Veritas Storage Foundation fixed issues](#)
- [Veritas Cluster Server fixed issues](#)
- [Veritas Storage Foundation and High Availability fixed issues](#)
- [Veritas Storage Foundation for Oracle RAC 6.0.3 fixed issues](#)
- [Veritas Storage Foundation Cluster File System High Availability fixed issues](#)

Installation and upgrades fixed issues

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

Installation and upgrades: issues fixed in 6.0.3

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

Table 1-5 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 1-6 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 1-7 Fixed issues related to installation and upgrades

Incident	Description
2329580	Unable to stop some SFCFSHA processes.
2873102	Perl module error on completion of SFHA installation
2627076	Incorrect server names sometimes display if there is a clock synchronization issue.
2622987	sfmh discovery issue when you upgrade your Veritas product to 6.0.1
2593148	cssd agent configuration failed with CPI when have two priv IP's in setup.
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 Dynamic Multi-pathing fixed issues

See [Veritas Volume Manager fixed issues](#) for the Veritas Dynamic Multi-pathing fixed issues in 6.0.3. [Veritas Volume Manager fixed issues](#) includes both the VxVM fixed issues and DMP issues.

Veritas Storage Foundation fixed issues

This section describes the Veritas Storage Foundation fixed issues in 6.0.3.

- [Veritas Volume Manager fixed issues](#)
- [Veritas File System fixed issues](#)
- [Veritas Storage Foundation for Databases \(SFDB\) tools fixed issues](#)

Veritas Volume Manager fixed issues

This section describes Veritas Volume Manager fixed issues.

Veritas Volume Manager: issues fixed in 6.0.3

[Table 1-8](#) describes the incidents that are fixed in Veritas Volume Manager in 6.0.3.

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues

Incident	Description
3002770	Accessing NULL pointer in <code>dmp_aa_recv_inquiry()</code> caused system panic.
2971746	For single-path device, <code>bdget()</code> 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	<code>vxassist dump core</code> with the <code>-o ordered</code> 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 <code>vxdisk -o thin, fssize list</code> 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 <code>fsck</code> 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 <code>voldco_update_itemq_chunk()</code> due to accessing invalid buffer
2930569	The LUNs in 'error' state in output of ' <code>vxdisk list</code> ' cannot be removed through DR(Dynamic Reconfiguration) Tool.
2928764	SCSI3 PGR registrations fail when <code>dmp_fast_recovery</code> is disabled.

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

Incident	Description
2919720	vxconfigd core in rec_lock1_5()
2919714	exit code from vxevac is zero when migrating on thin luns but FS is not mounted
2919627	Dynamic Reconfiguration tool should be enhanced to remove LUNs feasibly in bulk.
2919318	The I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing.
2916094	Enhancements have been made to the Dynamic Reconfiguration Tool(DR Tool) to create a separate log file every time DR Tool is started, display a message if a command takes longer time, and not to list the devices controlled by TPD (Third Party Driver) in 'Remove Luns' option of DR Tool.
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 <code>vradmin stoprep</code> comand.
2898547	vradmin 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 <code>tc ./scripts/ddl/dmpapm.tc#11</code> .
2886333	The <code>vxdg(1M) join</code> command should not allow mixing clone and non-clone disks in a DiskGroup.
2878876	vxconfigd dumps core in <code>vol_cbr_dolog()</code> 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 " <code>vxclustadm setmaster</code> " is used to select master.
2866059	Improving error messages hit during the <code>vxdisk resize</code> operation.
2859470	SRDF R2 with EFI label is not recognized by VxVM and showing in error state

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

Incident	Description
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.
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..

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

Incident	Description
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=ctrl" in default vxassist rulefile
2273190	Incorrect setting of the UNDISCOVERED flag can lead to database inconsistency.
2149922	Record the diskgroup import and deport events in syslog.
2000585	The <code>vxrecover -s</code> 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 <code>vxdmpadm listctrl</code> 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 6.0.1. This list includes Veritas Volume Replicator and Cluster Volume Manager fixed issues.

Table 1-9 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 <code>vxddmpadm list dmpnode dmpnodename=dmpnode name</code> .
2818840	Enhance the <code>vxddmprow</code> utility to support permission and "root:non-system" ownership to be set and make it persistent.
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 <code>svol_flush_srl_to_dv_start</code> fails to start.
2771452	IO hung because of hung port deletion.
2763206	The <code>vxdisk rm</code> command core dumps when list of disknames is very long.
2756059	Panic in <code>voldco_or_drl_to_pvm</code> 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 <code>vxsnap</code> operation.
2751102	Random panics seen in <code>vx_worklist_thr</code> on SLES11 and VxFS.
2747032	Write is taking long time to complete when read/write happen simultaneously.
2743926	DMP <code>restored</code> daemon fails to restart during system boot.
2741240	The <code>vxdbg join</code> transaction failed and did not rollback to the <code>sourcedg</code> .
2739709	Disk group rebuild related issues.
2739601	VVR: <code>repstatus</code> output occasionally reports abnormal timestamp.
2737420	The <code>vxconfigd</code> 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.

Table 1-9 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2726148	System unbootable after <code>/usr/lib/vxvm/bin/vxupdate1vm</code> script updates filter in <code>lvm.conf</code> 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.
2710147	Node panics in <code>dmp_pr_do_reg</code> 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 <code>vxconfigd</code> daemon during CVM startup.
2700486	The <code>vradmind</code> daemon core dumps 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 <code>vxassist mirror</code> command failed for thin LUN because <code>statvfs</code> 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 <code>vx.*</code> directories are left in the <code>/tmp</code> directory.
2680343	Node panic during <code>cur pri</code> path update in cluster while running I/O shipping.
2679917	Corrupt space optimized snapshot after a refresh with CVM master switching.
2675538	The <code>vxdisk resize</code> 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.

Table 1-9 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2660151	The <code>vxconfigd</code> 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 <code>vxnetd start</code> and <code>stop</code> operations causes panic.
2652485	Inactive snapshot LUNs cause trespassing.
2648176	Performance difference on Master versus Slave during recovery with Data Change Object (DCO).
2645196	Campus Cluster + Hot Relocation: When a disk failure is detected, the associated disks for that site are detached and ALL disks as marked as RLOC.
2644248	The <code>vxunroot</code> command fails as root partition "logvol" mounted on <code>/var/log</code> .
2643634	Message enhancement for a mixed (non-cloned and cloned) disk group import.
2627126	Lots of I/Os and paths are stuck in <code>dmp_delayq</code> and <code>dmp_path_delayq</code> respectively. DMP daemon did not wake up to process them.
2626199	The <code>vxddmpadm list dmpnode</code> printing incorrect path type.
2623182	<code>vxvm-boot</code> not cleaning up <code>/tmp/vx.*</code> directories whenever system reboot is done for Linux environment.
2620555	I/O 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 <code>dmpnode</code> 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 <code>volcvm_msg_rel_gslock()</code> .
2564092	Automate the LUN provisioning (addition) / removal steps using <code>vxdiskadm</code> .
2553729	Status of the EMC Clariion disk changed to "online clone_disk" after upgrade.
2486301	"VXFS" installation failed.
2441283	The <code>vxsnap addmir</code> command sometimes fails under heavy I/O load.

Table 1-9 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
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 <code>vxvg move</code> 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 <code>vxvg list dgname</code> command gives error 'state=new failed'.
1533134	DMP: deprecated SCSI <code>ioctl</code> use <code>sg_io</code> type of error.
1190117	<code>vxdisk -f init</code> can overwrite some of the public region contents.
2698035	Tunable values do not change as per the values applied through <code>vxvtune</code> .
2682491	The <code>vxvmconvert</code> command displays an error message while converting an LVM volume.

Veritas File System fixed issues

This section describes Veritas File System fixed issues.

Veritas File System: issues fixed in 6.0.3

[Table 1-10](#) describes the incidents that are fixed in Veritas File System in 6.0.3.

Table 1-10 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 <code>vxrestore</code> , only for files greater than 8192bytes.
2857751	The internal testing hits the assert " <code>f:vx_cbdnlc_enter:1a</code> ".

Table 1-10 Veritas File System 6.0.3 fixed issues (*continued*)

Incident	Description
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 1-11 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 fspadm 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.

Table 1-11 Veritas File System fixed issues (*continued*)

Incident	Description
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 Storage Foundation for Databases (SFDB) tools fixed issues

This section describes Veritas Storage Foundation for Databases (SFDB) tools fixed issues.

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

[Table 1-12](#) describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 6.0.3.

Table 1-12 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.

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

Table 1-13 describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-13 SFDB tools fixed issues

Incident	Description
2585643	<p>If you provide an incorrect host name with the <code>-r</code> option of <code>vxsfadm</code>, the command fails with an error message similar to one of the following:</p> <pre>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.</pre> <p>The error messages are unclear.</p>
2703881 (2534422)	<p>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:</p> <pre>SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.</pre>
2582694 (2580318)	<p>After you have done FlashSnap cloning using a snapplan, any further attempts to create a clone from the same snapplan using the <code>dbed_vmclonedb</code> continue to use the original clone SID, rather than the new SID specified using the <code>new_sid</code> 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.</p>
2579929	<p>The <code>sfae_auth_op -o auth_user</code> command, used for authorizing users, fails with the following error message:</p> <pre>SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username></pre> <p>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.</p>

Veritas Cluster Server fixed issues

This section describes the Veritas Cluster Server fixed issues.

Veritas Cluster Server: issues fixed in 6.0.3

Table 1-14 describes the incidents that are fixed in Veritas Cluster Server in 6.0.3.

Table 1-14 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
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 lcmp 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 1-15](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-15 LLT, GAB, and I/O fencing fixed issues

Incident	Description
2708619	If you set the <code>scsi3_disk_policy</code> attribute to <code>dmp</code> , you cannot enable the Veritas fencing module (VxFEN). The VxFEN source code is updated to pick up the <code>dmp</code> device path that contains the full disk name instead of a partition or slice.
2845244	<code>vxfen</code> startup script gives error <code>grep: can't open /etc/vxfen.d/data/cp_uid_db</code> . The error comes because <code>vxfen</code> startup script tries to read a file that might not be present. This error is typically seen when starting <code>vxfen</code> for the very first time after installation.
2554167	Setting <code>peerinact</code> value to 0 in the <code>/etc/llttab</code> file floods the system log file with large number of log messages.

Bundled agents fixed issues in 6.0.1

[Table 1-16](#) lists the fixed issues for bundled agents.

Table 1-16 Bundled agents fixed issues

Incident	Description
2850904	Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and <code>PanicSystemOnDGLoss</code> is set to 0.
2794175	Fire drill for Mount agent does not accurately portray a failover scenario.
2728802	If the <code>'httpd'</code> binary or the <code>'ab'</code> binary is not present at the location that you specified in the <code>'httpdDir'</code> attribute, the Apache agent cannot perform detail monitoring or start the HTTP server.
2850905	IMF registration for Mount resource for file systems type other than VxFS and NFS should be blocked.
2850916	Mount resource does not get registered with IMF if the attributes <code>BlockDevice</code> and/or <code>MountPoint</code> have a trailing slash in their values.
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 (<code>hastop -local -force</code>), if disk reservation resource is configured then user may observe a system panic.

Table 1-16 Bundled agents fixed issues (*continued*)

Incident	Description
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 1-17](#) lists the fixed issues for VCS engine.

Table 1-17 VCS engine fixed issues

Incident	Description
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility <code>hagrp</code> gives incorrect output with the "-clear", "-flush", "-state" options.
2741299	CmdSlave gets stuck in a tight loop when it gets an EBADF on a file descriptor(fd). The CmdSlave process keeps retrying on the FD and eventually dumps core.
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 <code>-nopro</code> option is selected.
2684818	If the following attributes are specified before SystemList attribute in <code>main.cf</code> , then the value got rejected when HAD started: <ul style="list-style-type: none"> ■ PreOnline ■ ContainerInfo ■ TriggersEnabled ■ SystemZones

Table 1-17 VCS engine fixed issues (*continued*)

Incident	Description
2696056	Memory leak occurs in the engine when <code>haclus -status <cluster></code> command is run.
2746802	When failover group is probed, VCS engine clears the <code>MigrateQ</code> and <code>TargetCount</code> .
2746816	The <code>syslog</code> call used in <code>gab_heartbeat_alarm_handler</code> and <code>gabsim_heartbeat_alarm_handler</code> functions is not <code>async signal safe</code> .

Enterprise agents fixed issues in 6.0.1

[Table 1-18](#) lists the fixed issues for enterprise agents.

Table 1-18 Enterprise agents fixed issues

Incident	Description
1985093	Ensure that the <code>ohasd</code> 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 1-19](#) lists the fixed issues for agent framework.

Table 1-19 Agent framework fixed issues

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

Veritas Storage Foundation and High Availability fixed issues

For fixed issues of Veritas Storage Foundation and High Availability in this release, see [Veritas Storage Foundation fixed issues](#) and [Veritas Cluster Server fixed issues](#).

Veritas Storage Foundation for Oracle RAC fixed issues

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

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 1-20](#) lists the issues fixed in 6.0.1.

Table 1-20 Issues fixed in 6.0.1

Incident	Description
2585899	<p>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:</p> <pre>The node sys1 doesn't seem to be part of the cluster, or CVM is not running on the node sys1.</pre>
2329580	<p>If you install and start SF, but later configure SF using <code>installvcs</code>, some drivers may not stop successfully when the installer attempts to stop and restart the SF drivers and processes. The reason the drivers do not stop is because some dependent SF processes may be in the running state.</p>
2392741	<p>Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool.</p> <p>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.</p>
2749412	<p>Setting the <code>UseVirtualIP</code> attribute to 1 overwrites the IP address of the virtual interface on some nodes in the cluster.</p>
2757032	<p>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.</p>

Table 1-20 Issues fixed in 6.0.1 (*continued*)

Incident	Description
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 <code>UseVirtualIP</code> attribute is set to 1.
2740150	The SF Oracle RAC installer fails to set the value of the CSSD resource attribute <code>OfflineWaitLimit</code> 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 <code>add_drv/rem_drv</code> race between our modules which are independent of each other.

