

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

6.0.3 Maintenance Release

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 7

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	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 and the uninstallmr scripts	10
The installmr script options	10
The uninstallmr script options	14
Overview of the installation and upgrade process	17
Changes introduced in 6.0.3	17
Changes introduced in Veritas Cluster Server 6.0.3	18
Changes introduced in SFDB 6.0.3	18
System requirements	18
Supported AIX operating systems	18
Database requirements	18
Veritas Storage Foundation memory requirements	19
Disk space requirements	19
Number of nodes supported	20
List of products	20
Fixed issues	20
Installation and upgrades fixed issues	20
Veritas Dynamic Multi-pathing fixed issues	21
Veritas Storage Foundation fixed issues	21
Veritas Cluster Server fixed issues	32
Veritas Storage Foundation and High Availability fixed issues	36
Veritas Storage Foundation for Oracle RAC fixed issues	36
Veritas Storage Foundation Cluster File System High Availability fixed issues	38
Known issues	39
Issues related to installation and upgrade	39
Veritas Dynamic Multi-pathing known issues	45
Veritas Storage Foundation known issues	48
Veritas Cluster Server known issues	76

	Veritas Storage Foundation and High Availability known issues	99
	Veritas Storage Foundation Cluster File System High Availability known issues	99
	Veritas Storage Foundation for Oracle RAC known issues	103
	Software limitations	106
	Limitations related to rolling upgrade	107
	Veritas Dynamic Multi-pathing software limitations	107
	Veritas Storage Foundation software limitations	108
	Veritas Cluster Server software limitations	111
	Veritas Storage Foundation and High Availability software limitations	121
	Veritas Storage Foundation Cluster File System High Availability software limitations	123
	Veritas Storage Foundation for Oracle RAC software limitations	123
	Documentation errata	124
	List of patches	125
	Downloading the 6.0.3 archive	125
Chapter 2	Installing the products for the first time	126
	Installing the Veritas software using the script-based installer	126
	Installing Veritas software using the Web-based installer	127
	Starting the Veritas Web-based installer	128
	Obtaining a security exception on Mozilla Firefox	128
	Installing 6.0.3 with the Veritas Web-based installer	128
Chapter 3	Upgrading to 6.0.3	130
	Prerequisites for upgrading to 6.0.3	130
	Downloading required software to upgrade to 6.0.3	130
	Supported upgrade paths	131
	Upgrading to 6.0.3	131
	Performing a full upgrade to 6.0.3 on a cluster	131
	Upgrading to 6.0.3 on a standalone system	141
	Performing a rolling upgrade using the <code>installmr</code> script	143
	Verifying software versions	150
Chapter 4	Uninstalling version 6.0.3	151
	About removing Veritas Storage Foundation and High Availability Solutions 6.0.3	151
	Rolling back using the <code>uninstallmr</code> script	151

- Rolling back manually 152
 - Rolling back Storage Foundation or Storage Foundation and High Availability manually 152
 - Rolling back Storage Foundation Cluster File System High Availability manually 155
 - Rolling back Storage Foundation for Oracle RAC manually 157
 - Rolling back Veritas Cluster Server manually 160
 - Rolling back Dynamic Multi-Pathing manually 163
 - 165

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [Introduction](#)
- [About the installmr and the uninstallmr scripts](#)
- [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 patches](#)
- [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

This Maintenance Release is available as 6.0.3.

About the installmr and the uninstallmr scripts

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

See “[Supported upgrade paths](#)” on page 131.

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
[-precheck]	Use the <code>-precheck</code> option to confirm that systems meet the products' installation requirements before the installation.
[-postcheck]	Use the <code>-postcheck</code> option after an installation or upgrade to help you determine installation-related problems and provide troubleshooting information.
[-responsefile <i>response_file</i>]	Use the <code>-responsefile</code> option to perform automated installations or uninstallations 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.
[-logpath <i>log_path</i>]	Use the <code>-logpath</code> option to select a directory other than <code>/opt/VRTS/install/logs</code> as the location where the installmr log files, summary file, and response file are saved.
[-tmppath <i>tmp_path</i>]	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.
[-timeout <i>timeout_value</i>]	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.
[-keyfile <i>ssh_key_file</i>]	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
[-hostfile <i>hostfile_path</i>]	Use the <code>-hostfile</code> option to specify the location of a file containing the system names for installer.
[-patchpath <i>patch_path</i>]	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> .
[-nim <i>LPP_SOURCE</i>]	Use to produce a NIM configuration file for installing with NIM. Refer to the product's <i>Installation Guide</i> for more information on using NIM.
[-serial -rsh -redirect -pkgset -pkgtable -pkginfo -listpatches]	Use the <code>-serial</code> option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion. 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. Use the <code>-redirect</code> option to display progress details without showing the progress bar. Use the <code>-pkgset</code> option to discover the package set installed on the systems specified. Use the <code>-pkgtable</code> option to display product rpms in correct installation order. 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>. Use the <code>-listpatches</code> option to display product patches in the correct installation order.

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

Command Line Option	Function
[-makeresponsefile -comcleanup -version -nolic]	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. 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. Use the <code>-version</code> option to check the status of installed products on the system. Use the <code>-nolic</code> option to install product filesets on systems without entering product licenses. Configuration, startup, or installation of license based features are not performed when using this option.

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

Command Line Option	Function
[-upgrade_kernelpkgs -upgrade_nonkernelpkgs -rolling_upgrade -rollingupgrade_phase1 -rollingupgrade_phase2]	The <code>-upgrade_kernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase1</code>. The <code>-upgrade_nonkernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase2</code>. 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. Use the <code>-rollingupgrade_phase1</code> option to to perform rolling upgrade phase 1. During this phase, the product kernel filesets will be upgraded to the latest version. Use the <code>-rollingupgrade_phase2</code> option perform rolling upgrade phase 2. During this phase, VCS and other agent filesets will be upgraded to the latest version. During this phase, product kernel drivers will be rolling-upgraded to the latest protocol version.

The uninstallmr script options

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

Symantec recommends that you use the uninstallation script. The `uninstallmr` script uninstalls all the patches associated with packages installed, and starts the processes. Do not use the `uninstallprodvers` script for rolling back, because it removes the entire stack.

Table 1-2 The command line options for the `uninstallmr` 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 1-2 The command line options for the `uninstallmr` script (*continued*)

Command Line Option	Function
[<code>-responsefile response_file</code>]	Use the <code>-responsefile</code> to perform automated installations or uninstallations 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 <code>uninstallmr</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>uninstallmr</code> . This destination is where initial logging is performed and where packages are copied on remote systems before installation.
[<code>-timeout <timeout_value></code>]	Use the <code>-timeout</code> 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.
[<code>-hostfile hostfile_path</code>]	Use the <code>-hostfile</code> option to specify the location of a file containing the system names for installer.

Table 1-2 The command line options for the `uninstallmr` script (*continued*)

Command Line Option	Function
<pre>[-serial -rsh -redirect -listpatches -makeresponsefile -comcleanup -version]</pre>	Use the <code>-serial</code> option to perform install, uninstall, start, and stop operations, typically performed simultaneously on all systems, in serial fashion. Use the <code>-rsh</code> option to have the script use <code>rsh</code> and <code>rcp</code> for communication between the systems. System communication using <code>rsh</code> and <code>rcp</code> is auto-detected by the script, so the <code>-rsh</code> option is only required when <code>ssh</code> and <code>scp</code> (default communication method) is also configured between the systems. Use the <code>-redirect</code> option to have the script display progress details without the use of advanced display functionality so that the output can be redirected to a file. Use the <code>-listpatches</code> option to display product patches in correct installation order. 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. Use the <code>-comcleanup</code> option to remove the <code>ssh</code> or <code>rsh</code> configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of <code>ssh</code> or <code>rsh</code> are abruptly terminated. Use the <code>-version</code> option to have the installer check and report the installed products and their versions. Identifies the installed and missing fileset and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing fileset and patches where applicable.

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.

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 AIX operating systems

This section lists the supported operating systems for this release of SFHA Solutions.

Product installation scripts verify the required update levels. The installation process terminates if the target systems do not meet the maintenance level requirements.

The minimum system requirements for this release are as follows:

For Power 7 or earlier processors at one of the following levels:

- AIX 7.1 TL0, AIX 7.1 TL1, AIX 7.1 TL2 with Service Pack 2, and AIX 7.1 TL3
- AIX 6.1 TL5 with Service Pack 1, AIX 6.1 TL6, AIX 6.1 TL7, AIX 6.1 TL8 with Service Pack 2, and AIX 6.1 TL9

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

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

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.

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

Additional Oracle support for SF Oracle RAC

Table 1-3 Oracle RAC versions that SF Oracle RAC supports

Oracle version	AIX 6.1	AIX 7.1
10gR2 10.2.0.5 (64-bit)	Yes	No
11gR2 11.2.0.2 (64-bit)	Yes	Yes
11gR2 11.2.0.3 (64-bit)	Yes	Yes

Note: 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 and the uninstallmr scripts”](#) 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)

Fixed issues

This section describes the issues fixed in 6.0.3.

See the `README_SYMC.xxxxxx-xx` files in the `/patches` 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 1-4 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 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 1-5 Fixed issues related to installation and upgrades

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

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 1-6](#) describes the incidents that are fixed in Veritas Volume Manager in 6.0.3.

Table 1-6 Veritas Volume Manager 6.0.3 fixed issues

Incident	Description
3002770	Accessing NULL pointer in dmp_aa_recv_inquiry() caused system panic.
2970368	Enhancing handling of SRDF-R2 WD devices in DMP.
2965910	vxassist dump core with the <code>-o ordered</code> option.
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.
2942609	Message displayed when user quits from Dynamic Reconfiguration Operations is shown as error message.
2940446	Full fsck hangs on I/O in VxVM when cache object size is very large
2935771	In the VVR environment, RLINK disconnects after the master is switched.
2934729	VM is claiming disks as 'online' in VIOS.
2933138	panic in voldco_update_itemq_chunk() due to accessing invalid buffer
2930569	The LUNs in 'error' state in output of 'vxdisk list' cannot be removed through DR(Dynamic Reconfiguration) Tool.
2919720	vxconfigd core in rec_lock1_5()
2919714	exit code from vxevac is zero when migrating on thin luns but FS is not mounted
2919627	Dynamic Reconfiguration tool should be enhanced to remove LUNs feasibly in bulk.

