

Veritas Storage Foundation™ and High Availability Solutions 6.0.3 Release Notes - HP-UX 11i v3

6.0.3 Maintenance Release

Veritas Storage Foundation™ and High Availability Solutions Release Notes

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

Product version: 6.0.3

Document version: 6.0.3 Rev 6

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apj@symantec.com

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

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

Documentation

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

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

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Contents

Technical Support	3
Chapter 1	
About this release	9
Introduction	9
Changes in this release	10
Changes in Veritas Storage Foundation High Availability	10
Changes in Veritas Cluster Server	10
System requirements	11
List of patches	11
Veritas Storage Foundation patches in 6.0.3	12
Veritas Cluster Server patches in 6.0.3	12
Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) patches in 6.0.3	13
Veritas Storage Foundation for Oracle RAC patches in 6.0.3	14
Fixed issues in this release	14
Installation and upgrade: Issues fixed in 6.0.3	14
Veritas Storage Foundation 6.0.3 fixed issues	15
Storage Foundation for Databases (SFDB) tools: Issues fixed in 6.0.3	20
Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues	20
Veritas Cluster Server 6.0.3 fixed issues	21
Veritas Storage Foundation for Oracle RAC 6.0.3 fixed issues	22
Software limitations in this release	22
Veritas Storage Foundation 6.0.3 software limitations	23
Veritas Storage Foundation 6.0.1 software limitations	23
Veritas Storage Foundation for Databases tools 6.0.3 software limitations	26
Veritas Storage Foundation for Databases (SFDB) tools 6.0.1 software limitations	26
Veritas Storage Foundation Cluster File System High Availability 6.0.3 software limitations	26
Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) 6.0.1 software limitations	27

	Veritas Cluster Server 6.0.3 software limitations	27
	Veritas Cluster Server 6.0.1 software limitations	27
	Veritas Storage Foundation for Oracle RAC 6.0.3 software limitations	31
	Known issues in this release	33
	Issues related to installation 6.0.3	33
	Issues related to any OS or supported technology	38
	Veritas Storage Foundation 6.0.3 known issues	38
	Veritas Storage Foundation 6.0.1 known issues	39
	Veritas Storage Foundation for Databases (SFDB) tools 6.0.3 known issues	55
	Veritas Storage Foundation for Databases (SFDB) tools 6.0.1 known issues	55
	Veritas Storage Foundation Cluster File System High Availability 6.0.3 known issues	61
	Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) 6.0.1 known issues	61
	Veritas Cluster Server 6.0.3 known issues	64
	Veritas Cluster Server 6.0.1 known issues	65
	Veritas Storage Foundation for Oracle RAC known issues	90
	Downloading the patches	93
Chapter 2	Upgrading to version 6.0.3	94
	About the installmr script	94
	About using the installer to upgrade when the root disk is encapsulated	98
	Performing a full upgrade to 6.0.3 on a cluster	98
	Performing a full upgrade to 6.0.3 for Veritas Cluster Server	99
	Performing a full upgrade to 6.0.3 on an SFHA cluster	100
	Performing a full upgrade to 6.0.3 on an SFCFSHA cluster	102
	Performing a full upgrade to version 6.0.3 on an SF Oracle RAC cluster	104
	Performing a full upgrade to 6.0.3 on a standalone system	110
	Performing a rolling upgrade to 6.0.3 on a cluster	112
	About rolling upgrades	112
	Prerequisites for rolling upgrades	112
	Performing a rolling upgrade using the script-based installer	113
Chapter 3	Uninstalling version 6.0.3	118
	About uninstalling Veritas Storage Foundation and High Availability Solutions 6.0.3	118
	About the uninstallmr script	119

Rolling back using the uninstallmr script	121
Uninstalling 6.0.3 with the Web-based installer	123

About this release

This chapter includes the following topics:

- [Introduction](#)
- [Changes in this release](#)
- [System requirements](#)
- [List of patches](#)
- [Fixed issues in this release](#)
- [Software limitations in this release](#)
- [Known issues in this release](#)
- [Downloading the patches](#)

Introduction

This document provides information about the products in Veritas Storage Foundation and High Availability Solutions 6.0.3 Maintenance Release (6.0.3 MR). Symantec strongly recommends installing the 6.0.3 immediately after installing Veritas Storage Foundation and High Availability Solutions 6.0.1.

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

<http://www.symantec.com/docs/TECH202209>

Review this entire document before installing and upgrading your Veritas Storage Foundation and High Availability product.