Veritas Storage Foundation Cluster File System High Availability fixed issues

This section describes the Veritas Storage Foundation Cluster File System High Availability fixed issues.

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

[Table 1-21](#) describes the Veritas Storage Foundation Cluster File System fixed issues in 6.0.3.

Table 1-21 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues

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

Table 1-21 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues (*continued*)

Incident	Description
2906018	The vx_iread errors are displayed after successful log replay and mount of the file system.
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 1-22 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.

Known issues

This section covers the known issues in 6.0.3 and 6.0.1.

- [Issues related to installation and upgrade](#)
- [Veritas Dynamic Multi-pathing known issues](#)
- [Veritas Storage Foundation known issues](#)

- [Veritas Cluster Server known issues](#)
- [Veritas Storage Foundation and High Availability known issues](#)
- [Veritas Storage Foundation Cluster File System High Availability known issues](#)
- [Veritas Storage Foundation for Oracle RAC known issues](#)
- [Veritas Storage Foundation for Sybase ASE CE known issues](#)
- [Symantec VirtualStore known issues](#)

Issues related to installation and upgrade

This section describes the known issues during installation and upgrade in 6.0.3 and 6.0.1.

During VCS upgrades on Linux from 6.0.2 to 6.0.3 using vCenter Integrated Installation, installer displays error messages (3328214)

During VCS upgrades on Linux from 6.0.2 to 6.0.3 using vCenter Integrated Installation, installer displays the following error message:

```
EULA file not found at C:\Program Files\Veritas\ApplicationHA\  
Console\portal\world\Installers\VCS_602\...
```

Workaround:

To resolve this issue:

- 1 Copy the Linux 6.0.2 Installation package on the Console host again.
- 2 If you haven't applied the following patch on the Console host, do so now:
<https://sort.symantec.com/patch/detail/7621>
- 3 To upgrade VCS to the 6.0.3 version, use the `CopyInstaller.bat -import` option.

For detailed steps about upgrading VCS with `CopyInstaller.bat -import`, see: <http://www.symantec.com/docs/TECH200772>.

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/install1.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/install1.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 SF 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 SF, 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>

Installing DMP with a keyless license or DMP-only license does not enable DMP native support for LVM root volumes (2874810)

When you install DMP with a keyless license or DMP-only license, the tunable parameter `dmp_native_support` is set to on. However, the DMP native support is not enabled for LVM root volumes. The DMP native support is enabled for non-root LVM volumes.

Workaround:

After package installation, use the following command to enable the DMP support for root LVM volumes.

```
# vxddmpadm settune dmp_native_support=on
```

Then reboot the system.

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/opensv`), 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 RPM `VRTSspbx`, `VRTSat`, and `VRTSicisco`. This causes NetBackup to stop working.

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the `/usr/opensv/netbackup/bin/version` file and `/usr/opensv/netbackup/version` file to the `/tmp` directory. If you have clustered NetBackup installed, you must also copy the `/usr/opensv/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/opensv/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/opensv/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 `VRTSpxbx`, `VRTSat`, and `VRTSicscso` RPM 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 partitioning on disk /dev/sda is not readable by
The partitioning 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.

During upgrade from 5.1SP1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name is used by a deported disk group (2280560)

During an upgrade from SF 5.1 SP1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name for the split operation is used by an existing deported disk group.

Workaround:

Specify a different disk group name as a target for the split operation.

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

```
# hagrps -list Frozen=1
```

- 2 Unfreeze all the frozen service groups:

```
# haconf -makerw  
# hagrps -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 not 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 whereas the IPv4RouteOptions value is used by the `ip route` command. The values to be configured for these two attributes are very specific to their respective commands.

Table 1-23 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	<p>In this case the <code>ifconfig</code> command is used. If <code>RouteOptions</code> is set, attribute value is used to add/delete routes using command <code>route</code>.</p> <p>As the <code>Options</code> attribute is configured, <code>IPv4RouteOptions</code> values are ignored.</p>	No need to configure <code>IPv4RouteOptions</code> .
Not configured	May or may not be configured	Must be configured	<p>In this case the <code>ip</code> command is used. <code>IPv4RouteOptions</code> must be configured and are used to add/delete routes using the <code>ip route</code> command. As <code>Options</code> attribute is not configured, <code>RouteOptions</code> value is ignored.</p>	<p>Configure <code>IPv4RouteOptions</code> and set the IP of default gateway. The value of this attribute typically resembles: <code>IPv4RouteOptions = "default via gateway_ip"</code></p> <p>For example: <code>IPv4RouteOptions = "default via 192.168.1.1"</code></p>

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 6.0.3 (3064919)

If you upgrade SF to 6.0.3 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 6.0.3. 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.

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

The `lvm.conf` file has the `obtain_device_list_from_udev` variable. If `obtain_device_list_from_udev` is set to 1, then LVM uses devices and symlinks specified by udev database, only.

Workaround:

In the `lvm.conf` file, set the `obtain_device_list_from_udev` variable to 0. This enables LVM to recognize `/dev/vx/dmp/` devices when performing a `vgscan`, which enables volume group migration to succeed.

Veritas Dynamic Multi-pathing known issues

This section describes the Veritas Dynamic Multi-pathing known issues in 6.0.3 and 6.0.1.

Migration of I/O fencing-enabled disks of VxVM disk group from EMC PowerPath TPD to VxVM DMP fails [3528561]

If I/O Fencing is enabled on some disks from VxVM disk group, migration of those disks from EMC PowerPath TPD to VxVM DMP fails with the following error messages:

```
VXFEN vxfenconfig NOTICE Driver will use SCSI-3 compliant disks.  
VXFEN vxfenconfig ERROR V-11-2-1090 Unable to register with a  
Majority of the coordination points.
```

Workaround: Restart the server.

Symantec has reported the issue to EMC PowerPath Engineering.

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.

```
# vxddmpadm 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 `vxctl 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:

```
# vxddmpadm 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

```
# vxddmpadm 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.

```
# vxddmpadm 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.

```
# vxddm adm 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:

```
# vxddm adm settune dmp_monitor_ownership=off
```

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.

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.

Veritas Storage Foundation known issues

This section describes the Veritas Storage Foundation known issues in 6.0.3 and 6.0.1.

- [Veritas Storage Foundation known issues](#)
- [Veritas Volume Manager known issues](#)
- [Veritas File System known issues](#)
- [Replication known issues](#)
- [Veritas Storage Foundation for Databases \(SFDB\) tools known issues](#)

Veritas Storage Foundation known issues

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

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 SFM GUI.

Workaround:

To resolve this known issue

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

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

An error message is received when you perform off-host clone for RAC and the off-host node is not part of the CVM cluster (1834860)

There is a known issue when you try to perform an off-host clone for RAC and the off-host node is not part of the CVM cluster. You may receive a similar error message:

```
Cannot open file /etc/vx/vxdba/racl1g1/.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'racl1g1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdg ERROR V-5-1-582 Disk group SNAP_racl1dgl: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_racl1dgl.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_racl1dgl failed.
```

Workaround:

Currently there is no workaound for this known issue. However, if the off-host node is part of the CVM cluster, then off-host clone for RAC works fine.

Also the `dbed_vmclonedb` command does not support `LOCAL_LISTENER` and `REMOTE_LISTENER` in the `init.ora` parameter file of the primary database.

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 dynamic storage tiering 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 Volume Manager known issues

The following are the Veritas Volume Manager known issues for this release.

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.conf` file 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 `vxddmpadm(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 `vxddmpadm(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 "vxdbg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set (2354560)

In Cluster Volume Manager environment, "vxdbg 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 "vxdbg 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 SF 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 `io_timeout` tunable to 600:

```
# vxddmpadm setattr enclosure encl1 recoveryoption=throttle \  
io_timeout=600
```

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

```
# vxddctl 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 5.x to 6.0.1 may fail for IBM XIV Series arrays (2715119)

Starting in the Veritas Storage Foundation 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 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 environment 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.

Complete site is detached, if plex detach operation is performed even after site consistency off (2845383)

By design, you cannot detach the last plex of a site on a site consistent volume without detaching the complete site. By default, attempting to detach the last plex causes an error. If you use the force detach option, then the complete site is detached to ensure site consistency. This behavior is seen even if you turn off the site consistent flag if the `allsites` flag is on.

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 multiple 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>

Veritas File System known issues

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

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/voll -  
blocks are currently in use.  
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume  
voll, 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 degradation is visible 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'.

Replication known issues

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

`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 `vxxrvg -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 `vradmin` daemon on all VVR hosts:

```
# /etc/init.d/vras-vradmin.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
```

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.

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:

```
# vxrvrg -g diskgroup stop rvrg
```
- 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 rvrg vol
```
- 7 Start the RVG. Enter the following:

```
# vxrvrg -g diskgroup start rvrg
```
- 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.

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.

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SF. 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 of 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.

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.

SFDB commands do not work with the ZHS16GBK character set (2715323)

SFDB commands do not work if the character set of the Oracle database is set to ZHS16GBK. This occurs because SFDB commands are not supported with multi-byte character sets except AL32UTF8 and ZHS16GBK is a multi-byte character set.

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 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 workaround at this point of time.

Checkpoint clone fails if the archive log destination is same as the datafiles destination (2869266)

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 6.0.3 release, create distinct archive and datafile mounts for the 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.

Checkpoint clone fails in CFS environment if cloned using same checkpoint and same clone name on both nodes (2869268)

The Checkpoint clone of an oracle database fails in a CFS environment, if you create a clone with a clone name and checkpoint name same as another clone up on a different CFS node.

Workaround:

There is no workaround. Create a clone with a different clone name.

Very long off-host cloning times for large number of datafiles (2849540)

When cloning off-host in certain Oracle database configurations, particularly with several hundred datafiles, the cloning can take a very long time, upto an hour or more. This problem does not cause the cloning to fail. The problem applies to all services such as FlashSnap, Space-optimized snapshots, FileSnap, and 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 amount: ERROR: V-3-26388: file system /rep has been mount
locked
SFORA sfua_rept_migrate ERROR V-81-5550 amount /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 6.0.3 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
```

Veritas Cluster Server known issues

This section describes the Veritas Cluster Server known issues in 6.0.3, 6.0.2, and 6.0.1.

- [Operational issues for VCS](#)
- [Issues related to the VCS engine](#)
- [Issues related to the bundled agents](#)
- [Issues related to the VCS database agents](#)
- [Issues related to the agent framework](#)
- [Issues related to global clusters](#)
- [LLT known issues](#)
- [GAB known issues](#)
- [I/O fencing known issues](#)
- [Veritas Cluster Server agents for Veritas Volume Replicator known issues](#)
- [Issues related to Intelligent Monitoring Framework \(IMF\)](#)
- [Issues related to the Cluster Manager \(Java Console\)](#)
- [Issues related to virtualization](#)

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.

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.

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 SF 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 dependency is disabled then the other resources do not come online and the following message is displayed:

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

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

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

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

Workaround: Terminate the agent process.

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

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

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

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

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

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

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`:

```
hagrpl -offline -force ClusterService -any
```

or

```
hagrpl -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.

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
```

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.

Issues related to the bundled agents

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 CFSSMount Errors [2564411]

When you try to bring the FireDrill service group online through Java Console or `hagr -online` command, the CFSSMount 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 environment.

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.

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

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

Workaround:

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

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 `rhevms -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.

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

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

Workaround: No workaround.

KVMGuest returns the resource state as OFFLINE for virtual machine in a "non-responding" state [2848003]

In case virtual machine goes into "non-responding" state, the KVMGuest agent returns the resource state as OFFLINE. This situation occurs if storage domain is in "inactive" state and data-center is in "down" state.

Workaround: Activate the storage domain in RHEV-M and make sure that the data-center is in "up" state.

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. SF 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.

Virtual machine may return the not-responding state when the storage domain is inactive and the data center is down (2848003)

In a Red Hat Enterprise Virtualization Environment, if the storage domain is in an inactive state and the data center is in down state, the virtual machine may return a not-responding state and the KVMGuest resource in OFFLINE state.

Workaround: To resolve this issue:

- 1 Activate the storage domain in RHEV-M.
- 2 Check that the data center is in the up state.

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.

Issues related to the VCS database agents

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 [1985055]

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2.

Workaround: Set MonitorOption attribute for Oracle resource to 0.

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

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

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

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

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

No health check monitoring for Oracle agent on SLES11 platform [1919203]

Oracle agent does not support health check monitoring on SLES11 platform.

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 heartbeat with the VCS engine under heavy load.

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

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

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

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

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

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

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

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

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

Workaround: No workaround.

Issues related to global clusters

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

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

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

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

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

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

LLT known issues

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

LLT 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 `llttab` 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 `llttab` 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 deinitd on user request
```

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

Workaround: There is no workaround for this issue.

I/O fencing known issues

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

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 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 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 coordinator disk group

Workaround: There is no workaround for this issue.

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

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

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

1 Rename cpsadm to cpsadmbin

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

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

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

3 Give executable permissions to the new file.

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

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

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

Workaround:

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

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.

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 `vxfereswap` 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 `vxfereswap` 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 `vxferenadm -d` command displays the following error:

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

Workaround: Start VxFEN again after some time.