Table 1-6 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

Incident	Description
2919318	The I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing.
2916094	Enhancements have been made to the Dynamic Reconfiguration Tool(DR Tool) to create a separate log file every time DR Tool is started, display a message if a command takes longer time, and not to list the devices controlled by TPD (Third Party Driver) in 'Remove Luns' option of DR Tool.
2915063	Rebooting VIS array having mirror volumes, master node panicked and other nodes CVM FAULTED
2911040	Restore from a cascaded snapshot when its source is DETACHED leaves the volume in unusable state
2910043	Avoid order 8 allocation by vxconfigd in node reconfig.
2899173	vxconfigd hang after executing the <code>vradmind stoprep</code> comand.
2898547	vradmind on VVR Secondary Site dumps core, when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2892983	vxvol dumps core if new links are added while the operation is in progress.
2886402	vxconfigd hang while executing <code>tc ./scripts/ddl/dmpapm.tc#11</code>
2886333	The <code>vxdbg(1M) join</code> command should not allow mixing clone and non-clone disks in a DiskGroup
2884225	vxconvert command fails to convert 1.5TB AIX LVM diskgroup to vxvm diskgroup
2882908	Machine failed to bootup with error "PReP-BOOT : Unable to load full PReP image"
2879248	vxdisk scandisks gets hung on VIO client with <code>dmp_native_support</code> enabled
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 1-6 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

Incident	Description
2858853	vxconfigd core dumps in dbf_fmt_tbl on the slave node after a Master Switch if you try to remove a disk from the DG
2851403	The vxportal and vxfs processes are failed to stop during first phase of rolling upgrade.
2851085	DMP doesn't detect implicit LUN ownership changes for some of the dmpnodes.
2837717	The vxdisk(1M) resize command fails if da name is specified.
2836798	Prevent DLE on simple/sliced disk with EFI label
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	vx dg adddisk should not allow mixing clone and 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 volumes 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	vxevac move error message needs to be enhanced to be less generic and give clear message for failure.
2619600	Live migration of virtual machine having SFHA/SFCFSHA stack with data disks fencing enabled, causes service groups configured on virtual machine to fault.

Table 1-6 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

Incident	Description
2567618	VRTSexplorer coredumps in vxcheckhbaapi/print_target_map_entry
2510928	Extended attributes for SRDF luns reported as Mirror with EMC (VMAX array)
2482308	Devices go into error state after unmanaging them from Powerpath
2398416	vxassist dumps core while creating volume after adding attribute "wantmirror=ctlr" 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	vx dg 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 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>vx dmpadm listctlr</code> command has to be enhanced to print the count of device paths seen through the controller.
1289985	vxconfigd core dumps upon running the <code>vx dctl enable</code> command.

Veritas Volume Manager: issues fixed in 6.0.1

Table 1-7 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.

Table 1-7 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2826958	The pwnn 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.
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.
2743926	DMP <code>restored</code> daemon fails to restart during system boot.
2741240	The <code>vxdg join</code> transaction failed and did not rollback to the sourcedg.
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.
2727590	The <code>vxconfigd</code> daemon dumped core after renaming iSCSI device from the OS side.
2713166	Suppress <code>vxio errpt</code> messages relating to write-disabled (WD) & NR devices (EMC SRDF-R2) & EMC BCV.
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.
2709767	Thin provisioning reclaim should not work with MPIO.
2700792	SEGV in <code>vxconfigd</code> daemon during CVM startup.

Table 1-7 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
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.
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.
2656803	Race between <code>vxnetd start</code> and <code>stop</code> operations causes panic.
2652485	Inactive snapshot LUNs cause trespassing.
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.
2612960	The <code>vxconfigd</code> daemon core dumps after upgrading, due to GPT/AIX label disk.
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_glock()</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.

Table 1-7 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2533248	SAN Boot: A node can reboot and bring volumes online with a failed path, but failed to boot up after restoration of the failed path again in next reboot.
2526606	Memory leaks in DMP.
2441283	The <code>vxsnap addmir</code> command sometimes fails under heavy I/O load.
2427894	Opaque disk support for VIS appliance.
2249445	Develop a tool to get the disk-related attributes like geometry, label, media capacity, partition info etc.
2240056	The <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'.
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>vxxtune</code> .

Veritas File System fixed issues

This section describes Veritas File System fixed issues.

Veritas File System: issues fixed in 6.0.3

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

Table 1-8 Veritas File System 6.0.3 fixed issues

Incident	Description
2978326	Changing value of the <code>dalloc_enable/dalloc_limit</code> tunables fails on a cluster mounted filesystem.
2895743	Accessing named attributes for some files seems to be slow.
2887423	Severe lock contention in <code>vx_sched</code> . HF1e needs to be addressed.
2885592	<code>vxdump</code> to the <code>vxcompress</code> file system is aborted.
2881211	File ACLs not preserved in checkpoints properly if file has hardlink.
2878164	VxFS consuming too much pinned heap.

Table 1-8 Veritas File System 6.0.3 fixed issues (*continued*)

Incident	Description
2874054	The vxconvert command fails to convert LVM diskgroup to vxvm diskgroup due to malloc issue
2858683	Reserve extent attributes changed after vxrestore, only for files greater than 8192bytes.
2857751	The internal testing hits the assert "f:vx_cbdnrc_enter:1a".
2857629	File system corruption can occur requiring a full fsck of the system.
2857568	Performance issues seen during back-up operations reading larges files sequentially.
2848948	VxFS buff cache consumption increased significantly after running over 248 days.
2806466	fsadm -R resulting in panic at LVM layer due to vx_ts.ts_length set to 2GB.
2624262	fsdedup.bin hit oops at vx_bc_do_brelse.
2590918	Delay in freeing unshared extents upon primary switch over.

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 1-9 Veritas File System fixed issues

Incident	Description
2764861	Uncompress by vxcompress ignores quota limitation.
2753944	The file creation threads can hang.
2735912	The performance of tier relocation using fsppadm enforce is poor when moving a large amount of files.
2712392	Threads hung in VxFS.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2682055	On AIX, If drefund_supported is set to 1 (i.e. drefund_supported=1) and drefund is not enabled(i.e. drefund_enable=0) then tuning of 'vmmbufs_resv_disable' to 1 fails.

Table 1-9 Veritas File System fixed issues (*continued*)

Incident	Description
2674639	The cp(1) command with the -p option may fail on a file system whose File Change Log (FCL) feature is enabled. The following error messages are displayed: cp: setting permissions for 'file_name': Input/output error cp: preserving permissions for 'file_name': No data available.
2670022	Duplicate file names can be seen in a directory.
2655788	Using cross-platform data sharing to convert a file system that has more than 32k nlinks does not update the vx_maxlink and maxlink_enable tunables.
2651922	Is -l command on local VxFS file system is running slow and high CPU usage is seen.
2650354	Allow 8MB and 4MB values for chunk_flush_size tunable on AIX.
2650330	Accessing a file with O_NSHARE mode by multiple process concurrently on AIX could cause file system hang.
2626390	Freeing a large number of pages at once can induce a small I/O latency.
2597347	fsck should not coredump when only one of the device record has been corrupted and the replica is intact.
2566875	The write(2) operation exceeding the quota limit fails with an EDQUOT error (Disc quota exceeded) before the user quota limit is reached.
2559450	Command fsck_vxfs(1m) may core-dump with SEGV_ACCERR error.
2536130	fscdsconv fails to convert FS between specific platforms if FCL is enabled.
2272072	GAB panics the box because VCS engine HAD did not respond. The lobolt wraps around.
2086902	Spinlock held too long on vxfs spinlock, and there is high contention for it.
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 1-10](#) describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 6.0.3.

Table 1-10 Veritas Storage Foundation for Databases (SFDB) tools 6.0.3 fixed issues

Incident	Description
3030663	<code>dbed_vmclonedb</code> does not read pfile supplied by <code>-p 'pfile_modification_file'</code> option.

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

[Table 1-11](#) describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-11 SFDB tools fixed issues

Incident	Description
2585643	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: <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> The error messages are unclear.
2703881 (2534422)	The FlashSnap validation operation fails with the following error if the mirrors for data volumes and archive log volumes share the same set of disks: <pre>SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.</pre>
2582694 (2580318)	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.

Table 1-11 SFDB tools fixed issues (*continued*)

Incident	Description
2579929	The <code>sfae_auth_op -o auth_user</code> command, used for authorizing users, fails with the following error message: <pre>SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username></pre> 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.

Veritas Cluster Server fixed issues

This section describes the Veritas Cluster Server fixed issues.

Veritas Cluster Server: issues fixed in 6.0.3

[Table 1-12](#) describes the incidents that are fixed in Veritas Cluster Server in 6.0.3.

Table 1-12 Veritas Cluster Server 6.0.3 fixed issues

Incident	Description
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 WPARs.
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
2737653	Incorrect descriptions about the RVGPrimary online script .
2736627	REMOTE CLUSTER STATE remains in INIT state and lcmp heartbeat status is UNKNOWN.
2619600	Use of Live Partition Mobility on an SFHA or SFCFSHA node with SCSI-3 fencing enabled for data disks causes service groups on that node to fault.

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 1-13](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-13 LLT, GAB, and I/O fencing fixed issues

Incident	Description
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 1-14](#) lists the fixed issues for bundled agents.

Table 1-14 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.
2773376	On a VCS node running the AIX operating system, the <code>Netlsnr</code> agent may fail to start a database if your system uses LDAP for authentication and if the default shell for <code>Netlsnr</code> users is CSH.
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.

Table 1-14 Bundled agents fixed issues (*continued*)

Incident	Description
2850900	VSS_HOME gets set to wrong directory inWPARentry point, because of which the WPAR agent displays errors.
2850902	The lpar_sysoffline scripts are not updated as part of the package update.
2779780	When the monitor of a MultiNICB resource and online of an IPMultiNICB resource are scheduled at the same time, one of the entry point fails. This leads to fault of one of the resources and may also trigger an incorrect failover.
2850858	Error observed when ContainerInfo attribute for service group gets updated while running the hawparsetup.pl script.
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.
2850903	Following error message appears in the engine log when Application agent is configured inside WPAR and the agent tries to online the application: VCS ERROR V-16-1-10600 Cannot connect to VCS engine

VCS engine fixed issues in 6.0.1

[Table 1-15](#) lists the fixed issues for VCS engine.