For further details, depending on the product for which you want to install this patch, refer to one of the following release notes:

- *Veritas Storage Foundation Release Notes (Version 6.0.1)*
- *Veritas Cluster Server Release Notes (Version 6.0.1)*
- *Veritas Storage Foundation Cluster File System High Availability Release Notes (Version 6.0.1)*
- *Veritas Dynamic Multi-Pathing Release Notes (Version 6.0.1)*
- *Veritas Storage Foundation for Oracle RAC Release Notes (Version 6.0.1)*

Apply this patch for the following Veritas Storage Foundation and High Availability Solutions products:

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation High Availability (SFHA)
- Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)
- Veritas Cluster Server (VCS)
- Veritas Dynamic Multi-Pathing (DMP)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)

Changes in this release

This section lists the changes introduced in Veritas Storage Foundation and High Availability Solutions 6.0.3 MR release.

- Changes in Veritas Storage Foundation High Availability
See [“Changes in Veritas Storage Foundation High Availability”](#) on page 10.
- Changes in Veritas Cluster Server
See [“Changes in Veritas Cluster Server”](#) on page 10.

Changes in Veritas Storage Foundation High Availability

This release supports HP Integrity Virtual Machines (IVM) 6.1 and 6.2 and vPars 6.2.

Note: SFRAC supports HP vPar(static) only and does not support HP vPar(dynamic) and HP IVM. SFRAC supports HP-UX vpar version 6.1 in this release.

Changes in Veritas Cluster Server

This release supports HP Integrity Virtual Machines (IVM) 6.1 and 6.2 and vPars 6.2.

Db2udb Agent support extended to DB2 10.1

The DB2udb Agent for VCS 6.0.3 now supports DB2 10.1.

Note: SFDB tools are also supported with HP Integrity Virtual Machines (IVM) and vPars 6.2.

System requirements

For information on system requirements, refer to the product documentation for Veritas Storage Foundation and High Availability Solutions 6.0.1.

Note: This release requires that Version 6.0.1 is installed on your systems.

The following table lists the supported operating systems for this release of Veritas Storage Foundation and High Availability Solutions.

Table 1-1 supported operating systems

Operating system	Operating system version	Architecture
HP-UX 11i Version 3 March 2011 Operating Environments Update Release or later	HP-UX B.11.31.1103 HP-UX B.11.31.1109 HP-UX B.11.31.1203 HP-UX B.11.31.1209 HP-UX B.11.31.1303	PA-RISC Itanium

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

For the latest information on supported hardware, visit:

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

Symantec recommends installing the latest HP-UX patches from HP.

List of patches

This section lists the patches included in this release.

- Veritas Storage Foundation
 - See [“Veritas Storage Foundation patches in 6.0.3”](#) on page 12.

- Veritas Cluster Server
 See [“Veritas Cluster Server patches in 6.0.3”](#) on page 12.
- Veritas Storage Foundation Cluster File System
 See [“Veritas Storage Foundation Cluster File System High Availability \(SFCFSHA\) patches in 6.0.3”](#) on page 13.
- Veritas Storage Foundation for Oracle RAC
 See [“Veritas Storage Foundation for Oracle RAC patches in 6.0.3”](#) on page 14.

Veritas Storage Foundation patches in 6.0.3

[Table 1-2](#) lists the Veritas Storage Foundation patches included in this release.

Table 1-2 Veritas Storage Foundation patches

Patch	Version	Description
PVCO_03972	1.0	VRTS 6.0.3 VRTSvxfs Command Patch (Veritas File System)
PVKL_03971	1.0	VRTS 6.0.3 VRTSvxfs Kernel Patch (Veritas File System)
PVCO_03974	1.0	VRTS 6.0.3 VRTSvxvm Command Patch (Veritas Volume Manager)
PVKL_03975	1.0	VRTS 6.0.3 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PVCO_03981	1.0	VRTS 6.0.3 VRTSdbed Command Patch (Veritas Storage Foundation for Databases Tools)
PVCO_03970	1.0	VRTS 6.0.300.000 VRTSsfcp601 Command Patch
PVCO_03982	1.0	VRTS 6.0.300.000 VRTSperl Command Patch (Perl Redistribution)

Veritas Cluster Server patches in 6.0.3

[Table 1-3](#) lists the Veritas Cluster Server patches included in this release.