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 `vxferesthdw` utility directly from the install media if VRTSvxfer package is not installed on the system (2858190)

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

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

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 6.0.3 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`.

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

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

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

It should be replaced as follows in main.cf:

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

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

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

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

Direct execution of linkamf displays syntax error [2858163]

Bash cannot interpret Perl when executed directly.

Workaround: Run linkamf as follows:

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

Error messages displayed during reboot cycles [2847950]

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

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

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

Workaround: No workaround.

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

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

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

Workaround: No Workaround.

Error message seen during system shutdown [2954309]

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

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

The system continues to proceed with the shutdown.

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

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

For example:

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

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

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

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

Workaround: Use the correct commands:

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

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

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.

Workaround: Restart the `libvirtd` process and run:

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

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

In a virtualization environment, SF detects the virtual machine migration initiated outside SF and changes the state accordingly. However, occasionally, SF 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 SF 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 because 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 `libvirtd` 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 acpihp  
# modprobe pci_hotplug
```

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

If a virtual machine is under SF control, and you revert to the virtual machine configuration snapshot, the SF configuration file moves to the read-only mode. During this time, if you run a SF 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.

Veritas Storage Foundation and High Availability known issues

For known issues of Veritas Storage Foundation and High Availability in this release, see [Veritas Storage Foundation known issues](#) and [Veritas Cluster Server known issues](#).

Veritas Storage Foundation Cluster File System High Availability known issues

This section describes the Veritas Storage Foundation Cluster File System High Availability known issues in 6.0.3 and 6.0.1.

NFS issues with VxFS checkpoint (2027492)

NFS clients mounting VxFS checkpoints that are NFS-exported by SFCFS 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 checkpoints not necessarily being the same on all SFCFS 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 checkpoint to remain the same on all cluster nodes.

For example:

To NFS-export a VxFS 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 checkpoint and add the `fsid share` option to force the `fsid` of an NFS-exported VxFS 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 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 issues with VxFS Storage Checkpoints (2027492)

NFS clients mounting VxFS Storage Checkpoints that are NFS-exported by SFCFSHA 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 cluster nodes.

Workaround: The following procedure resolves this issue.

To resolve this issue

- ◆ 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.

To resolve this issue

- ◆ 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 file systems.

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.

Multiple system panics upon unmounting a CFS file system (2107152)

There is a system panic when you unmount a `mntlock`-protected VxFS file system, if that device is duplicate mounted on different directories.

Workaround: There is no workaround for this issue.

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 SF 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 unmounted 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.

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 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 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 checkpoint or snapshot on the active node.

Flashsnap clone fails under some unusual archive log 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'
```

Veritas Storage Foundation for Oracle RAC known issues

This section describes the Veritas Storage Foundation for Oracle RAC known issues in 6.0.3 and 6.0.1.

- [Oracle RAC issues](#)
- [SF issues](#)

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/npoahsd",
O_WRONLY <unfinished ...>
```

For possible causes and workarounds, see the Oracle Metalink document: 1069182.1

SF issues

This section lists the known issues in SF 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 SF 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 SF 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 SF components) are not executed and the node being started does not join the SF cluster.

Workaround: If the rebooted node does not join the SF 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 `dmpmode`.

Rolling upgrade is not supported if you are upgrading from SF Oracle RAC 5.1 SP1 with fencing configured in `dmpmode`. 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 `dmpmode`.

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.

Veritas Storage Foundation for Sybase ASE CE known issues

This section describes the Veritas Storage Foundation for Sybase ASE CE known issues in 6.0.3 and 6.0.1.

SF issues

This section lists the known issues in SF 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: $((\text{number of retries} + 1) * (\text{quorum heartbeat interval})) + 5$.

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
/qrmnt/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 Symantec VirtualStore known issues in 6.0.3.

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.

Software limitations

This section covers the software limitations in 6.0.3, 6.0.2, and 6.0.1.

- [Veritas Dynamic Multi-pathing software limitations](#)
- [Veritas Storage Foundation software limitations](#)
- [Veritas Cluster Server software limitations](#)
- [Veritas Storage Foundation and High Availability software limitations](#)
- [Veritas Storage Foundation Cluster File System High Availability software limitations](#)
- [Veritas Storage Foundation for Oracle RAC software limitations](#)
- [Veritas Storage Foundation for Sybase ASE CE software limitations](#)
- [Symantec VirtualStore software limitations](#)

Limitations related to rolling upgrade

Rolling upgrade with responsefile to 6.0.1 is not supported for VCS, SFHA, SFCFSHA, SFRAC, SFSYBASECE, and SVS.

Veritas Dynamic Multi-pathing software limitations

This section describes the Veritas Dynamic Multi-pathing software limitations in 6.0.3 and 6.0.1.

DMP behavior on Linux SLES11 when connectivity to a path is lost (2049371)

On SLES 11, when the connectivity to a path is lost, the SLES 11 kernel removes the device path from its database. DMP reacts to the UDEV event that is raised in this process, and marks the device path as DISABLED[M]. DMP will not use the path for further I/Os. Unlike on other flavours of Linux, the path state is DISABLED[M] instead of DISABLED. Subsequently, if the path comes back online, DMP responds to the UDEV event to signal the addition of device path into SLES 11 kernel. DMP enables the path and changes its state to ENABLED.

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, change the default values for the DMP tunable parameters.

[Table 1-24](#) describes the DMP tunable parameters and the new values.

Table 1-24 DMP settings for NetApp storage attached environment

Parameter name	Definition	New value	Default value
dmp_restore_interval	DMP restore daemon cycle	60 seconds.	300 seconds.
dmp_path_age	DMP path aging tunable	120 seconds.	300 seconds.

The change is persistent across reboots.

To change the tunable parameters

- 1 Issue the following commands:

```
# vxddmpadm settune dmp_restore_interval=60  
  
# vxddmpadm settune dmp_path_age=120
```

- 2 To verify the new settings, use the following commands:

```
# vxddmpadm gettune dmp_restore_interval  
  
# vxddmpadm gettune dmp_path_age
```

LVM volume group in unusable state if last path is excluded from DMP (1976620)

When a DMP device is used by a native LVM volume group, do not exclude the last path to the device. This can put the LVM volume group in an unusable state.

DMP does not support devices in the same enclosure that are configured in different modes (2643506)

DMP does not support the configuration where two devices in the same enclosure are configured in different modes. For example, if one device is configured as ALUA and another one is configured as Active/Passive (A/P).

Veritas Storage Foundation software limitations

This section describes the Veritas Storage Foundation software limitations in 6.0.3 and 6.0.1.

Veritas Volume Manager software limitations

The following are software limitations in this release of Veritas Volume Manager.

The `vxlist` command cannot correctly display numbers greater than or equal to 1 EB

The `vxlist` command and all of the other commands that use the same library as the `vxlist` command cannot correctly display numbers greater than or equal to 1 EB.

You are unable to specify units in the tunables file

The CPI displays an error when you set the unit type, such as MB or GB, for tunable parameters in the tunables file that you specify with the `-tunablesfile file` option. To avoid the error, you must set the unit type manually.

Snapshot configuration with volumes in shared disk groups and private disk groups is not supported

A snapshot configuration with volumes in the shared disk groups and private disk groups is not a recommended configuration. In this release, this configuration is not supported.

Storage reclamation does not happen on volumes with break-off snapshot (2798523)

In this release, storage reclamation on a volume is prevented when it has a break-off type snapshot. If storage reclamation is allowed on such volumes, it can lead to the following undesired situation. Instant snapshot operations, including `vxsnap refresh` and `vxsnap restore` operations, lead to full synchronization of either the snapshot or the primary volume depending on the operation.

In this release, if the volume has a snapshot, the storage reclamation is silently prevented. The physical storage is not reduced. The reclaim command reports that the reclamation is done on the disks but the actual storage is not reclaimed for volumes with snapshots:

```
# vxdisk -o full reclaim dg1
Reclaiming storage on:
Disk xiv0_617 : Done.
Disk xiv0_616 : Done.
Disk xiv0_618 : Done.
Disk xiv0_612 : Done.
Disk xiv0_613 : Done.
Disk xiv0_614 : Done.
Disk xiv0_615 : Done
```

As shown in the following example output, the storage is not actually reclaimed.

```
# vxdisk -o thin list
DEVICE      SIZE(MB)  PHYS_ALLOC(MB)  GROUP  TYPE
xiv0_612    19313     2101            dg1    thinrclm
xiv0_613    19313     2108            dg1    thinrclm
xiv0_614    19313     35              dg1    thinrclm
xiv0_615    19313     32              dg1    thinrclm
xiv0_616    19313     31              dg1    thinrclm
```

```
xiv0_617 19313 31          dg1  thinrc1m
xiv0_618 19313 31          dg1  thinrc1m
```

SF does not support thin reclamation of space on a linked mirror volume (2729563)

The thin reclamation feature does not support thin reclamation for a linked mirror volume.

Veritas File System software limitations

The following are software limitations in the 6.0.3 release of Veritas Storage Foundation.

Linux I/O Scheduler for Database Workloads

Symantec recommends using the Linux deadline I/O scheduler for database workloads on both Red Hat and SUSE distributions.

To configure a system to use this scheduler, include the `elevator=deadline` parameter in the boot arguments of the GRUB or LILO configuration file.

The location of the appropriate configuration file depends on the system's architecture and Linux distribution:

Configuration File

```
/boot/grub/menu.lst
```

Architecture and Distribution

```
RHEL5 x86_64, RHEL6 x86_64, SLES10
x86_64, and SLES11 x86_64
```

For the GRUB configuration files, add the `elevator=deadline` parameter to the kernel command.

For example, for RHEL5, change:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2
    initrd /boot/initrd-2.6.18-128.el5.img
```

To:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2 \
    elevator=deadline
    initrd /boot/initrd-2.6.18-128.el5.img
```

For RHEL6, change:

```
title RHEL6
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.32-71.el6 ro root=/dev/sdb2
    initrd /boot/initrd-2.6.32-71.el6.img
```

To:

```
title RHEL6
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.32-71.el6 ro root=/dev/sdb2 \
    elevator=deadline
    initrd /boot/initrd-2.6.32-71.el6.img
```

A setting for the elevator parameter is always included by SUSE in its LILO and GRUB configuration files. In this case, change the parameter from `elevator=cfq` to `elevator=deadline`.

Reboot the system once the appropriate file has been modified.

See the Linux operating system documentation for more information on I/O schedulers.

Recommended limit of number of files in a directory

To maximize VxFS performance, do not exceed 100,000 files in the same directory. Use multiple directories instead.

Limitations with delayed allocation for extending writes feature

The following limitations apply to the delayed allocation for extending writes feature:

- In the cases where the file data must be written to disk immediately, delayed allocation is disabled on that file. Examples of such cases include Direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.
- Delayed allocation is not supported on memory mapped files.
- Delayed allocation is not supported with BSD quotas. When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.
- Delayed allocation is not supported for shared mounts in a cluster file system.

FlashBackup in NetBackup 7.1 and prior does not support disk layout Version 8 and 9

The FlashBackup feature of NetBackup 7.1 or prior does not support a VxFS file system with disk layout Version 8 or 9.

Compressed files that are backed up using NetBackup 7.1 or prior become uncompressed when you restore the files

The NetBackup 7.1 release and prior does not support the file compression feature. If you back up compressed files using NetBackup 7.1 or a prior release, the files become uncompressed when you restore the files.

Veritas Storage Foundation for Databases (SFDB) tools software limitations

The following are the SFDB tools software limitations in this release.

Oracle Data Guard in an Oracle RAC environment

Database snapshots and Database Storage Checkpoints are not supported in a Data Guard with Oracle RAC environment.

Upgrading to Oracle 10.2.0.5 is required if using SFDB tools

If you are running Oracle version 10.2.0.4 and upgrading a Storage Foundation product with SFDB tools to 6.0.3, you must upgrade the Oracle binaries and database to version 10.2.0.5, before upgrading to 6.0.3.

Parallel execution of `vxsfadm` is not supported (2515442)

Only one instance of the `vxsfadm` command can be run at a time. Running multiple instances of `vxsfadm` at a time is not supported.

Creating point-in-time copies during database structural changes is not supported (2496178)

SFDB tools do not support creating point-in-time copies while structural changes to the database are in progress, such as adding or dropping tablespaces and adding or dropping data files.

However, once a point-in-time copy is taken, you can create a clone at any time, regardless of the status of the database.

Veritas Cluster Server software limitations

This section describes the Veritas Cluster Server software limitations in 6.0.3, 6.0.2, and 6.0.1.

Limitations related to installing and upgrading VCS

Remote and target systems must have the same OS and architecture while using installer from a remote system [589334]

If you use the installer from a remote system, then the remote system must have the same operating system and architecture as that of the target systems where you want to install VCS.

Limitations related to bundled agents

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

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

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

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

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

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

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

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

Volume agent clean may forcibly stop volume resources

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

False concurrency violation when using PidFiles to monitor application resources

The PID files created by an application contain the PIDs for the processes that are monitored by Application agent. These files may continue to exist even after a node

running the application crashes. On restarting the node, the operating system may assign the PIDs listed in the PID files to other processes running on the node.

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

Mount agent limitations

The Mount agent has the following limitations:

- The Mount agent mounts a block device at only one mount point on a system. After a block device is mounted, the agent cannot mount another device at the same mount point.
- Mount agent does not support:
 - ext4 filesystem on SLES 11SP1, SLES 11SP2
 - ext4 filesystem configured on VxVM
 - xfs filesystem configured on VxVM

Share agent limitations

To ensure proper monitoring by the Share agent, verify that the `/var/lib/nfs/etab` file is clear upon system reboot. Clients in the Share agent must be specified as fully qualified host names to ensure seamless failover.

Driver requirements for DiskReservation agent

The VRTSvcsdr package ships the scsiutil utility. DiskReservation agent supports only those drivers supported by the scsiutil utility.

Application agent limitations

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

Limitations related to IMF

- IMF registration on Linux for “bind” file system type is not supported.
- In case of SLES11 SP1 and RHEL6.1:
 - IMF should not be enabled for the resources where the BlockDevice can get mounted on multiple MountPoints.
 - If FSType attribute value is nfs, then IMF registration for “nfs” file system type is not supported.

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

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

Workaround:

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