Table 1-15 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.

Table 1-15 VCS engine fixed issues (*continued*)

Incident	Description
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: ■ PreOnline ■ ContainerInfo ■ TriggersEnabled ■ SystemZones
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 MigrateQ and TargetCount.
2746816	The syslog call used in <code>gab_heartbeat_alarm_handler</code> and <code>gabsim_heartbeat_alarm_handler</code> functions is not async signal safe.

Enterprise agents fixed issues in 6.0.1

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

Table 1-16 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.
2773376	Oracle agent does not function when the user authentication is performed through LDAP and the default shell is CSH.
2831044	Sybase agent script entry points must handle large process command line.
2699800	Db2udb resource is reported OFFLINE unexpectedly by the monitor entry point of the Db2udb agent.

Agent framework fixed issues in 6.0.1

[Table 1-17](#) lists the fixed issues for agent framework.

Table 1-17 Agent framework fixed issues

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

Veritas Storage Foundation 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 1-18](#) lists the issues fixed in 6.0.1.

Table 1-18 Issues fixed in 6.0.1

Incident	Description
2585899	The SF Oracle RAC installer does not support the use of fully qualified domain names (FQDN). Specifying the fully qualified domain name of a system results in the following error: <pre>The node sys1 doesn't seem to be part of the cluster, or CVM is not running on the node sys1.</pre>

Table 1-18 Issues fixed in 6.0.1 (*continued*)

Incident	Description
2329580	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.
2392741	Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool. If the cardinality of a policy-managed Oracle RAC database is set to a number lesser than the number of nodes in the server pool, and if the Oracle agent tries to bring the database online on all the nodes in the server pool, the operation fails on some of the nodes in the server pool. The resource on respective nodes move to the faulted state.
2749412	Setting the <code>UseVirtualIP</code> attribute to 1 overwrites the IP address of the virtual interface on some nodes in the cluster.
2757032	The PrivNIC/MultiPrivNIC agents fail to match the exact IP address configured in the agent configuration with the IP address configured on the system. As a result, the agent detects the wrong interface as the active interface resulting in a resource fault.
2580393	Removal of SAN cable from any node in a global cluster setup takes application service groups offline on all nodes. In a replicated global cluster setup, the removal of SAN cable from any node in the cluster causes the CFS mount points to fault. As a result, dependent application groups are taken offline and replication to the secondary site is adversely affected.
2734745	The PrivNIC resource faults after the <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 1-19](#) describes the Veritas Storage Foundation Cluster File System fixed issues in 6.0.3.

Table 1-19 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues

Incident	Description
2977697	vx_idetach generated kernel core dump while filestore replication is running.
2942776	Mount fails when volumes in vset are not ready.
2923867	Internal test hits an assert "f:xted_set_msg_pri:1".
2923105	The upgrade VRTSvxfs5.0MP4HFaf hangs at vxfs preinstall scripts.
2916691	Customer experiencing hangs when doing dedups.
2906018	The vx_iread errors are displayed after successful log replay and mount of the file system.
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.

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 1-20 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](#)

Issues related to installation and upgrade

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

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.

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

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to Veritas Storage Foundation (SF) 6.0.1, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/openv`), then while upgrading to SF 6.0.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure fileset `VRTSspb`, `VRTSat`, and `VRTSicsco`. This causes NetBackup to stop working.

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

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

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

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

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

Performing an upgrade or rolling upgrade to SF 6.0.3 using NIM ADM may fail if the OS version is incorrect (2869221)

You may see the following error during an upgrade or rolling upgrade using NIM ADM:

```
CPI ERROR V-9-40-4782 Cannot install SFCFSHA on system
sfibmblch4-9-v07 since its oslevel is 6.1 TL 00. Upgrade the system
to 6.1 TL5 or later to install SFCFSHA
```

Workaround:

If you see the above error, upgrade the operating system to the correct technology level (TL5). To check the technology level prior to upgrading, run the `oslevel -s` command.

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
```

The VRTSvxvm fileset fails to install on a few cluster nodes because the template file is corrupted (2348780)

The installer debug log displays the failure of the `errupdate` command as following:
`errupdate -f /usr/lpp/VRTSvxvm/inst_root/VRTSvxvm.err.` The `errupdate` command gets invoked through `/usr/lib/instl/install` by the operating system. The command also fails for the VRTSvxfs, VRTSgIm, and VRTSgms s.

The `errupdate` command generally creates a `*.undo.err` file to remove entries from the Error Record Template Repository in case of failed installation or cleanup. However, in this case the `*.undo.err` file does not get generated as the `errupdate` command fails. Also, it is not possible to manually remove entries from the Error Record Template Repository in order to undo the changes made by the failed installation, because the file is corrupted.

Workaround: Save a copy of the `/var/adm/ras/errtmpl` and `/etc/trcfmt` files before you install the product. Replace `/var/adm/ras/errtmpl` and `/etc/trcfmt` files with the ones that you saved, when the installation fails because the template file is corrupted. Uninstall all the s you installed and reinstall.

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 vxconfig daemon recovery.

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

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

If the CVM protocol version does 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
```

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]
```

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

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

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

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

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

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

Workaround: Complete the rolling upgrade to VCS 6.0.

If VCS is configured on 6.0.3 but not on 6.0.1, HAD/fencing cannot start after rolling back from 6.0.3 to 6.0.1 (3065989)

If VCS is configured on 6.0.3 but not on 6.0.1, HAD/fencing cannot start after rolling back from 6.0.3 to 6.0.1.

Workaround:

1. Manually change the following files:
 - /etc/default/vcs
change VCS_START=0 to VCS_START=1
change VCS_STOP=0 to VCS_STOP=1
 - /etc/default/vxfen
change VXFEN_START=0 to VXFEN_STOP=1
change VXFEN_STOP=0 to VXFEN_START=1
2. Manually execute the following commands to start fencing and VCS:

```
# /etc/methods/vxfenext -start -dvxfen
# /etc/methods/vxfenext -start -dvxfend
# /etc/rc.d/rc2.d/S97vxfen start
# hstart
```

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.

Some paths in DMP can get DISABLED if LVM volume group is created on OS device path (1978941)

On AIX, when an LVM volume group is created directly on the OS device path, the SCSI driver performs SCSI2 reservation on the rest of the paths to that LUN. As a result, some of the paths of the corresponding DMP devices may be disabled, as shown by the `vxddmadm getsubpaths` command output. For some arrays, the `vxdisk list` command shows the device in the 'error' state.

This issue is not seen when LVM volume groups are created on the DMP devices.

Example of this issue:

```
# vxdisk list | grep emc0_00bc
emc0_00bc      auto:none      -              -              online invalid

# vxddmpadm getsubpaths dmpnodename=emc0_00bc
NAME          STATE[A]  PATH-TYPE[M]  CTLR-NAME  ENCLR-TYPE  ENCLR-NAME  ATTRS
=====
hdisk110      ENABLED(A)  -              fscsi0     EMC         emc0        -
hdisk123      ENABLED(A)  -              fscsi0     EMC         emc0        -
hdisk136      ENABLED(A)  -              fscsi1     EMC         emc0        -
hdisk149      ENABLED(A)  -              fscsi1     EMC         emc0        -

# vxdisk rm emc0_00bc

# mkvg -y dmxvg hdisk110
dmxvg

# lspv | egrep "hdisk110|hdisk123|hdisk136|hdisk149"
hdisk110      00c492ed6fbda6e3      dmxvg      active
hdisk123      none
hdisk136      none
hdisk149      none

# vxdisk scandisks

# vxddmpadm getsubpaths dmpnodename=emc0_00bc
NAME          STATE[A]  PATH-TYPE[M]  CTLR-NAME  ENCLR-TYPE  ENCLR-NAME  ATTRS
=====
hdisk110      ENABLED(A)  -              fscsi0     EMC         emc0        -
hdisk123      DISABLED    -              fscsi0     EMC         emc0        -
hdisk136      DISABLED    -              fscsi1     EMC         emc0        -
hdisk149      DISABLED    -              fscsi1     EMC         emc0        -
```

To recover from this situation

1 Varyoff the LVM volume group:

```
# varyoffvg dmxvg
```

2 Remove the disk from VxVM control.

```
# vxdisk rm emc0_00bc
```

3 Trigger DMP reconfiguration.

```
# vxdisk scandisks
```

4 The device which was in DISABLED state now appears as ENABLED.

```
# vxddmpadm getsubpaths dmpnodename=emc0_00bc
```

NAME	STATE [A]	PATH-TYPE [M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
hdisk110	ENABLED (A)	-	fscsi0	EMC	emc0	-
hdisk123	ENABLED (A)	-	fscsi0	EMC	emc0	-
hdisk136	ENABLED (A)	-	fscsi1	EMC	emc0	-
hdisk149	ENABLED (A)	-	fscsi1	EMC	emc0	-

DS4K series array limitations

In case of DS4K array series connected to AIX host(s), when all the paths to the storage are disconnected and reconnected back, the storage does not get discovered automatically. To discover the storage, run the `cfgmgr` OS command on all the affected hosts. After the `cfgmgr` command is run, the DMP restore daemon brings the paths back online automatically in the next path restore cycle. The time of next path restore cycle depends on the restore daemon interval specified (in seconds) by the tunable `dmp_restore_interval`.