Table 1-3 Veritas Cluster Server patches

Patch	Version	Description
PVCO_03976	1.0	VRTS 6.0.300.000 VRTSvcs Command Patch
PVCO_03977	1.0	VRTS 6.0.300.000 VRTSvcsag Command Patch

Table 1-3 Veritas Cluster Server patches (*continued*)

Patch	Version	Description
PVCO_03978	1.0	VRTS 6.0.300.000 VRTSvcsea Command Patch
PVKL_03979	1.0	VRTS 6.0.300.000 VRTSvxfen Kernel Patch
PVKL_03980	1.0	VRTS 6.0.300.000 VRTSamf Kernel Patch
PVCO_03970	1.0	VRTS 6.0.300.000 VRTSsfcp1601 Command Patch
PVCO_03982	1.0	VRTS 6.0.300.000 VRTSperl Command Patch (Perl Redistribution)

Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) patches in 6.0.3

[Table 1-4](#) lists the Veritas Storage Foundation Cluster File System patches included in this release.

Table 1-4 Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) patches

Patch	Version	Description
PVCO_03973	1.0	VRTS 6.0.3 VRTScavf Command Patch
PVCO_03981	1.0	VRTS 6.0.3 VRTSdbed Command Patch (Veritas Storage Foundation for Databases Tools)
PVCO_03978	1.0	VRTS 6.0.300.000 VRTSvcsea Command Patch
PVCO_03977	1.0	VRTS 6.0.300.000 VRTSvcsg Command Patch
PVCO_03976	1.0	VRTS 6.0.300.000 VRTSvcsc Command Patch
PVKL_03980	1.0	VRTS 6.0.300.000 VRTSamf Kernel Patch
PVKL_03979	1.0	VRTS 6.0.300.000 VRTSvxfen Kernel Patch
PVKL_03971	1.0	VRTS6.0.3 VRTSvxfs Kernel Patch (Veritas File System)
PVCO_03972	1.0	VRTS 6.0.3 VRTSvxfs Command Patch (Veritas File System)
PVCO_03974	1.0	VRTS 6.0.3 VRTSvxvm Command Patch (Veritas Volume Manager)

Table 1-4 Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) patches (*continued*)

Patch	Version	Description
PVKL_03975	1.0	VRTS 6.0.3 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PVCO_03970	1.0	VRTS 6.0.300.000 VRTSsfpci601 Command Patch
PVCO_03982	1.0	VRTS 6.0.300.000 VRTSperl Command Patch (Perl Redistribution)

Veritas Storage Foundation for Oracle RAC patches in 6.0.3

There are no updates to the SF Oracle RAC depots.

For the other patches included in this release:

See [“Veritas Storage Foundation Cluster File System High Availability \(SFCFSHA\) patches in 6.0.3”](#) on page 13.

Fixed issues in this release

This section describes issues fixed in this release.

- Veritas Storage Foundation
 See [“Veritas Storage Foundation 6.0.3 fixed issues”](#) on page 15.
- Veritas Storage Foundation for Databases (SFDB) tools
 See [“Storage Foundation for Databases \(SFDB\) tools: Issues fixed in 6.0.3”](#) on page 20.
- Veritas Cluster Server
 See [“Veritas Cluster Server 6.0.3 fixed issues”](#) on page 21.
- Veritas Storage Foundation Cluster File System
 See [“Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues”](#) on page 20.
- Veritas Storage Foundation for Oracle RAC
 See [“Veritas Storage Foundation for Oracle RAC 6.0.3 fixed issues”](#) on page 22.

Installation and upgrade: Issues fixed in 6.0.3

[Table 1-5](#) lists the install and upgrade issues fixed in this release.

Table 1-5 Install and upgrade 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.
2851403	The <code>vxportal</code> and <code>vxfs</code> processes are failed to stop during first phase of rolling upgrade.

Veritas Storage Foundation 6.0.3 fixed issues

[Table 1-6](#) lists the Veritas Volume Manager issues fixed in this release.

Table 1-6 Veritas Volume Manager fixed issues

Incident	Description
2858853	After master switch, <code>vxconfigd</code> dumps core on old master.
2779580	Secondary node gives configuration error 'no Primary RVG' after reboot of master node on primary site.
2866059	Improve error messages that are hit during <code>vxdisk</code> resize operation.
2149922	Record the diskgroup import and deport events in <code>syslog</code> .
2851403	System panic is seen while unloading "vxio" module. This happens whenever VxVM uses SmartMove feature and the "vxportal" module gets reloaded.
2930569	The LUNs in 'error' state in output of ' <code>vxdisk list</code> ' cannot be removed through DR (Dynamic Reconfiguration) Tool.
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.
1859018	Link <link-name> link detached from volume <volume-name> warnings are displayed when a linked-breakoff snapshot is created.
2715129	<code>Vxconfigd</code> hangs during Master takeover in a CVM (Clustered Volume Manager) environment.