```
# cd /opt/VRTSagents/ha/bin/<ResourceType>
# ln -s /opt/VRTSvcs/bin/Script51Agent <ResourceType>Agent
```
- 2 Run the following command to update the AgentFile attribute based on value of VCS_HOME.

- If VCS_HOME is /opt/VRTSvcs:

```
# hatype -modify <ResourceType> AgentFile
/opt/VRTSvcs/bin/<ResourceType>/<ResourceType>Agent
```

- If VCS_HOME is /opt/VRTSagents/ha:

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

Limitations related to the VCS database agents

The Db2udb agent does not support DB2 9.8

The Db2udb agent does not support DB2 version 9.8. This is because DB2 9.8 is designed especially for DB2 PureScale feature and it doesn't support some of the features in DB2 9.7 temporarily.

See

<http://www.ibm.com/developerworks/data/library/techarticle/dm-0909db2whicheedition/> for more information.

DB2 RestartLimit value [1234959]

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

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

Limitation with intentional offline functionality of VCS agent for Oracle

The Oracle resource never faults after an intentional offline.

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

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

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

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

Security-Enhanced Linux is not supported on SLES distributions

VCS does not support Security-Enhanced Linux (SELinux) on SLES10 and SLES11. [1056433]

Systems in a cluster must have same system locale setting

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

VxVM site for the disk group remains detached after node reboot in campus clusters with fire drill [1919317]

When you bring the DiskGroupSnap resource online, the DiskGroupSnap agent detaches the site from the target disk group defined. The DiskGroupSnap agent invokes VCS action entry points to run VxVM commands to detach the site. These commands must be run on the node where the disk group is imported, which is at the primary site.

If you attempt to shut down the node where the fire drill service group or the disk group is online, the node goes to a LEAVING state. The VCS engine attempts to take all the service groups offline on that node and rejects all action entry point requests. Therefore, the DiskGroupSnap agent cannot invoke the action to reattach the fire drill site to the target disk group. The agent logs a message that the node is in a leaving state and then removes the lock file. The agent's monitor function declares that the resource is offline. After the node restarts, the disk group site still remains detached.

Workaround:

You must take the fire drill service group offline using the `hagrp -offline` command before you shut down the node or before you stop VCS locally.

If the node has restarted, you must manually reattach the fire drill site to the disk group that is imported at the primary site.

If the secondary node has crashed or restarted, you must manually reattach the fire drill site to the target disk group that is imported at the primary site using the following command: `/opt/VRTSvcs/bin/hares -action $targetres joindg -actionargs $fdsitename $is_fenced -sys $targetsys`.

Limitations with DiskGroupSnap agent [1919329]

The DiskGroupSnap agent has the following limitations:

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

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

System reboot after panic

If the VCS kernel module issues a system panic, a system reboot is required [293447]. The supported Linux kernels do not automatically halt (CPU) processing. Set the Linux "panic" kernel parameter to a value other than zero to forcibly reboot the system. Append the following two lines at the end of the `/etc/sysctl.conf` file:

```
# force a reboot after 60 seconds
kernel.panic = 60
```

Host on RHEV-M and actual host must match [2827219]

You must configure the host in RHEV-M with the same name as in the `hostname` command on a particular host. This is mandatory for RHEV Manager to be able to search the host by hostname.

Cluster Manager (Java console) limitations

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

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

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

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

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

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

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

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the `UseFence` attribute is set to the default, "None".

Using the KDE desktop

Some menus and dialog boxes on Cluster Manager (Java Console) may appear misaligned or incorrectly sized on a KDE desktop. To ensure the proper appearance and functionality of the console on a KDE desktop, use the Sawfish window manager. You must explicitly select the Sawfish window manager even if it is supposed to appear as the default window manager on a KDE desktop.

Limited support from Cluster Manager (Java console)

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

are already supported in VOM. However, Java console may continue to work as expected with features of releases prior to VCS 6.0.

Port change required to connect to secure cluster [2615068]

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

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

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

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

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

Stopping systems in clusters with I/O fencing configured

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

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

keys are removed from coordination points and data disks to prevent possible difficulties with subsequent cluster startup.

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

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

Limitations related to global clusters

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

Limitations related to VMware virtual environment

Symantec High Availability Configuration Wizard does not support nested mount points [2874775]

Symantec High Availability Configuration Wizard does not discover nested mount points. As a result, you cannot use the wizard to monitor applications that access a database via a nested mount point.

Workaround: Use VCS commands to configure the mount resources for monitoring nested mount points. Ensure that you set appropriate resource dependencies to configure the nested mounts. For more information, see the *Veritas Cluster Server Administrator's Guide*.

Symantec High Availability tab does not provide an option to remove a node from a VCS cluster

The **Symantec High Availability** tab does not support removal of a system from a VCS cluster. (2944997)

Workaround: You can use VCS commands to remove the system from the VCS cluster. You can also unconfigure the entire VCS cluster from the **Symantec High Availability** tab.

VMware fault tolerance does not support adding or removing non-shared disks between virtual machines

VMware fault tolerance does not support adding or removing of non-shared disks between virtual machines. During a failover, disks that contain application data cannot be moved to alternate failover systems. Applications that are being monitored, thus cannot be brought online on the failover systems.

Veritas Storage Foundation and High Availability software limitations

This section describes the Veritas Storage Foundation and High Availability software limitations in 6.0.3 and 6.0.1.

Replication software limitations

The following are replication software limitations in this release of Veritas Storage Foundation.

VVR Replication in a shared environment

Currently, replication support is limited to 8-node cluster applications.

VVR IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.
- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.
- A replication configuration in which an IPv6 address is specified for the `local_host` attribute of a primary RLINK and an IPv4 address is specified for the `remote_host` attribute of the same RLINK.
- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.

- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 6.0 and the prior major releases of Storage Foundation (5.1 and 5.1SP1). Replication between versions is supported for disk group versions 150, 160, and 170 only. Both the Primary and Secondary hosts must be using a supported disk group version.

Softlink access and modification times are not replicated on RHEL5 and SLES10 for VFR jobs

When running a file replication job on RHEL5 and SLES10, softlink access and modification times are not replicated.

Thin reclamation requests are not redirected even when the ioship policy is enabled (2755982)

Reclamation requests fail from nodes that do not have local connectivity to the disks, even when the ioship policy is enabled. Reclamation I/Os are not redirected to another node.

Veritas Operations Manager does not support disk, disk group, and volume state information related to CVM I/O shipping feature (2781126)

The Veritas Operations Manager (VOM) does not support disk, disk group, and volume state information related to the I/O shipping feature introduced in this release of Cluster Volume Manager. New states such as Ifailed, lmissing or LDISABLED are introduced when I/O shipping is active because of storage disconnectivity.

Veritas Storage Foundation Cluster File System High Availability software limitations

This section describes the Veritas Storage Foundation Cluster File System High Availability software limitations in 6.0.3 and 6.0.1.

cfsmntadm command does not verify the mount options (2078634)

You must confirm the mount options are correct which are then passed to the `cfsmntadm` command. If the mount options are not correct, the mount fails and the CFSSMount resource will not come online. You can check the VCS engine log file for any mount failure messages.

Obtaining information about mounted file system states (1764098)

For accurate information about the state of mounted file systems on Linux, refer to the contents of `/proc/mounts`. The `mount` command may or may not reference this source of information depending on whether the regular `/etc/mtab` file has been replaced with a symbolic link to `/proc/mounts`. This change is made at the discretion of the system administrator and the benefits are discussed in the `mount` online manual page. A benefit of using `/proc/mounts` is that changes to SFCFS mount options are accurately displayed for all nodes.

Stale SCSI-3 PR keys remain on disk after stopping the cluster and deporting the disk group

When all nodes present in the SF cluster are removed from the cluster, the SCSI-3 Persistent Reservation (PR) keys on the data disks may not get preempted. As a result, the keys may be seen on the disks after stopping the cluster or after the nodes have booted up. The residual keys do not impact data disk fencing as they will be reused or replaced when the nodes rejoin the cluster. Alternatively, the keys can be cleared manually by running the `vxfcntlclearpre` utility.

For more information on the `vxfcntlclearpre` utility, see the *Veritas Storage Foundation Administrator's Guide*.

Veritas Storage Foundation for Oracle RAC software limitations

This section describes the Veritas Storage Foundation for Oracle RAC software limitations in 6.0.3 and 6.0.1.

Oracle Clusterware/Grid Infrastructure installation fails if the cluster name exceeds 14 characters

Setting the cluster name to a value that exceeds 14 characters during the installation of Oracle Clusterware/Grid Infrastructure causes unexpected cluster membership issues. As a result, the installation may fail.

Workaround: Restart the Oracle Clusterware/Grid Infrastructure installation and set the cluster name to a value of maximum 14 characters.

SELinux supported in disabled and permissive modes only

SELinux (Security Enhanced Linux) is supported only in "Disabled" and "Permissive" modes. After you configure SELinux in "Permissive" mode, you may see a few messages in the system log. You may ignore these messages.

Policy-managed databases not supported by CRSResource agent

The CRSResource agent supports only admin-managed database environments in this release. Policy-managed databases are not supported.

Health checks may fail on clusters that have more than 10 nodes

If there are more than 10 nodes in a cluster, the health check may fail with the following error:

```
vxgettext ERROR V-33-1000-10038  
Arguments exceed the maximum limit of 10
```

The health check script uses the `vxgettext` command, which does not support more than 10 arguments.[2142234]

Cached ODM not supported in SF environments

Cached ODM is not supported for files on Veritas local file systems and on Cluster File System.

Veritas Storage Foundation for Sybase ASE CE software limitations

This section describes the Veritas Storage Foundation for Sybase ASE CE software limitations in 6.0.3 and 6.0.1.

Only one Sybase instance is supported per node

In a Sybase ASE CE cluster, SF Sybase CE supports only one Sybase instance per node.

SF Sybase CE is not supported in the Campus cluster environment

SF Sybase CE does not support the Campus cluster. SF Sybase CE supports the following cluster configurations. Depending on your business needs, you may choose from the following setup models:

- Basic setup
- Secure setup
- Central management setup
- Global cluster setup

See the *Installation Guide* for more information.

Hardware-based replication technologies are not supported for replication in the SF Sybase CE environment

You can use Veritas Volume Replicator (VVR), which provides host-based volume replication. Using VVR you can replicate data volumes on a shared disk group in SF Sybase CE. Hardware-based replication is not supported at this time.

SF Sybase CE installation is not supported by Web installer

SF Sybase CE does not support the Web-based installer at this time. You can use one of the following methods to install and configure SF Sybase CE.

- Interactive installation and configuration using the script-based installer
- Silent installation using the response file

See the *Installation Guide* for more information.

Symantec VirtualStore software limitations

This section describes the Symantec VirtualStore software limitations in 6.0.3 and 6.0.1.

VMware vSphere extension for VirtualStore limitations

The following are the software limitations for VMware vSphere extension for VirtualStore that are known in this release.

F5 usage is not supported for wizard refreshing (2362940)

F5 usage is not supported for wizard refreshing.

Workaround

To get new or refreshed data, it is important to restart the wizard and not use the F5 key.

Virtual machines with VMware Snapshots cannot be used as golden images (2514969)

Any virtual machine (or template) which has VMware Snapshots stored, cannot be used as a golden image for making clones with the FileSnap wizard. To use such virtual machines (or templates), first delete the Snapshots, then use the FileSnap wizard.

Documentation errata

There are no documentation errata updates.

List of RPMs

This section lists the RPMs for 6.0.3.

Note: You can also view the following list using the `installmr` command, type:

```
./installmr -listpatches
```

Table 1-25 RPMs for Red Hat Enterprise Linux 5

Name	Version	Arch	Size in bytes
VRTSsamf	6.0.300.000	x86_64	1513407
VRTScavf	6.0.300.000	i386	198544
VRTSsdbed	6.0.300.000	x86_64	36606448
VRTSperl	5.14.2.8	x86_64	17923439
VRTSsfcp1601	6.0.300.000	noarch	810265
VRTSspt	6.0.300.000	noarch	21716879
VRTSvcsc	6.0.300.000	i686	64685562
VRTSvcscag	6.0.300.000	i686	16280232
VRTSvcsea	6.0.300.000	i686	3867611
VRTSvcsvmw	6.0.300.000	i686	10833977
VRTSvxfen	6.0.300.000	x86_64	490848
VRTSvxfs	6.0.300.000	x86_64	11014560
VRTSvxvm	6.0.300.000	x86_64	32590287

Table 1-26 RPMs for Red Hat Enterprise Linux 6

Name	Version	Arch	Size in bytes
VRTSsamf	6.0.300.000	x86_64	1030388
VRTScavf	6.0.300.000	i386	179464

Table 1-26 RPMs for Red Hat Enterprise Linux 6 (continued)

Name	Version	Arch	Size in bytes
VRTSdbed	6.0.300.000	x86_64	36606448
VRTSperl	5.14.2.8	x86_64	14331096
VRTSsfcp1601	6.0.300.000	noarch	810265
VRTSspt	6.0.300.000	noarch	21716879
VRTSvcsc	6.0.300.000	i686	48763888
VRTSvcscag	6.0.300.000	i686	12401472
VRTSvcscsea	6.0.300.000	i686	3431808
VRTSvcscvmw	6.0.300.000	i686	9491728
VRTSvxfen	6.0.300.000	x86_64	967468
VRTSvxfs	6.0.300.000	x86_64	7832224
VRTSvxvm	6.0.300.000	x86_64	22186248

Table 1-27 RPMs for SUSE Linux Enterprise Server 10

Name	Version	Arch	Size in bytes
VRTSsamf	6.0.300.000	x86_64	5824813
VRTScavf	6.0.300.000	i386	166361
VRTSdbed	6.0.300.000	x86_64	36433465
VRTSperl	5.14.2.8	x86_64	15456461
VRTSsfcp1601	6.0.300.000	noarch	810265
VRTSspt	6.0.300.000	noarch	21716879
VRTSvcsc	6.0.300.000	i586	59787590
VRTSvcscag	6.0.300.000	i586	15429020
VRTSvcscsea	6.0.300.000	i586	3642127
VRTSvcscvmw	6.0.300.000	i586	10075354
VRTSvxfen	6.0.300.000	x86_64	2406046
VRTSvxfs	6.0.300.000	x86_64	11178846

Table 1-27 RPMs for SUSE Linux Enterprise Server 10 (*continued*)

Name	Version	Arch	Size in bytes
VRTSvxvm	6.0.300.000	x86_64	31075897

Table 1-28 RPMs for SUSE Linux Enterprise Server 11

Name	Version	Arch	Size in bytes
VRTSsamf	6.0.300.000	x86_64	1516366
VRTScavf	6.0.300.000	i386	171823
VRTSdbed	6.0.300.000	x86_64	36433465
VRTSglm	6.0.100.100	x86_64	129646
VRTSgms	6.0.100.100	x86_64	31258
VRTSodm	6.0.100.100	x86_64	287009
VRTSperl	5.14.2.8	x86_64	13977816
VRTSsfpci601	6.0.300.000	noarch	810265
VRTSspt	6.0.300.000	noarch	21716879
VRTSvcsc	6.0.300.000	i686	48341306
VRTSvcscag	6.0.300.000	i686	12583704
VRTSvcsea	6.0.300.000	i686	2669668
VRTSvcsvm	6.0.300.000	i686	9063480
VRTSvxfen	6.0.300.000	x86_64	1012297
VRTSvxfs	6.0.300.000	x86_64	7858125
VRTSvxvm	6.0.300.000	x86_64	21351744

Downloading the 6.0.3 archive

The patches that are included in the 6.0.3 release are available for download from the Symantec website. After downloading the 6.0.3 rolling patch, use `gunzip` and `tar` commands to uncompress and extract it.

For the 6.0.3 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

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

Installing the products for the first time

This chapter includes the following topics:

- [Installing the Veritas software using the script-based installer](#)
- [Installing Veritas software using the Web-based installer](#)

Installing the Veritas software using the script-based installer

This section describes how to install a 6.0.3 Veritas Storage Foundation and High Availability Solutions product for the first time on a host. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 6.0.1 *Installation Guide* and *Release Notes* for your product for more information.

If you already have 6.0.1 or 6.0.2 installed, you must upgrade to 6.0.3.

See [“Upgrading to 6.0.3”](#) on page 166.

To install the Veritas software for the first time

- 1 Download Storage Foundation and High Availability Solutions 6.0.1 from <http://fileConnect.symantec.com>.
- 2 Extract the tar ball into a directory called `/tmp/sfha601`.
- 3 Check <http://sort.symantec.com/patches> to see if there are any patches available for the 6.0.1 Installer. Download applicable P-patches and extract them to the `/tmp` directory.

- 4 Change the directory to `/tmp/sfha601`:

```
# cd /tmp/sfha601
```
- 5 Run the installer to install SFHA 6.0.1. See the Installation Guide for instructions on installing the 6.0.1 version of this product.

```
# ./installer -require complete_path_to_601_installer_patch
```
- 6 Download SFHA 6.0.3 from <http://sort.symantec.com/patches>.
- 7 Extract it to a directory called `/tmp/sfha603`.
- 8 Check <http://sort.symantec.com/patches> to see if there are patches available for the 6.0.3 installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 9 Change the directory to `/tmp/sfha603`:

```
# cd /tmp/sfha603
```
- 10 Invoke the `installmr` script to install 6.0.3:

```
# installmr -require complete_path_to_603_installer_patch
```

See “[About the installmr script](#)” on page 10.
- 11 If you did not configure the product after the 6.0.1 installation, the installer prompts you to configure the product during RP installation. If you do not want to configure the product now, answer `n` when prompted. To configure the product in the future, run the product installation script from the 6.0.1 installation media or from `/opt/VRTS/install` directory with the `-configure` option

On SUSE 10 SP3, do not configure 5.1 product until 6.0.3 product is installed. As support for SUSE 10 SP3 was just recently released.

Installing Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 6.0.3 using the Web-based installer. For detailed instructions on how to install 6.0.1 using the Web-based installer, follow the procedures in the 6.0.1 Installation Guide and Release Notes for your products.

See “[Upgrading to 6.0.3](#)” on page 166.

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprt1wid`, on the installation server:

```
# ./webinstaller start
```

The `webinstaller` script displays a URL.

- 2 Start the Web browser on the system from which you want to perform the installation.
- 3 Navigate to the URL displayed from step 1.
- 4 The browser may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

- 5 When prompted, enter `root` and `root`'s password of the installation server.

Obtaining a security exception on Mozilla Firefox

You may need to get a security exception on Mozilla Firefox.

To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **Add Exception** button.
- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.
- 5 Click **Confirm Security Exception** button.
- 6 Enter `root` in User Name field and `root` password of the web server in the Password field.

Installing 6.0.3 with the Veritas Web-based installer

This section describes installing SF with the Veritas Web-based installer.

To install SF

- 1 The 6.0.1 version of the Veritas product must be installed before upgrading to 6.0.3.
See [“Prerequisites for upgrading to 6.0.3”](#) on page 164.
- 2 On the **Select a task and product** page, select **Install 6.0.3** from the **Task** drop-down list, and click **Next**.
- 3 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
- 4 You have the option to let the installer configure SSH or RSH communications between the systems. If you choose to allow this configuration, select the shell and provide the root passwords for each system.
- 5 After the validation completes successfully, click **Next** to install 6.0.3 patches on the selected system.
- 6 Select the checkbox to specify whether you want to send your installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future?

Click **Finish**.

Upgrading to 6.0.3

This chapter includes the following topics:

- [Prerequisites for upgrading to 6.0.3](#)
- [Downloading required software to upgrade to 6.0.3](#)
- [Supported upgrade paths](#)
- [About using the installer to upgrade when the root disk is encapsulated](#)
- [Upgrading to 6.0.3](#)
- [Verifying software versions](#)

Prerequisites for upgrading to 6.0.3

The following list describes prerequisites for upgrading to the 6.0.3 release:

- For any product in the Veritas Storage Foundation stack, you must have the 6.0.1 installed before you can upgrade that product to the 6.0.3 release.
- Each system must have sufficient free space to accommodate patches.
- For RHEL6, un-install the VRTSaslapm package if it is already installed.
- The full list of prerequisites can be obtained by running `./installmr -precheck`.
- Make sure to download the latest patches for the installer.
See [“Downloading required software to upgrade to 6.0.3”](#) on page 164.

Downloading required software to upgrade to 6.0.3

This section describes how to download the latest patches for the installer.

To download required software to upgrade to 6.0.3

- 1 Download SFHA 6.0.3 from <http://sort.symantec.com/patches>.
- 2 Extract it to a directory such as `/tmp/sfha603`.
- 3 Check <http://sort.symantec.com/patches> to see if there are patches available for the 6.0.3 installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 4 When you run the `installmr` script, use the `-require` option and specify the location where you downloaded the 6.0.3 installer patches.

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 6.0.1 to 6.0.3
- 6.0.2 to 6.0.3

Note: If you want to deploy VCS on virtual machines in a VMware environment, you must first upgrade to VCS 6.0.2, and then upgrade to VCS 6.0.3. This step-wise upgrade is required because VCS 6.0.2 packages include the `VRTSvcsvmw` package. This package is critical as it enables you to configure and administer VCS from the VMware vSphere Client.

Upgrading using VMware vSphere Client menu:

If you have installed Veritas Cluster Server 6.0.2, you can install VCS maintenance releases and updated application-agents by using the VMware vSphere Client menu.

You can use this installation option to upgrade to the following releases:

- VCS Maintenance Release 6.0.3 and later
- Agent Pack 4Q2012 and later

For more information, see <http://www.symantec.com/docs/TECH200772>.

About using the installer to upgrade when the root disk is encapsulated

When you use the installer to upgrade from a previous version of SF and the system where you plan to upgrade has an encapsulated root disk, you may have to unencapsulate it.

Table 3-1 Upgrading using installer when the root disk is encapsulated

Starting version	Ending version	Action required
5.0 MP3	6.0.3	You need to unencapsulate the root disk. The installer exits.
5.1 or 5.1 RPx	6.0.3	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 SP1 or 5.1 SP1 RPx	6.0.3	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0 or 6.0 RP1	6.0.3	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Upgrading to 6.0.3

This section describes how to upgrade from 6.0.1 or 6.0.2 to 6.0.3 on a cluster or a standalone system.

- [Performing a full upgrade to 6.0.3 on a cluster](#)
Use the procedures to perform a full upgrade to 6.0.3 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System High Availability (SFCFSHA), Veritas Storage Foundation for Oracle RAC (SFRAC), Veritas Storage Foundation for Sybase ASE CE (SF Sybase CE), or Symantec VirtualStore (SVS) installed and configured.
- [Upgrading to 6.0.3 on a standalone system](#)
Use the procedure to upgrade to 6.0.3 on a system that has SF installed.
- [Upgrading to 6.0.3 on a system that has encapsulated boot disk](#)
Use the procedure to upgrade your Veritas product on a system that has encapsulated boot disk
- [Performing a rolling upgrade using the `installmr` script](#)
Use the procedure to upgrade your Veritas product with a rolling upgrade.
- [Performing a phased upgrade to SFCFSHA and SFRAC](#)
Use the procedure to perform a phase upgrade using the `installmr` script.
- [Upgrading the operating system](#)
Use the procedure to upgrade the operating system.

See [“Installing the Veritas software using the script-based installer”](#) on page 160.

Performing a full upgrade to 6.0.3 on a cluster

Performing a full upgrade on a cluster requires stopping cluster failover functionality during the entire procedure. However, if you use Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) and Cluster Volume Manager (CVM), the SFCFSHA and CVM services remain available.

Depending on your cluster’s configuration, select one of the following procedures to upgrade to 6.0.3:

- [Performing a full upgrade to 6.0.3 on a Veritas Cluster Server](#)
- [Performing a full upgrade to 6.0.3 on an SFHA cluster](#)
- [Performing a full upgrade to 6.0.3 on an SFCFSHA cluster](#)
- [Performing a full upgrade to 6.0.3 on an SF Oracle RAC cluster](#)
- [Performing a full upgrade to 6.0.3 on an SF Sybase ASE CE cluster](#)
See [“Downloading required software to upgrade to 6.0.3 ”](#) on page 164.

Performing a full upgrade to 6.0.3 on a Veritas Cluster Server

The following procedure describes performing a full upgrade on a Veritas Cluster Server (VCS) cluster.

Note: You need to make sure that IPv4RouteOptions attribute is configured for MultiNICA resources, otherwise network connection may be interrupted.

To upgrade VCS

- 1 Make sure you have downloaded the latest software required for the upgrade.
See [“Downloading required software to upgrade to 6.0.3 ”](#) on page 164.
- 2 Log in as superuser.

Note: Upgrade the Operating System and reboot the systems if required.

See [“System requirements”](#) on page 17.

- 3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 6.0.3 rolling patch binaries, change to the directory that contains the `installmr` script. Start the pre-upgrade check:

```
# ./installmr -precheck node1 node2 ... nodeN
```

See [“About the installmr script”](#) on page 10.

- 4 Resolve any issues that the precheck finds.
- 5 Start the upgrade:

```
# ./installmr node1 node2 ... nodeN
```

See [“About the installmr script”](#) on page 10.

- 6 After the upgrade, review the log files for any issues.

Performing a full upgrade to 6.0.3 on an SFHA cluster

The following procedure describes performing a full upgrade on an SFHA and VCS cluster.

To perform a full upgrade to 6.0.3 on an SFHA cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
See [“Downloading required software to upgrade to 6.0.3”](#) on page 164.
- 2 Log in as superuser.
- 3 From the directory that contains the extracted and untarred 6.0.3 rolling patch binaries, change to the directory that contains the `installmr` script. Check the readiness of the nodes where you plan to upgrade. Start the pre-upgrade check:

```
# ./installmr -precheck node1 node2... nodeN
```

where `node1`, `node2` and `nodeN` are nodes to be upgraded.

- 4 Start the upgrade:

```
# ./installmr node1 node2... nodeN
```

where `node1`, `node2` and `nodeN` are nodes to be upgraded.

Performing a full upgrade to 6.0.3 on an SFCFSHA cluster

The following procedure describes performing a full upgrade on an SFCFSHA cluster.

To perform a full upgrade to 6.0.3 on an SFCFSHA cluster

- 1 Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 6.0.3”](#) on page 164.

- 2 Log in as superuser.

- 3 On each node, enter the following command to check if any Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

If any Storage Checkpoints are mounted, on each node in the cluster unmount all Storage Checkpoints.