```
# vxddmpadm gettune dmp_restore_interval
```

Tunable	Current Value	Default Value
dmp_restore_interval	300	300

On DS4K array series connected to AIX host(s) DMP is supported in conjunction with RDAC. DMP is not supported on DS4K series arrays connected to AIX hosts in MPIO environment.

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.

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:

```
# vxddmpadm settune dmp_monitor_ownership=off
```

The `dmp_monitor_fabric` is not persistent across reboot and upgrades (2975623)

The `dmp_monitor_fabric` parameter is not persistent across reboot and upgrades. Even if you have changed its value, it will be changed back to the previous value after system reboot or product upgrade.

Workaround:

Change it again after system reboot or product upgrade.

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).

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.

AT Server crashes when authenticating unixpwd user multiple times (1705860)

There is a known issue in the AIX kernel code that causes 'getgrent_r' function to corrupt the heap. This issue is present in AIX 5.3 and AIX 6.1 Refer to IBM's Web site for more information:

<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52585>

AT uses `getgrent_r` function to get the groups of the authenticated user.

IBM has released the fix as a patch to fileset `bos.rte.libc`. There are different patches available for different version of `bos.rte.libc`. You need to check the version of `bos.rte.libc` (For example: `lsbpp -l | grep bos.rte.libc`) and apply the appropriate IBM patch:

- For version 6.1.3.1:
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52959>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52959/>
- For version 6.1.2.4:
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52720>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52720/>
- For version 6.1.2.5 :
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52975>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52975/>

There are IBM patches for only certain version of `bos.rte.libc` that are available. If your system has a different `bos.rte.libc` version, you may have to upgrade to a higher version where the fix is available. If your version is not available, you may have to contact IBM.

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

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

Running AIX 6.1, you may receive the following error message when using sqlplus:

```
$ sqlplus " / as sysdba"
SQL> startup nomount
SQL> ORA 0-0-0-0
```

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 Deployed 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/rac11g1/.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'rac11g1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdg ERROR V-5-1-582 Disk group SNAP_rac11dgl: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dgl.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dgl 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.

Upgrading operating system Technology Levels along with Storage Foundation using an alternate disk fails (2162945)

Upgrading the operating system Technology Levels (TL) along with Storage Foundation using an alternate disk fails occasionally with the following error:

```
alt_disk_copy: 0505-224 ATTENTION:
An error occurred during installation of
one or more software components.
Modifying ODM on cloned disk.
Building boot image on cloned disk.
forced unmount of /alt_inst/var/adm/ras/platform
forced unmount of /alt_inst/var
umount: error unmounting /dev/alt_hd2: Device busy
0505-144 alt_disk_install: Unable to unmount alt_inst filesystems.
```

No issues have been observed with Storage Foundation in the cause of the failure.

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.

The cluster may hang if a node goes down (1835718)

The cluster may hang if a node goes down while one array is disabled or offline in a mirror=enclosure configuration.

This may occur, if a node panics or loses power while one array of a mirror=enclosure configuration is offline or disabled, then the cluster, fencing, I/O loads, and VxVM transactions hang.

Workaround: There is no workaround for this issue.

vxconvert failures if PowerPath disks are formatted as simple disks (857504)

If a PowerPath disk is formatted as a simple disk (a foreign device), then the `vxconvert` utility may fail during conversion of LVM to VxVM. To view the format of the disk, use the `vxdisk list` command. This issue may also occur if the `/etc/vx/darecs` file contains an `hdiskpower` disk entry. This entry may be present if PowerPath disks were configured as foreign disks in Storage Foundation 4.0, and the entry was not changed after subsequent upgrades.

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
```

Co-existence check might fail for CDS disks

In Veritas Volume Manager (VxVM) 5.1 SP1, VxVM introduces the ability to support Cross-platform Data Sharing (CDS) on disks larger than 1 TB. VxVM uses the SUN VTOC Table to initialize the cdsdisk layout on devices up to 1 TB. VxVM uses the GUID Partition Table (GPT) to initialize the cdsdisk layout on devices larger than 1 TB.

In layouts where SUN VTOC Table is used for initialization (typically, when the disk size has never exceeded 1 TB), the AIX co-existence label can be found at sector 7 and VxVM ID block (also known as HP co-existence label) can be found at sector 16.

In layouts where GPT is used for initialization (typically, when the disk size is currently greater than or had earlier exceeded 1 TB), the AIX co-existence label is placed at sector 55 and VxVM ID block (also known as HP co-existence label) is placed at sector 64. Consequently, AIX utilities would not be able to recognize a cdsdisk initialized using GPT to be a valid VxVM disk. Symantec is working with IBM and third party OEMs to enhance the co-existence check in these utilities.

Workaround:

There is no workaround for this issue.

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

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

Workaround:

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

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

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In 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.

Recovery and rollback to original configuration may not succeed if the system reboots while the online migration setup is in partial state (2611423)

During online migration from LVM to VxVM volumes, if there is a system reboot when the migration setup is in partial state, that is, the start operation has not completed successfully, then the recover and abort operations might not be able to recover and rollback the configuration.

Workaround: This needs manual intervention for cleanup, depending on the state, to restore the original configuration.

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.

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.

Upgrading SF to version 6.0 marks vSCSI disks as cloned disks (2434444)

This issue is seen when you upgrade from a previous version of SF which has vSCSI disks included in a disk group. After upgrading SF to 6.0, the vSCSI disks that were included in a disk group are marked as cloned disks.

Workaround:

Use the following procedure to clear the clone disk flag.

To clear the clone disk flag

- 1 Remove the vSCSI devices that are in error state (ibm_vscsi#_#) using the following command:

```
# vxdisk rm device_name
```

- 2 Deport the disk group.

```
# vxdg deport dg_name
```

- 3 Re-import the disk group with a new udid.

```
# vxdg -o updateid import dg_name
```

- 4 Display the devices that are part of the disk group.

```
# vxdisk -g dg_name list
```

- 5 Clear the clone_disk tag from these devices.

```
# vxdisk set device_name clone=off
```

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`.

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
```

vxconfigd hang with path removal operation while IO is in-progress (1932829)

In AIX with HBA firmware version SF240_320, `vxdisk scandisks` (device discovery) takes a long time when a path is disabled from the switch or from the array.

Workaround:

To resolve this issue, upgrade the HBA firmware version to SF240_382.

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.

Required attributes of LUNs for DMP devices with cluster set-up having fencing enabled (2521801)

When cluster set-up has fencing enabled, the following attributes are required to be set on the LUNs.

Set the following attributes for LUNs**1 Set the following attributes:**

- If the path has the `reserve_policy` attribute set, change the `reserve_policy` attribute to `no_reserve` for all the paths.

```
# lsattr -El hdisk557 | grep res  
reserve_policy single_path  
Reserve Policy True
```

```
# chdev -l hdisk557 -a reserve_policy=no_reserve -P
hdisk557 changed
```

- If the path has the `reserve_lock` attribute set, change the `reserve_lock` attribute to `no`.

```
# lsattr -El hdisk558 | grep reserve_lock
reserve_lock    yes
Reserve Device  on open True

# chdev -l hdisk558 -a reserve_lock=no -P
hdisk558 changed
```

2 Reboot the system for the changes to take effect.

If SF is installed in global system and WPAR is created with the `-t` option, `/usr/sbin/sync` gets removed from WPAR (3007163)

With AIX 6.1 TL8 and AIX 7.1 TL2, when WPARs are created with the new option `-t`, all the file systems are copied from global to WPAR. Also the files and filesets that are not visible in WPAR are deleted, so `sync` file gets deleted as part of this operation.

Any utilities that use `sync` internally might get the following error:

```
# oslevel -s
rpm_share: 0645-025 Unable to locate command module sync.
rpm_share: 0645-007 ATTENTION: init_baselib() returned an
unexpected result.
7100-02-01-1245
```

Workaround :

Use one of the two ways to avoid this issue:

- Append `-T preserve_private=yes` option while creating WPAR with the new option `-t`.
- Copy `/usr/sbin/sync` from global system to WPAR.

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.

Veritas File System known issues

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

Cannot use some commands from inside an automounted Storage Checkpoint (2490709)

If your current work directory is inside an automounted Storage Checkpoint, for example `/mnt1/.checkpoint/clone1`, some commands display the following error:

```
can't find current directory
```

This issue is verified with the following commands:

- `cp -r`
- `du`

However, this issue might occur with other commands.

Workaround: Run the command from a different directory.

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.

Performance on a VxFS file system can be slower than on a JFS file system (2511432)

At times, the performance on a VxFS file system can be slower than on a JFS file system.

Workaround:

There is no workaround for this issue.

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.

You are unable to unmount the NFS exported file system on the server if you run the fsmigadm command on the client (2355258)

Unmounting the NFS-exported file system on the server fails with the "Device busy" error when you use the `fsmigadm` command on the NFS client.

Workaround:

Unexport the file system prior to unmounting.

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/rdisk/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.

The vxcompress operation with the multithreaded option takes longer than the single threaded one (3031878)

The `vxcompress` operation with the multithreaded option takes longer than the single threaded one.

Workaround:

There is no workaround for this issue.

The `fsvoladm` command fails to clear the `metadataok` flag (2999560)

The `fsvoladm` command fails to clear the `metadataok` flag on some volumes of a VXFS file system, which is mounted on a volume set with 6 volumes.

Workaround:

There is no workaround for this issue.

Replication known issues

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

`vradmin syncvol` command compatibility with IPv6 addresses (2075307)

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

Workaround:

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

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

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

Workaround:**To resolve this issue**

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

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

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

Workaround:

Destroy the instant snapshots manually using the `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 -V 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 -V vxfs /dev/vx/dsk/dg/data_volume
```

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

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

environment, `vradmin createpri` may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

Workaround:

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

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

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

```
vradmin 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 `vradmin` on all the cluster nodes using the following commands:

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

While vradmin commands are running, vradmin 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, `vradmin` may temporarily lose heartbeats, and the commands terminate with the following error message:

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

Workaround:

To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

vxassist relayout removes the DCM (145413)

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

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

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

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

Workaround:

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

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

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

- 3 Stop the affected RVG. Enter the following:

```
# 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 rvg vol
```

- 7 Start the RVG. Enter the following:

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

- 8 Resume or start the applications.

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

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

```
VxVM VVR vxrsync ERROR V-5-52-2222 [from host]: VxVM in.vxrsyncd
ERROR V-5-36-2125 Server volume access error during [assign volids]
volume path: [/dev/vx/dsk/dg/snapshot_volume] reason: [this could be
because a target volume is disabled or an rlink associated with a
target volume is not detached during sync operation].
```

Workaround:

There are two workarounds for this issue.

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

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

When VVR is used for asynchronous replication in shared disk group environment, one of the nodes of the cluster at the primary site is chosen as the logowner. When the logowner node is on a node which is a slave node for the underlying CVM cluster, in the presence of heavy I/O from a node that is not the logowner, it is

possible to get into a replication hang. This is due to an internal defect which will be fixed in later releases.

Workaround:

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

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

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

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

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

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

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

Workaround:

Manually recover the layered volumes using the vxvol utility, as follows:

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

vxassist and vxresize operations do not work with layered volumes that are associated to an RVG (2162579)

This issue occurs when you try a resize operation on a volume that is associated to an RVG and has a striped-mirror layout.

Workaround:

To resize layered volumes that are associated to an RVG

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


```
# vxrlink -g diskgroup status rlink
```
- 3 Stop the affected RVG. Enter the following:


```
# vxrvrg -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:


```
# vxrvrg -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 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.

dbed_update command failed after upgrading a Storage Foundation product from 5.1SP1RP1 to 6.0.1 on AIX 6.1 (2846434)

dbed_update might fail under some Oracle configurations, even when the database is up and running, with the following error message.