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

Incident	Description
2753954	When a cable is disconnected from one port of a dual-port FC HBA, the paths via another port are marked as SUSPECT PATH.
2886402	When re-configuring devices, <code>vxconfigd</code> hang is observed.
2878876	<code>vxconfigd</code> dumps core in <code>vol_cbr_dolog()</code> due to race between two threads processing requests from the same client.
1973983	<code>vxunreloc</code> fails when DCO (Data Change Object) plex is in disabled state.
2801962	Growing a volume takes significantly large time when the volume has version 20 DCO (Data Change Object) attached to it.
2886333	<code>vx dg (1M) join</code> command allowed mixing clone and non-clone disk group. Subsequent import of new joined disk group fails.
2000585	<code>vxrecover</code> does not start remaining volumes, if, one of the volumes is removed during <code>vxrecover</code> command run.
2834046	NFS migration failed due to device reminoring.
2567618	VRTSexplorer core dumps in <code>checkhbaapi/print_target_map_entry</code> .
1765916	VxVM socket files do not have proper write protection.
1903700	Removing mirror using <code>vxassist</code> does not work.
2919627	Dynamic Reconfiguration tool should be enhanced to remove LUNs feasibly in bulk.
2892983	<code>vxvol</code> dumps core if new links are added while the operation is in progress.
1982965	<code>vx dg import DGNAME <da-name..></code> fails when "da-name" used as an input to <code>vx dg</code> command is based on naming scheme which is different from the prevailing naming scheme on the host.
2899173	<code>vxconfigd</code> hangs after executing <code>vradmin stoprep</code> command.
2910043	Frequent swapin/swapout of pages seen due to higher order memory requests.
2876256	<code>vx disk -f -g <dg1> set <da_name> mediatype=ssd</code> command fails with new naming scheme.

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

Incident	Description
2859470	EMC SRDF (Symmetrix Remote Data Facility) R2 disk with EFI label is not recognized by VxVM (Veritas Volume Manager) and its shown in error state.
2836798	In VxVM, resizing simple EFI disk fails and causes system panic/hang.
2919720	<code>vxconfigd</code> dumps core in <code>rec_lock1_5()</code> function.
2940446	I/O can hang on volume with space optimized snapshot if the underlying cache object is of very large size. It can also lead to data corruption in cache-object.
1725593	The <code>vxdumpadm listctlr</code> command has to be enhanced to print the count of device paths seen through the controller.
2970368	Enhancing handling of SRDF-R2 WD devices in DMP.
2510928	The extended attributes reported by <code>vxdisk -e list</code> for the EMC SRDF LUNs are reported as <code>tdev mirror</code> , instead of <code>tdev srdf-r1</code> .
2942609	Message displayed when user quits from Dynamic Reconfiguration Operations is shown as error message.
2911040	Restore from a cascaded snapshot leaves the volume in unusable state if any cascaded snapshot is in detached state.
2744004	<code>vxconfigd</code> hangs on the VVR (Veritas Volume Replicator) secondary node during VVR (Veritas Volume Replicator) configuration.
2919318	During CVM (Cluster Volume Manager) node join and shared disk group operation, the I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing.
2833498	<code>vxconfigd</code> daemon hangs in <code>vol_ktrans_commit()</code> while reclaim operation is in progress on volumes having instant snapshots.
2857827	During early boot, recovery of linked volume resize fails due to <code>/usr</code> not mounted.
2815517	<code>vxvg adddisk</code> allows mixing of clone and non-clone disks in a DiskGroup.
2915063	System panics during detaching plex of volume in CVM (Cluster Volume Manager) environment.
2837717	<code>vxdisk(1M) resize</code> command fails if 'da name' is specified.