```
# umount /checkpoint_name
```

- 4 On each node, enter the following command to check if any VxFS file systems are mounted:

```
# mount | grep vxfs
```

If any VxFS file systems are present, on each node in the cluster unmount all of the VxFS file systems:

```
# umount /filesystem
```

Note: If file system is CFS monted then use `cfsumount` command.

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
 - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
 - Use the `vrxvg stop` command to stop each RVG individually:

```
# vxrvvg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

Caution: To avoid data corruption, do not proceed until all RLINKs are up-to-date.

- 6 Stop activity to all VxVM volumes.

For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes

- 7 Stop all VxVM volumes by entering the following command for each disk group on master node:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

- 8 If required, apply the OS kernel patches.
- 9 From the directory that contains the extracted and untarred 6.0.3 rolling patch binaries, change to the directory that contains the `installmr` script. Start the upgrade.

```
# ./installmr node1 node2
```

where *node1* and *node2* are nodes which are to be upgraded.

- 10 After all nodes in the cluster are upgraded, the processes will be restarted automatically. Should there be any problem, the `installmr` script will ask you to reboot the system. Then the application failover capability will be available.
- 11 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.
- 12 Restart all the volumes by entering the following command on the master node, for each disk group:

```
# vxvol -g diskgroup startall
```

- 13 If you stopped any RVGs in step 5, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

- 14 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

- 15 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

Performing a full upgrade to 6.0.3 on an SF Oracle RAC cluster

The following procedure describes performing a full upgrade on an SF for Oracle RAC cluster.

To upgrade to 6.0.3 on an SF Oracle RAC cluster

- 1 Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 6.0.3”](#) on page 164.

- 2 Log in as superuser.

- 3 Verify that `/opt/VRTSvcs/bin` is in your `PATH` so that you can execute all product commands.

- 4 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 5 Enter the following command to freeze HA service group operations on each node:

```
# hasys -freeze -persistent nodename
```

- 6 Make the configuration read-only:

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

- 7 If Oracle Clusterware is not controlled by VCS, enter the following command on each node of the cluster to stop Oracle Clusterware:

```
# $CRS_HOME/bin/crsctl stop crs
```

- 8 Stop VCS.

```
# hastop -all
```

- 9 If required, apply the OS kernel patches.
See “[System requirements](#)” on page 17.
See *Oracle’s* documentation for the procedures.
- 10 From the directory that contains the extracted and untarred 6.0.3 rolling patch binaries, change to the directory that contains the installmr script. If ssh key authentication is configured then enter:

```
# ./installmr node1 node2
```

If ssh is not configured then enter:

```
# ./installmr -rsh node1 node2
```

where `node1` and `node2` are nodes which are to be upgraded.

- 11 Follow the instructions from the installer. If there is some module load/unload issue, reboot all of the nodes of the cluster.

```
# /sbin/shutdown -r now
```

- 12 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.
- 13 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 14 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.
- 15 Relink the SF Oracle RAC libraries with Oracle.

Refer to *Veritas Storage Foundation for Oracle RAC 6.0.1 or later Installation and Configuration Guide* for more information.

- 16 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 17 Enter the following command on each node to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

- 18 Make the configuration read-only:

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

- 19 Enter the following command on each node to take service groups online:

```
# hagrpl -online service_group -sys nodename
```

- 20 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 21 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

- 22 If Oracle Clusterware is not controlled by VCS, enter the following command on each node to start Oracle Clusterware.

```
# $CRS_HOME/bin/crsctl start crs
```

- 23 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

Performing a full upgrade to 6.0.3 on an SF Sybase ASE CE cluster

The following procedure describes performing a full upgrade on an SF for Oracle RAC cluster.

To upgrade to 6.0.3 on an SF Sybase ASE CE cluster

- 1 Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 6.0.3”](#) on page 164.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTSvcs/bin` is in your `PATH` so that you can execute all product commands.
- 4 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 5 Enter the following command to freeze HA service group operations on each node:

```
# hasys -freeze -persistent nodename
```

- 6 Make the configuration read-only:

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

- 7 Take the Sybase database group offline:

```
# hagrps -offline sybasece -sys node_name
```

- 8 Stop VCS.

```
# hastop -all
```

- 9 If required, apply the OS kernel patches.

See [“System requirements”](#) on page 17.

See *Sybase* documentation for the procedures.

- 10 From the directory that contains the extracted and untarred 6.0.3 rolling patch binaries, change to the directory that contains the installmr script. If ssh key authentication is configured then enter:

```
# ./installmr node1 node2
```

If ssh is not configured then enter:

```
# ./installmr -rsh node1 node2
```

where *node1* and *node2* are nodes which are to be upgraded.

- 11 Follow the instructions from the installer. If there is some module load/unload issue, reboot all of the nodes of the cluster.

```
# /sbin/shutdown -r now
```

- 12 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.

- 13 Manually mount the VxFS and CFS file systems that are not managed by VCS.

- 14 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.

- 15 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 16 Enter the following command on each node to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

- 17 Make the configuration read-only:

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

- 18 Enter the following command on each node to take service groups online:

```
# hagr -online service_group -sys nodename
```

- 19 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 20 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

- 21 Bring the Sybasece resource group online.

```
# hagr -online sybasece -sys node_name
```

- 22 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

Upgrading to 6.0.3 on a standalone system

You can use this procedure to upgrade on a standalone system that runs SF.

To upgrade to 6.0.3 on a standalone system

- 1 Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 6.0.3”](#) on page 164.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 4 If required, apply the OS kernel patches.
- 5 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

- 6 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

- 7 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
 - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
 - Use the `vrxvg stop` command to stop each RVG individually:

```
# vxrvg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

Caution: To avoid data corruption, do not proceed until all RLINKs are up-to-date.

- 8 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

- 9 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

- 10 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 11 Navigate to the folder that contains the installation program. Run the `installmr` script:

```
# ./installmr nodename
```

- 12 If necessary, reinstate any missing mount points in the `/etc/filesystems` file.

- 13 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 14 If you stopped any RVGs in step 7, restart each RVG:

```
# vxrvrg -g diskgroup start rvg_name
```

- 15 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
```

```
# mount /checkpoint_name
```

- 16 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

Upgrading to 6.0.3 on a system that has encapsulated boot disk

You can use this procedure to upgrade to 6.0.3 on a system that has encapsulated boot disk.

Note: Upgrading with encapsulated boot disk from 6.0.1 to 6.0.3 requires multiple reboots.

To upgrade to 6.0.3 on a system that has encapsulated boot disk

- 1 Manually unmount file systems and stop open volumes.
- 2 If required, manually break the mirror.
- 3 Upgrade to 6.0.3 using `installmr` command.
- 4 After upgrading, reboot the system to have the new VM drivers take effect.
- 5 If the mirrors of boot disk are split manually in step 2, re-join the mirrors manually when systems reboot after upgrading to 6.0.3.

Performing a rolling upgrade using the `installmr` script

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

- [About rolling upgrades](#)
- [Prerequisites for a rolling upgrades](#)
- [Performing a rolling upgrade using the installer](#)

About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime.

Rolling upgrades take two discrete phases. In the first, you upgrade the kernel packages with exception of VCS packages and agent packages. In the second, you upgrade the non-kernel packages, which are VCS packages and agents packages.

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System and High Availability
- Storage Foundation for Oracle RAC

- Symantec VirtualStore
- Storage Foundation for Sybase CE

You can perform a rolling upgrade from 6.0.1 to 6.0.3.

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrade.
- Split up your cluster into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.
- For SF Oracle RAC, stop Oracle Clusterware before upgrading Kernel packages on any node, which is not managed by VCS.
- Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 6.0.3”](#) on page 164.

Limitation: During VCS and agents upgrade, you must bring down the application High Availability (HA) for several minutes. This does not affect the application running on the cluster. You can restore the application's high availability after VCS and the agent packages are upgraded.

Performing a rolling upgrade using the installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

Performing a rolling upgrade on kernel packages for VCS, SFHA, SFSYBASECE, SVS and SFCFSHA: phase 1

Note that in the following instructions a sub-cluster can represent one or more nodes in a full cluster, but is represented by nodeA.

To perform the rolling upgrade on kernel packages: phase 1

- 1 Stop all the applications that access volumes which are not under VCS control in the sub-cluster to be upgraded.
- 2 Unmount all the file systems managed by SF which are not under VCS control in the sub-cluster to be upgraded.

- 3 On the first sub-cluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

```
# ./installmr -rollingupgrade_phase1 nodeA
```

- 4 Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 5 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 6 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.
- 7 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.
- 8 After the installer shows the package list, it detects if there are online failover service groups on the nodes to be upgraded. If there are online failover service groups, the installer prompts you to do one of the following:
 - Manually switch service groups
 - Use the CPI to automatically switch service groupsThe downtime is the time that it normally takes for the service group's failover.
- 9 After switching failover service group, installer detects if there are online parallel service groups on the node to be upgraded. If there are online parallel service groups, installer will ask user to use CPI to automatically offline parallel service groups.
- 10 The installer stops relevant processes, uninstalls old kernel RPM, and installs the new RPM. When prompted, enable replication or global cluster capabilities, if required, and register the software.
- 11 The installer performs the upgrade configuration and re-starts processes.
If some processes fail to start, you may need to reboot the nodes and manually check the cluster's status.
- 12 If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step [13](#).
- 13 Before you proceed to phase 2, complete step [1](#) to step [11](#) on the second subcluster.

Performing a rolling upgrade on non-kernel packages for VCS, SFHA, SFSYBASECE, SVS and SFCFSHA: phase 2

In this phase installer installs all non-kernel patches on all the nodes in cluster and restarts VCS cluster.

To perform the rolling upgrade on non-kernel packages: phase 2

- 1 When phase 1 of the rolling upgrade completes, begin phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.
- 2 Reboot the nodes if the installer requires.
- 3 The installer determines the remaining RPM to upgrade. Press **Enter** to continue.
- 4 The installer stops Veritas Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs prechecks, uninstalls old RPM, and installs the new RPM. It performs post-installation tasks, and the configuration for the upgrade.
- 5 Type **y** or **n** to help Symantec improve the automated installation.
- 6 If you have network connection to the Internet, the installer checks for updates.
- 7 A prompt message appears to ask if the user would like to read the summary file. You can choose **y** if you want to read the install summary file.
- 8 Verify the cluster's status:

```
# hastatus -sum
```

- 9 If you want to upgrade CP server systems that use VCS or SFHA to 6.0.3, make sure that you upgraded all application clusters to version 6.0.3. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Performing a rolling upgrade on kernel packages for SFRAC: phase 1

Note that in the following instructions that a sub-cluster can represent one or more nodes in a full cluster, but is represented by nodeA.

To perform the rolling upgrade on kernel packages: phase 1

- 1 On the first subcluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

```
# ./installmr -rollingupgrade_phase1 nodeA
```

- 2 Note that if the boot disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 3 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 4 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.
- 5 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.
- 6 After the installer shows the package list, it detects if there are online failover service groups on the nodes to be upgraded. If there are online failover service groups, the installer prompts you to do one of the following:
 - Manually switch service groups
 - Use the CPI to automatically switch service groups

The downtime is the time that it normally takes for the service group's failover.

- 7 After switching failover service group, installer detects if there are online parallel service groups on the node to be upgraded. If there are online parallel service groups, installer will ask user to use CPI to automatically offline parallel service groups.
- 8 The installer stops relevant processes, uninstalls old kernel RPM, and installs the new RPM. When prompted, enable replication or global cluster capabilities, if required, and register the software.
- 9 The installer performs the upgrade configuration and re-starts processes.

If some processes fail to start, you may need to reboot the nodes and manually check the cluster's status.

Note: The Oracle service group is offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically. The service group is started later in the process.

- 10 In case of SFRAC, refer to the “Performing post-upgrade Tasks” section to relink Oracle RAC libraries with SF Oracle RAC from 6.0.1 Installation and Configuration guide.
- 11 Bring the Oracle database service group online.
 - If VCS manages the Oracle database:


```
# hagrps -online oracle_group -sys node_name
```
 - If VCS does not manage the Oracle database:


```
# srvctl start database -d db_name
```
- 12 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 13 Start all applications that VCS does not manage. Use native application commands to start the applications.
- 14 If the boot disk is encapsulated, reboot the first sub-cluster's system.
- 15 Before you proceed to phase 2, complete step 1 to 14 on the second subcluster.
- 16
 - If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online when VCS starts:


```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```
 - If VCS does not manage the Oracle database, change the management policy for the database to automatic:


```
$ srvctl modify database -d db-name -y AUTOMATIC
```
- 17 Migrate the SFDB repository database.

Performing a rolling upgrade on non-kernel packages for SFRAC: phase 2

In this phase installer installs all non-kernel s on all the nodes in cluster and restarts VCS cluster.

To perform the rolling upgrade on non-kernel packages: phase 2

- 1 When phase 1 of the rolling upgrade completes, begin phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.
- 2 Reboot the nodes if the installer requires.

- 3 The installer determines the remaining RPM to upgrade. Press **Enter** to continue.
- 4 The installer stops Veritas Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.
 The installer performs prechecks, uninstalls old RPM, and installs the new RPM. It performs post-installation tasks, and the configuration for the upgrade.
- 5 Type **y** or **n** to help Symantec improve the automated installation.
- 6 If you have network connection to the Internet, the installer checks for updates.
- 7 A prompt message appears to ask if the user would like to read the summary file. You can choose **y** if you want to read the install summary file.
- 8 Verify the cluster's status:

```
# hastatus -sum
```
- 9 If you want to upgrade CP server systems that use VCS or SFHA to 6.0.3, make sure that you upgraded all application clusters to version 6.0.3. Then, upgrade VCS or SFHA on the CP server systems.
 For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Performing a phased upgrade to SFCFSHA and SFRAC

This section describes how to perform a phased upgrade to SFCFSHA and SFRAC.

Performing a phased upgrade of SFCFSHA

Performing a phased upgrade involves the following tasks:

- Preparing the first subcluster
- Upgrading the SFCFSHA stack on the first subcluster
- Preparing the second subcluster
- Activating the first subcluster
- Upgrading the operating system on the second subcluster
- Upgrading the second subcluster
- Finishing the phased upgrade

Before you start the upgrade on the first half of the cluster, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade. Note that your applications have downtime during this procedure.

Planning for a phased upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group. Some basic guidelines follow:

- Split the cluster in half. If the cluster has an odd number of nodes, calculate $(n+1)/2$, and start the upgrade with the even number of nodes.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.

Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules. Also, do not add or remove service groups to any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.
- You can perform a phased upgrade when the root disk is encapsulated.

Preparing the first subcluster

To move the service groups to the second subcluster

- 1 Switch failover groups from the first half of the cluster to one of the nodes in the second half of the cluster. In this procedure, `sys1` is a node in the first half of the cluster and `sys4` is a node in the second half of the cluster. Enter the following:

```
# hagrps -switch failover_group -to sys4
```

- 2 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the applications.

- 3 On the first half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS.