```
dbed_update -S apr1 -H /opt/oracle/app/oracle/product/11.2/db_1
```

```
SFDB vxsfadm ERROR V-81-0564 Oracle returned error.
```

```
Reason: ORA-01041: internal error. hostdef extension doesn't exist  
(DBD ERROR: OCISessionBegin)
```

You are able to connect to the database manually using `sqlplus`. This problem is because the version of `DBD::Oracle` perl module, used by the SFDB tools, uses somewhat older Oracle instant client libraries. With these, the SFDB tools are unable to connect to the Oracle database even when the database is up and running.

Workaround: There is no workaround at this point of time.

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, up to 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 umount: ERROR: V-3-26388: file system /rep has been mount  
locked  
SFORA sfua_rept_migrate ERROR V-81-5550 umount /dev/vx/dsk/repdg/repvol  
failed.  
SFORA sfua_rept_migrate ERROR V-81-9162 Failed to umount repository.
```

Workaround:

The error does not hamper the upgrade. The repository migration works fine, but the old repository volume does not get unmounted. Unmount the mount using the manual option.

For example, use `/opt/VRTS/bin/umount -o mntunlock=VCS /rep`.

For more information, see [TECH64812](#).

Veritas Cluster Server known issues

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

- [Operational issues for VCS](#)
- [Issues related to the VCS engine](#)
- [Issues related to the VCS database agents](#)
- [Issues related to the agent framework](#)
- [Issues related to global clusters](#)
- [LLT 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 \(Java Console\)](#)

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

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

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

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

Connecting to the database outside VCS control using sqlplus takes too long to respond [704069]

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

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.

Issues related to the VCS engine

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

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

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

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

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

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

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

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

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

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

If the top level resource which does not have any parent 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 `hargp -online -sys` command [2871892]

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

Workaround: No workaround.

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

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

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

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

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

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

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

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

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

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

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

Workaround

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

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

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

Workaround: No workaround.

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

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

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

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

or

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

GCO clusters remain in INIT state [2848006]

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

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

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

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.

After rollback, fencing/HAD fails to start if CPS based fencing is configured (3074644)

Fencing/HAD fails to start in the following scenarios:

- Scenario 1:
 1. CPS based fencing is configured on 6.0.1.
 2. Upgrade to 6.0.3.
 3. Rollback from 6.0.3 to 6.0.1.
- Scenario 2:
 1. CPS based fencing is not configured on 6.0.1.
 2. Upgrade to 6.0.3.
 3. Rollback from 6.0.3 to 6.0.1.
 4. The system is not rebooted.
 5. CPS based fencing is configured on 6.0.1.

This issue is caused because the vxfeed module doesn't get loaded after rollback with CPI.

Workaround:

After rollback,

For Scenario 1, manually execute the following commands on each system of the cluster:

```
# /etc/methods/vxfenext -start -dvxfen
# /etc/methods/vxfenext -start -dvxfend
# /etc/rc.d/rc2.d/S97vxfen start
# /opt/VRTSvcs/bin/hastart
```

For Scenario 2, manually execute the following commands on each system of the cluster:

```
# /etc/methods/vxfenext -start -dvxfen
# /etc/methods/vxfenext -start -dvxfend
```

Issues related to the bundled agents

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

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

Workaround: Set the LinkTestRatio attribute to 0.

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

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

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

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.

Bringing the LPAR resource offline may fail [2418615]

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

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

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

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

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

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

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

RemoteGroup agent does not failover in case of network cable pull [2588807]

A RemoteGroup resource with ControlMode set to OnOff may not fail over to another node in the cluster in case of network cable pull. The state of the RemoteGroup resource becomes UNKNOWN if it is unable to connect to a remote cluster.

Workaround:

- Connect to the remote cluster and try taking offline the RemoteGroup resource.
- If connection to the remote cluster is not possible and you want to bring down the local service group, change the ControlMode option of the RemoteGroup resource to MonitorOnly. Then try taking offline the RemoteGroup resource. Once the resource is offline, change the ControlMode option of the resource to OnOff.

CoordPoint agent remains in faulted state [2852872]

The CoordPoint agent remains in faulted state because it detects `rfsm` to be in replaying state.

Workaround: After HAD has stopped, reconfigure fencing.

MemCPUAllocator agent fails to come online if DLPAR name and hostname do not match [2954312]

If hostname of the DLPAR and name of DLPAR as seen from HMC are different, the MemCPUAllocator agent is unable to provide CPU or memory to the DLPAR.

Workaround: Change the name of DLPAR from HMC to match the hostname.

HA commands inside WPAR agent get stuck due to the login/password prompt [2431884]

After upgrade of secure clusters from VCS versions lower than VCS 6.0, the HA commands that run from within the WPAR display login/password prompts. Hence, agents trying to run HA commands inside WPAR get stuck because of the prompt, as the WPAR credentials are not upgraded because of change of architecture of VxAT in VCS 6.0.

Workaround: Run `hawparsetup.pl` again for each WPAR resource. This will create new credentials for the WPAR which can be used by HA commands in VCS 6.0.

NFS client reports I/O error because of network split brain [2564517]

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

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

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


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

- 2 Enable PREONLINE trigger for the service group.

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

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

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

Workaround:

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

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

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

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

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

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

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

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

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

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

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.

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

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

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

Workaround: No workaround.

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

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

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

When two IPs of different subnets are assigned to a virtual NIC, the NIC resource might go into faulted state (2919101)

When two IPs of different subnets are assigned to a virtual NIC, the NIC resource might go into faulted state.

Workaround:

Change the order of plumbing the IPs on the interface. The base IP of the interface which belongs to the same subnet as the NetworkHosts (of NIC resource) should be plumbed last.

For example, for the following configuration of `nic1` resource of NIC agent:

```
NIC nic1(  
    Device = en0  
    NetworkHosts = { "10.209.76.1" }  
    Protocol = IPv4  
)
```

The IPs "10.209.78.46" and "192.168.1.29" are plumbed on `en0`. The order of plumbing IPs on the device should be "192.168.1.29" first and "10.209.78.46" later. This is because the NetworkHosts IP (10.209.76.1) and 10.209.78.46 are of the same subnet.

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.

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.

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

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

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

Issues related to the agent framework

Issues with configuration of resource values (1718043)

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

```
VCS WARNING V-16-2-13806 Thread(1437547408) ArgListValues overflow;  
  
Cannot append values more than upper limit of (425).
```

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

Agent may fail to heartbeat under heavy load [2073018]

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

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

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

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

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

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

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

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

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

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

Workaround: No workaround.

Issues related to global clusters

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

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

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

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

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

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

LLT known issues

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

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

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

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

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

```
# lsattr -El SEA
```

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

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

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

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

GAB known issues

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

Cluster panics during reconfiguration (2590413)

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

Workaround: There is no workaround for this issue.

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

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

```
GAB INFO V-15-1-20239
```

```
Client GabTestDriver with refcount 2 forcibly deinited on user request
```

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

Workaround: There is no workaround for this issue.

I/O fencing known issues

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

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 vxfsnwap utility the CoordPoint agent may fault (3462738)

After you run the `vxfsnwap` 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 vxfsnwap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)

The `vxfsnwap` utility runs the `vxfsnconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxfsnwap` 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, `vxfsnwap` 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 `vxfsnswap` 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 `vxfsnswap` 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 vxfsnconfig ERROR V-11-2-1007 Vxfen already configured
```

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

```
VXFEN vxfsnadm 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 vxfsnsthdw utility directly from the install media if VRTSvxfsn package is not installed on the system (2858190)

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

Workaround: Install VRTSvxfsn 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.

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

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

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

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

Issues related to Intelligent Monitoring Framework (IMF)

Process offline monitoring issues with Asynchronous Monitoring Framework[3540463]

Process offline monitoring registrations with Asynchronous Monitoring Framework (AMF) freeze nodes of Symantec Cluster Server (VCS) on AIX 7.1 TL3 SP3 and AIX 6.1 TL9 SP3 versions.

Workaround: Install the IBM authorized program analysis report (APAR) IV63274 on all nodes to fix this issue. For information related to APAR, contact IBM Support.

You can also refer to the following TechNote:

<http://www.symantec.com/docs/TECH222286>

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 CFSSMount 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.

Issues related to the Cluster Manager (Java Console)

This section covers the issues related to the Cluster Manager (Java Console).

Some Cluster Manager features fail to work in a firewall setup [1392406]

In certain environments with firewall configurations between the Cluster Manager and the VCS cluster, the Cluster Manager fails with the following error message:

```
V-16-10-13 Could not create CmdClient. Command Server  
may not be running on this system.
```

Workaround: You must open port 14150 on all the cluster nodes.

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.

The vxfsckd resource fails to start when vxfsckd is killed manually and the cluster node is rebooted (2720034)

If you kill the `vxfsckd` resource manually and reboot the node, `vxfsckd` does not come up and the cvm services are faulted.

Workaround:

Use the following commands for this situation:

```
hastop -local  
rm /var/adm/cfs/vxfsckd-pid
```

Kill all `vxfsckd` processes:

```
fsclustadm cfsdeinit  
hastart
```

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.

There is no workaround at this time.

The mount command may hang when there are large number of inodes with extops and a small vxfs_ninode, or a full fsck cannot fix the link count table corruptions (2689326)

You might encounter one of the following issues:

- If there are large number of inodes having extended operations (extops), then the number of inodes used by the `mount` command reaches the maximum number of inodes that can be created in core. As a result, the `mount` command will not get any new inodes, which causes the `mount` command to run slowly and sometimes hang.

Workaround: Increase the value of `vxfs_ninode`.

- The link count table (LCT) file can get damaged such that the flag is set, but the attribute inode is already freed. In this case, the mount command tries to free an inode that has been already freed thereby marking the file system for a full structural file system check.

Workaround: There is no workaround for this issue.

You sometimes receive shell error messages (2172138)

You sometimes receive shell error messages while adding a node into an existing SF cluster. The following is a sample of the shell error message that you can receive:

```
sh[2]: sw: not found
```

You can safely ignore these error messages.

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.

The installer may fail to mount some share disk groups (2167226)

The `installer` fails to mount some share disk groups if its name is a substring of other disk groups.

Workaround

You need to manually add those share disk groups to the newly added nodes. Or avoid naming your share disk groups that could be substring of others.

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)

Stale NFS file handle

Workaround: There is no workaround for this issue.

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.

Inode access and modification times are not getting updated on the primary node when a file owned by the primary node is accessed from a secondary node (2170318)

The inode access times and inode modification itimes (collectively known as itimes) are not getting updated on the primary node when a file owned by the primary node is accessed from a secondary node. The primary node has a stale value for those itimes. A cluster file system requires consistent itimes on all the nodes at the same time. The system performance has a minimal impact even if itimes are not same on all nodes.

Workaround: There is no workaround for this issue.

File system check daemon fails to restart after abnormal termination (2942464)

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

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

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

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

3. Remove the `vxfsckd-pid` file:

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

4. Bring the `vxfsckd` resource online:

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

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.

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

Oracle VIP Configuration Assistant fails with an error message (1182220)

During Oracle RAC 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

```
The given interface(s), "" is not public.  
Public interfaces should be used to configure virtual IPs.
```

This message appears only when the VIP is not from the regular public IP range (for example, 200.).

Workaround: Invoke the `vipca` utility manually as the superuser.

```
# export DISPLAY=nebula:0.0  
# $CRS_HOME/bin/vipca
```

Oracle Cluster Verification utility displays a warning message

During the final stage of Oracle RAC 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification utility.

For example:

```
Utility  
=====
```

OUI-25031: Some of the configuration assistants failed. It is strongly recommended that you retry the configuration assistants at this time. Not successfully running any "Recommended" assistants means your system will not be correctly configured.

1. Check the Details panel on the Configuration Assistant Screen to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button to retry them.

=====

Workaround: You may safely ignore this message if the cluster is operating satisfactorily.

11.2.0.2 is not supported on AIX 7.1 TL2 SP1 (3063257)

Without IBM fix APAR IV35893 pack on AIX7.1 TL2 SP1, 11.2.0.3 CRS configuration may fail with the following CRS messages on the second node :

```
CRS-4000: Command Start failed, or completed with errors.
The exclusive mode cluster start failed, see Clusterware alert
log for more information
Initial cluster configuration failed. See
/crsbin/crshome/cfgtoollogs/crsconfig/rootcrs_swaxp720-1-v02.log
for details
/crsbin/crshome/perl/bin/perl -I/crsbin/crshome/perl/lib
-I/crsbin/crshome/crs/install /crsbin/crshome/crs/install/rootcrs.pl
execution failed
```

Workaround:

To resolve this issue, use any of the following solutions:

- Apply the Oracle 11gr2 patch to fix the issue [ID 1352887.1].
- Prepare the setup again with CRS 11.0.2.3.
- Install IBM fix APAR IV35893 pack on AIX7.1 TL2 SP1.

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.

Software limitations

This section covers the software limitations in 6.0.3 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](#)

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.

Limitation with device renaming on AIX 6.1TL6

If you rename an operating system (OS) path with the `rendev` command on AIX 6.1TL6, the operation might remove the paths from DMP control. DMP cannot discover these paths.

Upgrade of secure clusters not supported using native operating system tools

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

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 1-21](#) describes the DMP tunable parameters and the new values.

Table 1-21 DMP settings for NetApp storage attached environment

Parameter name	Definition	New value	Default value
<code>dmp_restore_interval</code>	DMP restore daemon cycle	60 seconds.	300 seconds.
<code>dmp_path_age</code>	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
```

DMP support in AIX virtualization environment (2138060)

DMP does not support exporting paths to the same LUN through both vSCSI and NPIV interfaces.

DMP treats the same LUN seen through vSCSI and NPIV interfaces as two separate LUNs, because the behavior of the LUN at the VIOC level is different due to the intermediate SCSI interface at the VIOS level for vSCSI devices.

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	thinrc1m
xiv0_613	19313	2108	dg1	thinrc1m
xiv0_614	19313	35	dg1	thinrc1m
xiv0_615	19313	32	dg1	thinrc1m
xiv0_616	19313	31	dg1	thinrc1m
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.

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.

The shell cannot handle 64-bit inode numbers inside the .checkpoint directory when uniqueino is enabled

Due to a limitation with the AIX operating system, the shell cannot handle the 64-bit inode numbers inside the `.checkpoint` directory when the `uniqueino` mount option is enabled. Some shell functions such as auto-complete and globs, for example `rm *`, do not function properly in the `.checkpoint` directory. This also affects 32-bit applications that try to read the contents of the `.checkpoint` directory or any of its subdirectories. This does not affect any 64-bit applications.

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.

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 and 6.0.1.

Limitations related to installing and upgrading VCS

Limitation on upgrading to 6.0.3 on a Veritas Storage Foundation and High Availability cluster

Veritas Storage Foundation (SF) 6.0.3 requires the AIX operating system to be at 6.1 TL5 or above. To upgrade SF to 6.0.3 from a release prior to 5.0 MP3 RP1, you must first upgrade SF to the 5.0 MP3 RP1 release. If upgrading to 5.0 MP3 RP1 requires an intermediate operating system upgrade, the operating system level cannot exceed 6.1 TL1. After upgrading to 5.0 MP3 RP1, you must upgrade the operating system to AIX 6.1 TL5, which is the minimum requirement for the 6.0.3 release. You must upgrade SF to 5.0 MP3 RP1 to avoid a system panic or crash that can occur when a node running AIX 6.1 TL2 or above with a release prior to 5.0 MP3 RP1 is removed from the Veritas Storage Foundation and High Availability cluster. Removing the node causes file system threads to exit. The panic is caused by a check introduced from AIX 6.1 TL2 that validates the lockcount values when a kernel-thread-call exits.

For more information, see the following TechNote:

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

Limitations related to bundled agents

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/netsvc.conf

Volume agent clean may forcibly stop volume resources

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

False concurrency violation when using PidFiles to monitor application resources

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

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

WPAR agent registered to IMF for Directory Online event

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

Application agent limitations

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

Limitations related to IMF

- If a process is registered with IMF for offline monitoring, IMF may not detect the process being executed if the length of the process and related arguments exceed 80 characters. This limitation affects Application agent and Process agent. Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for more information. (2768558)

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-0909db2whichedition/> 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]

Systems in a cluster must have same system locale setting

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

Limitations with DiskGroupSnap agent [1919329]

The DiskGroupSnap agent has the following limitations:

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

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

Virtualizing shared storage using VIO servers and client partitions

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

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

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

Supported storage

Refer to the IBM data sheet:

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

Disk Restrictions

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

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

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

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

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

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

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

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

Get the serial number of the disk of interest:

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

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

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

List the virtual SCSI adapters.

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

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

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

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

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

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

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

Next create a virtual device to map hdisk20 to vhost1.

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

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

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

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

Cluster Manager (Java console) limitations

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

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

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

See the *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".

Limited support from Cluster Manager (Java console)

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

Port change required to connect to secure cluster [2615068]

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

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

In a dual-stack configuration, when you use packet counts and the IPv6 network is disabled, the NIC agent might not detect a faulted NIC. It might not detect a fault

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

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

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

Limitations related to LLT

This section covers LLT-related software limitations.

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

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

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

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

LLT does not start automatically after system reboot (2058752)

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

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

Workaround: To resolve the LLT startup issue

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

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

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

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.

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

For multi-pathing to connected storage, AIX uses the RDAC driver for FASTT arrays. Since it is an active/passive array, only the current active path is exposed to clients. The I/O fencing driver, vxfen, can use only a single active path and has no foreknowledge of the passive paths to the coordinator disks on an array. If the single active path fails, all nodes in the cluster lose access to the coordinator disks.

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

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

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *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.

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.

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`, `Imissing` 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 CFSMount resource will not come online. You can check the VCS engine log file for any mount failure messages.

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.

Web installation program does not support phased upgrade and native operating system upgrade mechanisms

The Web installation program does not support phased upgrade and native operating system upgrade mechanisms such as Alternate Disk Installation (ADI) and Network Installation Manager Alternate Disk Migration (NIMADM).

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.

Documentation errata

There are no documentation errata updates.

List of patches

This section lists the patches for 6.0.3.

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

Table 1-22 Patches for AIX

BFF file	size in bytes	Patches	Version
VRTSamf.bff	12134400	VRTSamf	06.00.0300.0000
VRTScavf.bff	972800	VRTScavf	06.00.0300.0000
VRTSdbed.bff	126771200	VRTSdbed	06.00.0300.0000
VRTSperl.bff	37836800	VRTSperl	05.14.0002.0008
VRTSsfcp601.bff	4915200	VRTSsfcp601	06.00.0300.0000
VRTSvcsv.bff	360448000	VRTSvcsv	06.00.0300.0000
VRTSvcsvsag.bff	19712000	VRTSvcsvsag	06.00.0300.0000
VRTSvcsea.bff	6502400	VRTSvcsea	06.00.0300.0000
VRTSvxfen.bff	3174400	VRTSvxfen	06.00.0300.0000
VRTSvxfs.bff	42905600	VRTSvxfs	06.00.0300.0000
VRTSvxvm.bff	285286400	VRTSvxvm	06.00.0300.0000

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 Veritas Storage Foundation and High Availability Solutions product for the first time on a host and install 6.0.3. 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 installed, you must upgrade to 6.0.3.

See [“Upgrading to 6.0.3”](#) on page 131.

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 and the uninstallmr scripts](#)” 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

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 131.

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 `xprtlwid`, 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 130.

- 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](#)
- [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.
- 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 130.

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

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

Upgrading to 6.0.3

This section describes how to upgrade from 6.0.1 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), or Veritas Storage Foundation for Oracle RAC (SFRAC) 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.
- [Performing a rolling upgrade using the `installmr` script](#)
Use the procedure to upgrade your Veritas product with a rolling upgrade.

See “[Installing the Veritas software using the script-based installer](#)” on page 126.

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 a Veritas Cluster Server

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

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 130.
- 2 Log in as superuser.

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

See [“System requirements”](#) on page 18.

- 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 and the uninstallmr scripts”](#) on page 10.

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

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

See [“About the installmr and the uninstallmr scripts”](#) 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 130.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 4 On each node, enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 5 Unmount all Storage Checkpoints and file systems:

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

- 6 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 `vxrvg 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.

- 7 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.

- 8 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
```

- 9 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
```

- 10 Log in as superuser.

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

See [“System requirements”](#) on page 18.

- 11 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 [-rsh] node1node2 ... nodeN
```

The program proceeds in a noninteractive mode to examine the systems for licenses, filesets, disk space, system-to-system communications, etc.