```
# mount | grep vxfs
```

Verify that no processes use the VxFS or CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS or CFS mount point with the mechanism provided by the application.

Unmount the VxFS or CFS file system. Enter the following:

```
# umount /mount_point
```

- 4 On the first half of the cluster, bring all the VCS service groups offline including CVM group. Enter the following:

```
# hagrps -offline group_name -sys sys1
```

When the CVM group becomes OFFLINE, all the parallel service groups such as the CFS file system will also become OFFLINE on the first half of the cluster nodes.

- 5 Verify that the VCS service groups are offline on all the nodes in first half of the cluster. Enter the following:

```
# hagrps -state group_name
```

- 6 Freeze the nodes in the first half of the cluster. Enter the following:

```
# haconf -makerw
```

```
# hasys -freeze -persistent sys1
```

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

- 7 If I/O fencing is enabled, then on each node of the first half of the cluster, change the contents of the `/etc/vxfenmode` file to configure I/O fencing in disabled mode. Enter the following:

```
# cp /etc/vxfen.d/vxfenmode_disabled /etc/vxfenmode
# cat /etc/vxfenmode
[root@swlx08 ~]# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3          - use scsi3 persistent reservation disks
# customized    - use script based customized fencing
# disabled      - run the driver but don't do any actual fencing
#
vxfen_mode=disabled
```

- 8 If the cluster-wide attribute `UseFence` is set to `SCSI3`, then reset the value to `NONE` in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.
- 9 Verify that only GAB ports a, b, d and h are open. Enter the following:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port h gen ada40f membership 01
```

Do not stop VCS. Port h should be up and running.

- 10 In the first half of the cluster, stop all VxVM and CVM volumes. Enter the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open. Enter the following:

```
# vxprint -Aht -e v_open
```

- 11 On first half of the cluster, upgrade the operating system on all the nodes, if applicable. For instructions, see the upgrade paths for the operating system.

Upgrading the SFCFSHA stack on the first subcluster

To upgrade the SFCFSHA stack on the first subcluster

- ◆ **Note:** This procedure is based on an "in-place" upgrade path; that is, if the operating system is upgraded, the release will be the same, but only the path level will change. If you are moving from major operating system release to another, you must uninstall the SFCFSHA stack before you upgrade the operating system. After the operating system is upgraded, you must reinstall SFCFSHA.
-

On the first half of the cluster, upgrade SFCFSHA by using the `installmr` script. For example use the `installmr` script as shown below:

```
# ./installmr sys1
```

where `<sys1>` is the node on the first subcluster.

After the upgrade for first half of the cluster is complete, no GAB ports will be shown in `gabconfig -a` output.

To upgrade your operating system, follow the normal procedures for your platform.

Note: After the installation completes, you can safely ignore any instructions that the installer displays.

Preparing the second subcluster

To prepare the second subcluster

- 1 On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application. [Downtime starts now.]
- 2 On the second half of the cluster, unmount the VxFS and CFS file systems that are not managed by VCS. Enter the following:

```
# mount | grep vxfs
```

Verify that no processes use the VxFS and CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS and CFS mount point with the mechanism provided by the application.

Unmount the VxFS and CFS file system. Enter the following:

```
# umount /mount_point
```

- 3 On the second half of the cluster, unfreeze all the VCS service groups on all the nodes using the following commands:

```
# haconf -makerw
```

```
# hagr -unfreeze group_name -persistent
```

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

- 4 On the second half of the cluster, bring all the VCS service groups offline, including CVM group. Enter the following:

```
# hagr -offline group_name -sys sys4
```

- 5 On the second half of the cluster, verify that the VCS service groups are offline. Enter the following:

```
# hagr -state group_name
```

- 6 Stop VCS on the second half of the cluster. Enter the following:

```
# hastop -local
```

- 7 On each node of the second half of the cluster, change the contents of the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cp /etc/vxfen.d/vxfenmode_disabled /etc/vxfenmode
# cat /etc/vxfenmode#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=disabled
```

- 8 If the cluster-wide attribute `UseFence` is set to `SCSI3`, reset the value to `NONE` in the `/etc/VRTSvcs/conf/config/main.cf` file, in second half of the cluster.
- 9 On the second half on cluster, stop the following SFCFSA modules: `VCS`, `VxFEN`, `ODM`, `GAB`, and `LLT`. Enter the following:

```
# /etc/init.d/vxglm stop
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- 10** On each node in the first half of the cluster, enable fencing. Enter the following:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# sybase     - use scsi3 disks in kernel but coordinate membership
# with Sybase ASE
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=scsi3
#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
scsi3_disk_policy=dmp
```

- 11** If the cluster-wide attribute UseFence is set to NONE, reset the value to SCSI3 in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

Activating the first subcluster

To activate the first subcluster

- 1 Restart the upgraded nodes in the first half of the cluster:

```
# /sbin/shutdown -r now
```

When the first half of the cluster nodes come up, no GAB ports are OPEN. The following command does not show any GAB ports:

```
# /sbin/gabconfig -a
```

```
GAB Port Memberships
```

```
=====
```

- 2 If required, force gab to form a cluster after the upgraded nodes are rebooted in first half of the cluster.

```
# /sbin/gabconfig -x
```

GAB ports a, b, d and h appear in `gabconfig -a` command output.

Note: If port b and h are not up, you need to bring fencing and VCS manually online.

- 3 On the first half of the cluster, bring the VCS service groups online. Enter the following:

```
# hagrps -online group_name -sys node_name
```

After you bring the CVM service group ONLINE, all the GAB ports u, v, w and f come ONLINE and all the CFS mounts service groups also come ONLINE automatically. Only failover service groups need to be brought ONLINE manually.

- 4 Manually mount the VxFS and CFS file systems that are not managed by VCS. [Downtime ends now.]

Upgrading the operating system on the second subcluster

To upgrade the operating system on the second subcluster

- ◆ Enter the following.

```
# chkconfig vcs off
# chkconfig vxfen off
# chkconfig gab off
# chkconfig llc off
```

On the second half of the cluster, upgrade the operating system, if applicable. For instructions, see the upgrade paths for the operating system.

See [“Supported Linux operating systems”](#) on page 17.

Upgrading the second subcluster

To upgrade the second subcluster

- ◆ Enter the following:

```
# ./installmr node_name
```

Completing the phased upgrade

To complete the phased upgrade

- 1 Verify that the cluster UUID on the nodes in the second subcluster is the same as the cluster UUID on the nodes in the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -display nodename
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first subcluster to the nodes in the second subcluster. For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -copy -from_sys \  
node01 -to_sys node03 node04
```

- 2 On each node in the second half of the cluster, enable fencing. Enter the following:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode  
# cat /etc/vxfenmode  
#  
# vxfen_mode determines in what mode VCS I/O Fencing should work.  
#  
# available options:  
# scsi3      - use scsi3 persistent reservation disks  
# customized - use script based customized fencing  
# sybase     - use scsi3 disks in kernel but coordinate  
#             membership with Sybase ASE  
# disabled   - run the driver but don't do any actual fencing  
#  
vxfen_mode=scsi3  
#  
# scsi3_disk_policy determines the way in which I/O Fencing  
# communicates with the coordination disks.  
#  
# available options:  
# dmp - use dynamic multipathing  
# raw - connect to disks using the native interface  
#  
scsi3_disk_policy=dmp
```

3 Enter the following before reboot of nodes:

```
# chkconfig vcs on
# chkconfig vxfs on
# chkconfig gab on
# chkconfig lltd on
```

4 Restart the upgraded nodes in the second half of the cluster:

```
# /sbin/shutdown -r now
```

When second half of the nodes come up, all the GAB ports a, b, d, h, u, v, w and f are ONLINE. Also all the CFS mounts service groups come online automatically.

5 Manually mount the VxFS and CFS file systems that are not managed by VCS in the second half of the cluster.

6 Find out which node is the CVM master. Enter the following:

```
# vxctl -c mode
```

Performing phased upgrade of SF Oracle RAC to version 6.0.3

Table 3-2 illustrates the phased upgrade process. Each column describes the steps to be performed on the corresponding subcluster and the status of the subcluster when operations are performed on the other subcluster.

Table 3-2 Summary of phased upgrade

First half of the cluster	Second half of the cluster
---------------------------	----------------------------

SF Oracle RAC cluster before the upgrade:




Table 3-2 Summary of phased upgrade (*continued*)

First half of the cluster	Second half of the cluster
<p>STEP 1: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> Switch failover applications. Stop all parallel applications. <p>See “Step 1: Performing pre-upgrade tasks on the first half of the cluster” on page 197.</p> <p>STEP 2: Upgrade SF Oracle RAC.</p> <p>See “Step 2: Upgrading the first half of the cluster” on page 199.</p>	<p>The second half of the cluster is up and running.</p> 
<p>The first half of the cluster is not running.</p> 	<p>STEP 3: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> Stop all parallel and failover applications. Stop SF Oracle RAC. <p>See “Step 3: Performing pre-upgrade tasks on the second half of the cluster” on page 200.</p> <p>The downtime starts now.</p>
<p>STEP 4: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> Start SF Oracle RAC. Start all applications. <p>See “Step 4: Performing post-upgrade tasks on the first half of the cluster” on page 201.</p> <p>The downtime ends here.</p>	<p>The second half of the cluster is not running.</p> 
<p>The first half of the cluster is up and running.</p> 	<p>STEP 5: Upgrade SF Oracle RAC.</p> <p>See “Step 5: Upgrading the second half of the cluster” on page 202.</p> <p>STEP 6: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> Start SF Oracle RAC. Start all applications. <p>See “Step 6: Performing post-upgrade tasks on the second half of the cluster” on page 203.</p>

Table 3-2 Summary of phased upgrade (*continued*)

First half of the cluster	Second half of the cluster
---------------------------	----------------------------

The phased upgrade is complete and both the first and the second half of the cluster are running.



Step 1: Performing pre-upgrade tasks on the first half of the cluster

Perform the following pre-upgrade steps on the first half of the cluster.

To perform the pre-upgrade tasks on the first half of the cluster

- 1 Back up the following configuration files: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```

# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.save
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
/etc/VRTSvcs/conf/config/OracleTypes.cf.save
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
/etc/VRTSvcs/conf/config/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
/etc/VRTSvcs/conf/config/MultiPrivNIC.cf.save
    
```

- 2 Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

If not, a warning message is displayed.

Warning: Backup `/etc/vx/cbr/bk` directory.

- 3 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

4 Stop the applications configured under VCS. Stop the Oracle RAC database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrpl -offline oracle_group -sys sys1
# hagrpl -offline oracle_group -sys sys2
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the first half of the cluster and shut down the instances:

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

- #### 5
- If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# hagrpl -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

```
$ srvctl modify database -d db_name -y manual
```

6 Unmount the CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs | grep cluster

# fuser -cu /mount_point
```

- Unmount the non-system CFS file system:

```
# umount /mount_point
```

- 7 Stop the parallel service groups and switch over failover service groups on each of the nodes in the first half of the cluster:

```
# hastop -local -evacuate
```
- 8 Unmount the VxFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
```

```
# fuser -cu /mount_point
```
 - Unmount the non-system VxFS file system:

```
# umount /mount_point
```
- 9 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.
- 10 If you plan to upgrade the operating system, stop all ports.

```
# /opt/VRTS/install/installsfrac -stop sys1 sys2
```

Step 2: Upgrading the first half of the cluster

Perform the following steps to upgrade the first half of the cluster.

To upgrade the first half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```
- 2 Upgrade the operating system, if required.
For instructions, see the operating system documentation.
- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```
- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 Upgrade SF Oracle RAC. Navigate to the product directory on the installation media. When you invoke the installer, select the Full Upgrade option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

```
# cd product folder

# cd /dvd_mount/specific_os/

# ./installmr sys1 sys2
```

Note: After you complete the upgrade of the first half of the cluster, no GAB ports will be shown in the output when you run the `gabconfig -a` command.

Step 3: Performing pre-upgrade tasks on the second half of the cluster

Perform the following pre-upgrade steps on the second half of the cluster.

To perform the pre-upgrade tasks on the second half of the cluster

- 1 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

Note: The downtime starts now.