- 12 Review the output as the program displays the results of the check and saves the results of the check in a log file.
- 13 Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved.

See [“Prerequisites for upgrading to 6.0.3”](#) on page 130.

- 14 Start the upgrade.

```
# ./installmr [-rsh] node1 node2 ... nodeN
```

Review the output.

- 15 Restart all the volumes. Enter the following command for each disk group:

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

- 16 If you stopped any RVGs in step 6, restart each RVG:

```
# vxrvvg -g diskgroup start rvg_name
```

- 17 Remount all VxFS file systems on all nodes in the selected group:

```
# mount /filesystem
```

- 18 Remount all Storage Checkpoints on all nodes in the selected group:

```
# mount /checkpoint_name
```

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 130.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 4 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
```

- 5 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 that are not under VCS control:

```
# umount /filesystem
```

- 6 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 `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -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.

- 7 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.

- 8 On each node, 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
```

- 9 If required, apply the OS kernel patches.

See [“System requirements”](#) on page 18.

See IBM’s documentation for the procedures.

- 10 On each node, 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 Log in as superuser.

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

See [“System requirements”](#) on page 18.

- 12 From the directory that contains the extracted and untarred 6.0.3 rolling patch binaries, change to the directory that contains the `installmr` script.

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

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

- 13 After all the nodes in the cluster are upgraded, the processes restart. If the `installmr` script finds issues, it may require you to reboot the nodes.
- 14 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.
- 15 Bring the CVM service group online on each node:

```
# hagrps -online cvm -sys nodename
```

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

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

- 17 If you stopped any RVGs in step 6 , restart each RVG:

```
# vxrvrg -g diskgroup start rvrg_name
```

- 18 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

- 19 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 a 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 130.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your **PATH** so that you can execute all product commands.
- 4 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.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
```

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

```
# haconf -makerw
```

And set the the AutoStart attribute of Oracle Agent to 0:

```
# hagrps -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- 6 If the Oracle DB is not managed by VCS, prevent auto startup of Oracle DB:

```
# srvctl modify database -d db_name -y manual
```

- 7 Stop Oracle database on the cluster:
 - If the Oracle RAC instance is managed by VCS:

```
# hagrpl -offline oracle_group -sys galaxy
# hagrpl -offline oracle_group -sys nebula
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and shut down the instances:

```
$ srvctl stop database -d db_name
```

- 8 Stop all applications on the cluster that are not configured under VCS. Use native application commands to stop the application.
- 9 Unmount the VxFS and CFS file systems that are not managed by VCS.

- Ensure that no processes are running that make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point, enter the following commands:

```
# mount | grep vxfs
# fuser -cu /mount_point
```

- Unmount the VxFS or CFS file system:

```
# umount /mount_point
```

- 10 Stop all VxVM and CVM volumes for each diskgroup that are not managed by VCS on the cluster:

```
# vxvol -g disk_group stopall
```

Verify that no volumes remain open:

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

- 11 Stop VCS.

```
# hastop -all
```

- 12 If you plan to upgrade the operating system,

- Stop all processes with below command:

```
# ./installsfrac -stop
```

- Rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- Upgrade the operating system on all nodes in the cluster.
For instructions, see the operating system documentation.
- After the system restarts, restore the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 13 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 ...
```

- 14 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 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 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.
- 17 Start Oracle Group on All nodes.

```
# hagrps -online oracle_group -any
```

- 18 If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and start the instances:

```
$ srvctl start database -d db_name
```

- 19 ■ 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
# hagrps -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
```

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 130.

- 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.

See [“System requirements”](#) on page 18.

See IBM’s documentation for the procedures.

- 5 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | 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:

```
# vrxvg -g diskgroup stop rvg_name
```

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

```
# vxlink -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 Copy the patch archive downloaded from the patch central to temporary location and untar the archive and browse to the directory containing the installmr installer script. Enter 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
```

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

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 130.

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 and SFCFSA: 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 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 5 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.
- 6 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.
- 7 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.

- 8 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.
- 9 The installer stops relevant processes, uninstalls old kernel fileset, and installs the new fileset. When prompted, enable replication or global cluster capabilities, if required, and register the software.
- 10 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.
- 11 Before you proceed to phase 2, complete step 1 to step 10 on the second subcluster.

Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSA: 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 fileset 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 fileset, and installs the new fileset. 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 SF Oracle RAC: phase 1

Note that in the following instructions that a subcluster can represent one or more nodes in a full cluster, but is represented by nodeA, nodeB as subcluster1 and nodeC, nodeD as subcluster2.

To perform the rolling upgrade on kernel: phase 1

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: main.cf, types.cf, CVMTypes.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 \
    /var/tmp/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
    /var/tmp/MultiPrivNIC.cf.save
```

- 3 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
```

If the Oracle database is not managed by VCS, change the management policy for the database to manual. Execute the command with oracle database user credentials.

```
$ srvctl modify database -d db_name -y manual
```

- 4 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 6.0.3.
- 5 If the applications are not under VCS control, stop the applications that use VxFS or VxVM disk groups on each node of subcluster, whether local or CFS. Use native application commands to stop the application.
- 6 If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrpl -offline grp_name -sys node_name
```

If the database instances are not managed by VCS, stop the Oracle RAC database resources on each node of subcluster, run the following from each node of subcluster. Execute the command with oracle database user credentials.

```
$ srvctl stop instance -d db_name -i instance_name
```

- 7 Unmount all the VxFS file system which is not under VCS control on each node of subcluster.

```
# mount |grep vxfs
# fuser -c /mount_point
# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -c /mount_point
```

- 8 On subcluster, stop all VxVM and CVM volumes (for each disk group) that are not managed by VCS:

```
# vxvol -g disk_group stopall
```

Verify that no volumes remain open under the diskgroups which are not managed by VCS:

```
# vxprint -g disk_group -ht -e v_open
```

- 9 Take all the VCS service groups offline:

```
# hagrps -offline grp_name -sys sys_name
```

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

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

- 11 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 12 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.
- 13 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.
- 14 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.

- 15 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.
- 16 The installer stops relevant processes, uninstalls old kernel files, and installs the new files. When prompted, enable replication or global cluster capabilities, if required, and register the software.

- 17 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 at this point will be offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically.

- 18 Relink the SF Oracle RAC libraries with Oracle on upgraded subcluster by using the `/opt/VRTS/install/installsfrac -configure` command and choosing option **Post Oracle Installation tasks**, then select option **Relink Oracle database Binary** from the program menu.
- 19 Mount all the Veritas File Systems which are not managed by VCS.
Bring the Oracle database service group online.
 - If the Oracle database is managed by VCS:

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

```
$ srvctl start instance -d db_name-I instance_name
```
- 20 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 21 Before you proceed to phase 2, complete step 4 to step 20 on the remaining subcluster.
- 22 Perform one of the following steps:
 - If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

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

```
$ srvctl modify database -d db-name -y AUTOMATIC
```

23 Migrate the SFDB repository database.

Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2

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

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

- 1 Start the installer for the rolling upgrade with the `-rollingupgrade_phase2` option. Specify all the nodes in the cluster:

```
./installmr -rollingupgrade_phase2 nodeA nodeB nodeC nodeD
```

- 2 The installer checks system communications, fileset versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3 Installer will start HA daemon (had) on all nodes, HA will be available once HA daemon is up.
- 4 Verify the cluster's status:

```
# hastatus -sum
```

- 5 If you want to upgrade CP server systems that use VCS or SFHA to 6.0.1, make sure that you upgraded all application clusters to version 6.0.1. 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.

Verifying software versions

To list the Veritas filesets installed on your system, enter the following command:

```
# ls1pp -L VRTS\*
```

The output version for 6.0.3 is 6.0.300.0.

Uninstalling version 6.0.3

This chapter includes the following topics:

- [About removing Veritas Storage Foundation and High Availability Solutions 6.0.3](#)
- [Rolling back using the `uninstallmr` script](#)
- [Rolling back manually](#)
-

About removing Veritas Storage Foundation and High Availability Solutions 6.0.3

This section describes how to roll back either by using the `uninstallmr` script or manually.

Rolling back using the `uninstallmr` script

Use the following procedure to roll back from any Veritas product to the previous version using the `uninstallmr` script.

To roll back

- 1 Browse to the directory that contains the `uninstallmr` script.
- 2 Stop all the processes and services accessing the file systems. For each disk group enter:

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

Verify that no volumes remain open.

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

- 3 Unmount all the Storage Checkpoints and the file systems.

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

Verify that you unmounted the the Storage Checkpoints and the file systems.

```
# mount | grep vxfs
```

- 4 Run the `uninstallmr` script to rollback patches, type:

```
# ./uninstallmr
```

- 5 The `uninstallmr` script checks whether the patches are at 6.0.1 committed level, and 6.0.3 applied level. If this is not the case, error messages showing the list of packages and commit levels will be shown.
- 6 The `uninstallmr` script removes 6.0.3 patches. After patch rollback completes, modules are loaded and processes are restarted. `uninstallmr` will also report any warning happened during uninstallation.

Rolling back manually

Use one of the following procedures to roll back to 6.0.1 manually.

- [Rolling back Storage Foundation or Storage Foundation and High Availability manually](#)
- [Rolling back Storage Foundation Cluster File System High Availability manually](#)
- [Rolling back Storage Foundation for Oracle RAC manually](#)
- [Rolling back Veritas Cluster Server manually](#)
- [Rolling back Dynamic Multi-Pathing manually](#)

Note: You must reboot systems that you roll back manually at the end of the roll back procedure.

Rolling back Storage Foundation or Storage Foundation and High Availability manually

Use the following procedure to roll back to 6.0.1 manually.