- 2 Stop all applications that are configured under VCS. Stop the Oracle RAC database:
 - If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys sys3
# hagrps -offline oracle_group -sys sys4
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the second half of the cluster and shut down the instances:

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \  
-i instance_name
```

- 3 Unmount the CFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs | grep cluster  
  
# fuser -cu /mount_point
```
 - Unmount the non-system VxFS file system:

```
# umount /mount_point
```
- 4 Stop VCS on each of the nodes in the second half of the cluster:

```
# hastop -local
```
- 5 Unmount the VxFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs  
  
# fuser -cu /mount_point
```
 - Unmount the non-system VxFS file system:

```
# umount /mount_point
```
- 6 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.
- 7 Stop all ports.

```
# /opt/VRTS/install/installsfrac -stop sys3 sys4
```

Step 4: Performing post-upgrade tasks on the first half of the cluster

Perform the following post-upgrade steps on the first half of the cluster.

To perform the post-upgrade tasks on the first half of the cluster

- 1 On any one node on the first half of the cluster, force GAB to form a cluster.

```
# /etc/init.d/llt start
# /etc/init.d/gab start

# gabconfig -x
```

- 2 On the first half of the cluster, start SF Oracle RAC:

```
# cd /opt/VRTS/install
# ./installsfrac -start sys1 sys2
```

- 3 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.

- 4 Relink the SF Oracle RAC libraries with Oracle.

- 5 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys node_name
```

If the Oracle database is not managed by VCS:

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \
-i instance_name
```

Note: The downtime ends here.

- 6 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

Step 5: Upgrading the second half of the cluster

Perform the following steps to upgrade the second half of the cluster.

To upgrade the second half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.
For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.

- 6 On the second half of the cluster, upgrade SF Oracle RAC. Navigate to the product directory on the installation media.

When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

```
# cd /dvd_mount/specific_os
```

```
# ./installmr sys3 sys4
```

Step 6: Performing post-upgrade tasks on the second half of the cluster

Perform the following post-upgrade steps on the second half of the cluster.

To perform the post-upgrade tasks on the second half of the cluster

- 1 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 2 On the second half of the cluster, start SF Oracle RAC:

```
# cd /opt/VRTS/install
```

```
# ./installsfrac -start sys3 sys4
```

- 3 Relink the SF Oracle RAC libraries with Oracle.

4 Upgrade VxVM disk group version.

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk group version to 180.

Check the existing disk group version:

```
# vxvg list dg_name | grep -i version
```

If the disk group version is not 180, run the following command on the master node to upgrade the version:

```
# vxvg -T 180 upgrade dg_name
```

5 Upgrade disk layout version.

In this release, you can create and mount only file systems with disk layout Version 7, 8, and 9. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Disk layout Version 6 has been deprecated and you cannot cluster mount an existing file system that has disk layout Version 6. To upgrade a cluster file system with disk layout Version 6, you must local mount the file system and then upgrade the file system using the `vxupgrade` utility to a later version.

See the `vxupgrade(1M)` manual page.

Support for disk layout Version 4 and 5 has been removed. You must upgrade any existing file systems with disk layout Version 4 or 5 to disk layout Version 7 or later using the `vxfscconvert` command.

See the `vxfscconvert(1M)` manual page.

Note: Symantec recommends that you upgrade existing file systems to the highest supported disk layout version prior to upgrading to this release.

You can check which disk layout version your file system has by using the following command:

```
# fstyp -v /dev/vx/dsk/dg1/vol1 | grep -i version
```

For more information about disk layout versions, see the *Veritas Storage Foundation Administrator's Guide*.

6 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrpl -online oracle_group -sys sys3
# hagrpl -online oracle_group -sys sys4
```

If the Oracle database is not managed by VCS:

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \
-i instance_name
```

7 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrpl -modify oracle_group AutoStart 1
# haconf -dump -makero
```

■ If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

8 Start all applications that are not managed by VCS. Use native application commands to start the applications.

9 Set or change the product license level, if required.

10 Migrate the SFDB repository database.

11 If you want to upgrade CP server systems that use VCS or SFHA to 6.0.3, make sure that you upgraded all application clusters to version 6.0.3. Then, upgrade VCS or SFHA on the CP server systems. For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Upgrading the operating system

This section describes how to upgrade the operating system on a Veritas Storage Foundation and High Availability Solutions (SFHA) node where you plan to upgrade to 6.0.3 for Red Hat Enterprise Linux (RHEL) 5, RHEL 6, SUSE Linux Enterprise (SLES) 11, Oracle Linux (OL) 5, and OL 6.

Note: If you are upgrading to RHEL 6 Update 4, you must install an additional SFHA package.

See [“Support for RHEL 6.4, RHEL 6.5, OL 6.4, and OL 6.5”](#) on page 19.

To upgrade the operating system to a later version

- 1 Stop Storage Foundation.
- 2 Upgrade to the latest operating system.
- 3 Upgrade to 6.0.3.
See [“Upgrading to 6.0.3”](#) on page 166.
- 4 Start Storage Foundation.

Verifying software versions

To list the Veritas RPMs installed on your system, enter the following command:

```
# rpm -qa | egrep VRTS
```

Uninstalling version 6.0.3

This chapter includes the following topics:

- [About removing Veritas Storage Foundation and High Availability Solutions 6.0.3](#)
- [Uninstalling Veritas Storage Foundation Cluster File System 6.0.3](#)
- [Uninstalling Veritas Storage Foundation for Oracle RAC 6.0.3](#)
- [Uninstalling Veritas Storage Foundation for Sybase CE 6.0.3](#)

About removing Veritas Storage Foundation and High Availability Solutions 6.0.3

Symantec recommends that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the previous release.

For extensive uninstallation and reinstallation procedures, refer to the appropriate product's Installation Guide.

Uninstalling Veritas Storage Foundation Cluster File System 6.0.3

This section provides procedures for uninstalling Veritas Storage Foundation Cluster File System High Availability (SFCFSHA). You must complete the preparatory tasks before you uninstall SFCFSHA.

Preparing to uninstall Veritas Storage Foundation Cluster File System High Availability 6.0.3

The following procedure prepares your system for uninstalling Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

To prepare to uninstall Veritas Storage Foundation Cluster File System

- 1 Log in as the root user on any node in the cluster.
- 2 Verify that the following directories are set in your `PATH` environment variable:

```
/opt/VRTS/bin
/opt/VRTSvcs/bin
```

- 3 Determine if each node's root disk is under VxVM control and proceed as follows.
 - Check if each node's root disk is under VxVM control:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

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

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

- 4 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
 - Use the `vrxvg stop` command to stop each RVG individually:


```
# vxrvg -g diskgroup stop rvg_name
```
 - On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:


```
# vxrlink -g diskgroup status rlink_name
```
- 5 Check if any VxFS file systems or Storage Checkpoints are mounted:
- ```
df -T | grep vxfs
```
- 6 Unmount all Storage Checkpoints and file systems:
- ```
# umount /checkpoint1
# umount /filesystem1
```
- If file system is mounted in a cluster, then use `cfsumount` command.
- 7 Stop all VxVM volumes by entering the following command for each disk group:
- ```
vxvol -g dg1 stopall
```
- To verify that no volumes are open:
- ```
# vxprint -Aht -e v_open
```
- 8 Stop VCS:
- ```
hastop -all
```

## Uninstalling Veritas Storage Foundation Cluster File System High Availability

The following procedure uninstalls Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

**To uninstall Veritas Storage Foundation Cluster File System High Availability**

- 1 Log in as the root user on any node in the cluster.
- 2 Navigate to the directory that contains the uninstallation program:

```
cd /opt/VRTS/install
```

- 3 Run the uninstallation program while specifying each node in the cluster:

```
./uninstallsfcfsha601 node1 node2
```

- 4 Confirm the uninstallation:

```
Are you sure you want to uninstall SFCFSHA [y,n,q] (y)
```

The installer stops the Veritas Storage Foundation Cluster File System processes and uninstalls the packages.

After uninstalling the Veritas Storage Foundation Cluster File System, refer to the *Veritas Storage Foundation Cluster File System 6.0.1 Installation Guide* to reinstall the 6.0.1 software.

## Uninstalling Veritas Storage Foundation for Oracle RAC 6.0.3

The following procedure uninstalls Veritas Storage Foundation for Oracle RAC (SFRAC).

---

**Note:** This procedure will remove the complete SFRAC stack from all nodes.

---

**To uninstall Veritas Storage Foundation for Oracle RAC**

- 1 On each node, take the Oracle resources in the VCS configuration file (`main.cf`) offline.

```
hagrpl -offline oracle_group -sys node_name
```

If the database is not managed by VCS, stop the Oracle database as follows:

```
$ srvctl stop database -d db_name
```

- 2 If Oracle Clusterware is not under VCS Control, then enter the following command on each node of the cluster to stop Oracle Clusterware.
  - For 10gR2 or 11gR1:

```
/etc/init.d/init.crs stop
```

- For 11gR2:

```
$ GRID_HOME/bin/crsctl stop crs
```

### 3 Stop the applications that use CVM or CFS that are not under VCS control

- Using native application commands, stop the applications that use CVM or CFS on all nodes.
- Verify that no processes use the CFS mount point:

```
fuser -c mount_point
```

### 4 Unmount CFS file systems that are not under VCS control

- Determine the file systems that need to be unmounted by checking the output of mount command.

```
mount -v | grep vxfs | grep cluster
```

- Unmount each file system that is not controlled by VCS on each node:

```
umount mount_point
```

### 5 Stop VCS to take the service groups on all nodes offline

On any one node execute following command to stop VCS:

```
hastop -all
```

### 6 Stopping the applications that use VxVM or VxFS that are not under VCS control

- Using native application commands, stop the applications that use VxVM or VxFS.
- Verify that no processes use the VxFS mount point:

```
fuser -c mount_point
```

### 7 Unmounting VxFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

```
mount -v | grep vxfs
```

Unmount each file system that is not controlled by VCS on each node:

```
umount mount_point
```

## 8 Remove SF for Oracle RAC.

- On any one node, navigate to the directory that contains the `uninstallsfrac` program:

```
cd /opt/VRTS/install
```

- Start the `uninstallsfrac` program:

```
./uninstallsfrac601
```

- 9 After uninstalling the SFRAC, refer to the *Veritas Storage Foundation for Oracle RAC 6.0.1 Installation and Configuration Guide* document to reinstall the SFRAC 6.0.1 software.

# Uninstalling Veritas Storage Foundation for Sybase CE 6.0.3

You can use the following procedure to uninstall SF Sybase CE 6.0.3 and re-install 6.0.1.

## To uninstall Veritas Storage Foundation for Sybase CE 6.0.3

- 1 Log in as the superuser on one of the nodes in the cluster.
- 2 On each node, take the Sybase resources in the VCS configuration file (`main.cf`) offline.

```
hagrps -offline Sybase_group -sys node_name
```

For example:

```
/opt/VRTSvcs/bin/hagrps -offline sybasece -sys sys1
```

```
/opt/VRTSvcs/bin/hagrps -offline sybasece -sys sys2
```

These commands stop the Sybase resources under VCS control.

- 3 Verify that the state of the Sybase and CVM service groups are offline and online respectively.

```
/opt/VRTSvcs/bin/hagrp -state
Group Attribute System Value
binmnt State sys1 |ONLINE|
binmnt State sys2 |ONLINE|
cvm State sys1 |ONLINE|
cvm State sys2 |ONLINE|
sybasece State sys1 |OFFLINE|
sybasece State sys2 |OFFLINE|
```

- 4 Backing up the Sybase database.

If you plan to retain the Sybase database, you must back up the Sybase database. For instructions on backing up the Sybase database, see the Sybase documentation.

- 5 Stop the applications that use CVM volumes or CFS mount points not controlled by VCS.

To stop the applications that use CVM or CFS (outside of VCS control):

- Stop the applications that use a CFS mount point. The procedure varies for different applications. Use the procedure appropriate for the application.
- Verify that no processes use the CFS mount point:

```
fuser -c mount_point
```

- 6 Unmount CFS file systems that are not under VCS control on all nodes.

To unmount CFS file systems not under VCS control:

- Determine the file systems that need to be unmounted by checking the output of the mount command. The command lists all the mounted clustered file systems. Consult the main.cf file for identifying the files that are under VCS control.

```
mount -v | grep vxfs | grep cluster
```

- Unmount each file system that is not controlled by VCS:

```
umount mount_point
```

- 7 Stop VCS to take the service groups on all nodes offline.

```
hastop -all
```

- 8 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped. In this command output, the VCS engine or high availability daemon (HAD) port `h` is not displayed. This output indicates that VCS has been stopped.

```
/sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen 5c3d0b membership 01
Port b gen 5c3d10 membership 01
```

- 9 Stop all applications that use VxVM volumes or VxFS mount points not under VCS control.

To stop the applications that use VxVM or VxFS (outside of VCS control):

- Stop the applications that use a VxFS mount point. The procedure varies for different applications. Use the procedure that is appropriate for your application.
- Verify that no processes use the VxFS mount point:

```
fuser -c mount_point
```

- 10 Unmount VxFS file systems that are not under VCS control on all nodes.

---

**Note:** To avoid issues on rebooting, you must remove all entries of VxFS from the `/etc/vfstab` directory.

---

To unmount VxFS file systems not under VCS control:

- Determine the file systems that need to be unmounted by checking the output of the `mount` command. The command lists all the mounted file systems.

```
mount -v | grep vxfs
```

- Unmount each file system that is not under VCS control:

```
umount mount_point
```

- 11 Remove the SF Sybase CE:

- On any one node, navigate to the directory that contains the `uninstallsybasece` program:

```
cd /opt/VRTS/install
```

- Start the `uninstallsfsybasece` program:

```
./uninstallsfsybasece601
```

- 12 After uninstalling the SF Sybase CE, refer to the *Veritas Storage Foundation for Sybase ASE CE 6.0.1 Installation and Configuration Guide* document to reinstall the SF for Sybase CE 6.0.1 software.