To roll back SF or SFHA

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 Unmount all Storage Checkpoints and file systems:

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

- 4 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 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 `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -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.
- 7 Stop all VxVM volumes by entering the following command for each disk group:

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

To verify that no volumes remain open, enter the following command:

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

8 Stop VCS and its modules manually.

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

9 Stop the AMF kernel driver:

```
# /etc/rc.d/rc2.d/S93amf stop
```

10 Stop I/O fencing on each node:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

11 Stop GAB:

```
# /etc/rc.d/rc2.d/S92gab stop
```

12 Stop LLT:

```
# /etc/rc.d/rc2.d/S701lt stop
```

13 Unmount `/dev/odm`:

```
# umount /dev/odm
```

14 Unload the ODM module:

```
# genkex | grep odm  
# vxkextadm vxodm unload
```

15 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

16 Remove the Storage Foundation or Storage Foundation and High Availability 6.0.3 patches.

- Create a file that contains all the 6.0.3 patches. In this example, it is called `/reject.list`.
- Reject each patch from the patch list file, for example:

```
# installp -rBf /reject.list
```

- 17 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

Rolling back Storage Foundation Cluster File System High Availability manually

Use the following procedure to roll back to 6.0.1 manually.

To roll back SFCFSA manually

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 Unmount all Storage Checkpoints and file systems:

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

- 4 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 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 `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -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 that are not under VCS control by entering the following command for each disk group:

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

To verify that no volumes remain open, enter the following command:

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

- 8 Stop VCS along with all the resources. Then, stop the remaining resources manually:

```
# /etc/rc.d/rc2.d/S99vcs stop
```

- 9 Unmount /dev/odm:

```
# umount /dev/odm
```

- 10 Unload the ODM module:

```
# genkex | grep odm  
# vxkextadm vxodm unload
```

- 11 Stop I/O fencing on each node:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

- 12 Stop GAB:

```
# /etc/rc.d/rc2.d/S92gab stop
```

- 13 Stop LLT:

```
# /etc/rc.d/rc2.d/S7011t stop
```

- 14 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
```

15 Remove the Storage Foundation Cluster File System 6.0.3 patches.

- Create a file that contains all the 6.0.3 patches. In this example, it is called `/reject.list`.
- Reject each patch from the patch list file, for example:

```
# installp -rBf /reject.list
```

16 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

Rolling back Storage Foundation for Oracle RAC manually

Use the following procedure to roll back to 6.0.1 manually.

To roll back SF for Oracle RAC manually**1** Stop Oracle and CRS on each node of the cluster.

- If Oracle Clusterware is controlled by VCS, log in as superuser on each system in the cluster and enter the following command:

```
# hastop -all
```

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

```
# crsctl stop crs
```

Unmount all VxFS file system used by a database or application and enter the following command to each node of the cluster:

```
# hastop -local
```

- 2 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped. In the `gabconfig -a` command output, the VCS engine or high availability daemon (HAD) port h is not displayed. This indicates that VCS has been stopped.

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

Sample output:

```
GAB Port Memberships
=====
Port a gen 5c3d0b membership 01
Port b gen 5c3d10 membership 01
Port d gen 5c3d0c membership 01
Port o gen 5c3d0f membership 01
```

3 Bring down the rest of the stack:

Stop vcsmm:

```
# /etc/rc.d/rc2.d/S98vcsmm stop
```

Stop lmx:

```
# /etc/rc.d/rc2.d/S71lmx stop
# /usr/lib/methods/lmxtxt -stop
```

Stop odm:

```
# /etc/rc.d/rc2.d/S99odm stop
```

Stop vxgms:

```
# /etc/methods/gmskextadm unload
```

Stop vxglm:

```
# /etc/methods/glmkextadm unload
```

Stop vxfen:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

Stop gab:

```
# /sbin/gabconfig -U
# /etc/methods/gabkext -stop
```

Stop llt:

```
# /sbin/lltconfig -U
```

4 Remove the Storage Foundation for Oracle RAC 6.0.3 patches.

- Create a file that contains all the 6.0.3 patches. In this example, it is called `/reject.list`:

You can use the following list as the reject list for Storage Foundation for Oracle components:

```
VRTSamf VRTScavf VRTSdbed
VRTSperl VRTSsfcp1601 VRTSvcS
VRTSvcSag VRTSvcsea
VRTSvxfen VRTSvxfs VRTSvxvm
```

- Reject each patch from the patch list file, for example:

```
# installp -rBf /reject.list
```

- 5 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 6.0.3 to VCS 6.0.1 on your cluster manually. To uninstall VCS, see the *Veritas Cluster Server Installation Guide*.

Note: Use this procedure only when rolling back VCS. Do not roll back VCS when it is part of other products that rely on VCS, for example Storage Foundation Clustered File System or Storage Foundation for Oracle RAC.

To roll back 6.0.3:

- 1 Verify that all of the VCS 6.0.3 patches are in the APPLIED state. Create a text file called `filesets.to.reject` that contains the name and version of each fileset, one per line, exactly as shown below.

```
VRTSamf                6.0.300.0
VRTSvcS                 6.0.300.0
VRTSvcSag               6.0.300.0
VRTSvcsea               6.0.300.0
VRTSvcxfen              6.0.300.0
```

- 2 On each node, make a local copy of `filesets.to.reject` and then type:

```
# nohdr='^Z$'
# while read pkg ver; do
  lslpp -l $pkg | egrep -v "$nohdr"
  nohdr='^ Fileset +Level State '
done < filesets.to.reject
```

Note: Any updates that are in COMMITTED state cannot be rejected (undone). You must remove each one and then re-install it.

- 3 List the service groups in your cluster and their status. On any node, type:

```
# hagrps -state
```

- 4 Take the ClusterService service group offline if it is running. On any node, type:

```
# hagrps -offline -force ClusterService -any
```

- 5 Make the VCS configuration writable. On any node, type:

```
# haconf -makerw
```

- 6 Freeze all service groups except the ClusterService service group. On any node, type:

```
# hagrps -list | sort -u +0b -1 | \
    while read grp sys ; do
        hagrps -freeze $grp -persistent
    done
```

You can safely ignore the warning about the failure to freeze the ClusterService group.

- 7 Save the configuration (main.cf) file with the groups frozen. On any node, type:

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

- 8 Make a backup copy of the current main.cf and all types.cf configuration files. For example, on one node in the cluster, type:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
    /etc/VRTSvcs/conf/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
    /etc/VRTSvcs/conf/types.cf.save
```

- 9 Shut down VCS. On any node, type:

```
# /opt/VRTSvcs/bin/hastop -all -force
```

- 10 Shut down CmdServer. On each node, type:

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

- 11 Verify that VCS has shut down.

- On any node, type:

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

The output resembles:

GAB Port Memberships

```
Port a gen 23dc0001 membership 01
```

Output for membership for port h does not appear.

- On each node, run the command:

```
# ps -ef | egrep "had|hashadow|CmdServer"
```

Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.

12 Stop AMF, fencing, GAB, and LLT.

```
# /etc/rc.d/rc2.d/S93amf stop
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/methods/vxfenext -stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/methods/gabkext -stop
# /etc/rc.d/rc2.d/S70llt stop
```

13 Preview the patch removal selection and validity tests. On each node, type:

```
# installp -pr -gXv -f filesets.to.reject
```

Confirm that the patches to be removed are exactly the same as those listed in the filesets.to.reject file that you created in [step 1](#).

14 Perform the patch removal. On each node, type:

```
# installp -r -gXv -f filesets.to.reject
```

Review the summaries at the end of each run and confirm that all of the intended patches removed successfully.

15 Reboot all nodes in the cluster.**16 After VCS has started, perform the following steps:**

- Verify all resources have been probed. On any node, type:

```
# hastatus -summary
```

- Unfreeze all service groups. On any node, type:

```
# haconf -makerw
# hagr -list | sort -u +0b -1 | \
  while read grp sys ; do
    hagr -unfreeze $grp -persistent
```

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

You can safely ignore the warning about the failure to unfreeze the ClusterService group.

- 17 Bring the ClusterService service group online, if necessary. On any node, type:

```
# hagrps -online ClusterService -sys system
```

where system is the node name.

Rolling back Dynamic Multi-Pathing manually

Use the following procedure to roll back to 6.0.1 manually.

To roll back DMP manually

- 1 Stop activity to all VxVM volumes.
- 2 Stop all VxVM volumes that are not under VCS control by entering the following command for each disk group:

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

To verify that no volumes remain open, enter the following command:

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

- 3 Perform the following commands to determine whether root support or DMP native support is enabled.

```
■ # vxdmpadm gettune dmp_native_support
```

If the command returns an "on" value, DMP native support is enabled on the system. If the command returns any other value, DMP native support is disabled.

```
■ # vxdmpadm native list vgname=rootvg
```

If the output is a list of hdisks, root support is enabled on this system. If the command returns any other value, root support is disabled.

- Once you have determined if root support or DMP native support is enabled, go to step 4.
- Once you have determined that root support and DMP native support is not enabled, go to step 5.

- 4 If root support or DMP native support is enabled:

- You must disable DMP native support.
Run the following command to disable DMP native support and to disable root support:

```
# vxddmpadm settune dmp_native_support=off
```

- If only root support is enabled, run the following command to disable root support:

```
# vxddmpadm native disable vgname=rootvg
```

- Reboot the system:

```
# shutdown -r now
```

- Before backing out patch, stop the VEA server's vxsvc process:

```
# /opt/VRTSob/bin/vxsvcctl stop
```

- Create a file that contains all the 6.0.3 patches. In this example, it is called `/reject.list`:

```
# /reject.list
```

- Reject each patch from the patch list file, for example:

```
# installp -rBf /reject.list
```

- Reboot the system:

```
# shutdown -r now
```

5 If root support or DMP native support is not enabled:

- Before you back out the patch, kill the VEA Server's vxsvc process:

```
# /opt/VRTSob/bin/vxsvcctl stop
```

- To reject the patch if it is in `APPLIED` state

```
# installp -r patch_name
```

- Reboot the system:

```
# shutdown -r now
```

- 6 Enable DMP native support (this also enables root support) if it was enabled before applying the patch:

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

- Reboot the system:

```
# shutdown -r now
```

- Verify DMP native or root support is enabled:

```
# vxddmpadm gettune dmp_native_support
```

- 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/hagrps -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
```