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

Incident	Description
2826125	VxVM (Veritas Volume Manager) script daemon is terminated abnormally on its invocation.
2919714	On a THIN LUN, <code>vxevac</code> returns 0 without migrating unmounted VxFS (Veritas FileSystem) volumes.
2692012	When moving subdisks, using <code>vxevac</code> command fails with a generic message which does not convey exactly why the operation failed.
2851085	DMP does not detect implicit LUN ownership changes for some of the DMP nodes.
2898547	<code>vradmind</code> dumps core on Veritas Volume Replicator secondary Site , when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2798673	System panics in <code>voldco_alloc_layout()</code> function while creating volume with instant DCO (Dynamic Change Object).
2933138	System panics in <code>voldco_update_itemq_chunk()</code> function due to accessing invalid buffer.
2619600	Live migration of virtual machine having SFHA/SFCFSHA stack with data disks fencing enabled, causes service groups configured on virtual machine to fault.
2149922	Record the diskgroup import and deport events in syslog.
1901838	Incorrect setting of "No license" flag leads to DMP (Dynamic Multi -Pathing) database inconsistency.
2273190	Incorrect setting of UNDISCOVERED flag can lead to database inconsistency.
2898547	<code>vradmind</code> on VVR (Veritas Volume Replicator) Secondary Site dumps core, when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2962262	Uninstallation of DMP (Dynamic Multi - Pathing) fails in presence of other multipathing solutions.
2851085	DMP (Dynamic Multi - Pathing) does not detect implicit LUN ownership changes for some of the DMP nodes.

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

Incident	Description
2948172	Execution of <code>vxdisk -o thin, fssize list</code> command can result in panic.
2935771	In VVR (Veritas Volume Replicator) environment, RLINK diconnects after switching a master on the primary.
3002770	Accessing NULL pointer in <code>dmp_aa_recv_inquiry()</code> causes system panic.
2869594	Master node panics due to corruption, if, space optimized snapshots are refreshed and master node is selected using <code>vxclustadm setmaster</code> .
2965910	When <code>-o ordered</code> is used, <code>vxassist</code> handles non-disk parameters in a differenn way. This scenario may result in invalid comparison, leading to a core dump.
2398416	<code>vxassist</code> dumps core while creating volume when attribute <code>wantmirror=ctrlr</code> is added to the <code>/etc/default/vxassist</code> file.
2851403	System panic is seen while unloading <code>vxio</code> module. This happens whenever VxVM uses SmartMove feature and the <code>vxportal</code> module gets reloaded (For example, during VxFS package upgrade).

Table 1-7 lists the Veritas File System issues fixed in this release.

Table 1-7 Veritas File System fixed issues

Incident	Description
2895743	Accessing named attributes for some files seems to be slow.
2881211	File ACLs not preserved in checkpoints properly if file has hardlink.
2756779	Write and read performance concerns on CFS when running applications that rely on posix file-record locking (fcntl).
2858683	Reserve extent attributes changed after <code>vxrestore</code> , only for files greater than 8192 bytes.
2806466	A reclaim operation on a file system mounted on a Logical Volume Manager (LVM) volume using the <code>fsadm(LM)</code> command with the 'R' option may panic the system.
2624262	Panic hit in <code>vx_bc_do_brelse()</code> function while executing dedup functionality.

Table 1-7 Veritas File System fixed issues (*continued*)

Incident	Description
2616622	The performance of the <code>mmap()</code> function is slow when the file system block size is 8 KB and the page size is 4 KB.
2857751	The internal testing hits the assert <code>f:vx_cbdnlc_enter:1a</code> .
2730759	The sequential read performance is poor because of the read-ahead issues.
2850730	LM conformance hits an assert <code>f:vx_do_getpage:6b,3</code> and panics.
2417858	VxFS quotas do not support 64 bit limits.
2857629	File system corruption can occur requiring a full <code>fsck</code> of the system.
2590918	Delay in freeing unshared extents upon primary switch over.
2885592	<code>vxdump</code> to the <code>vxcompress</code> file system is aborted.

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

[Table 1-8](#) lists the Veritas Storage Foundation for Databases (SFDB) tools fixed issues in this release.

Table 1-8 Veritas Storage Foundation for Databases (SFDB) tools fixed issues

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

Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues

[Table 1-9](#) lists the Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) issues fixed in this release.

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

Incident	Description
2715175	<code>cfsumount</code> command runs slowly on large file system and for large file systems, reconfiguration process takes a longer time.

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

Incident	Description
2750860	Performance write issue observed due to CFS (Cluster File System) fragmentation in CFS (Cluster File System) cluster.
2942776	CFS mount fails with the error ENXIO or EIO on volume vset device when the volumes in vset is not ready.
2923105	The upgrade <code>VRTSvxfs5.0MP4HFaf</code> hangs at vxfs (Veritas File System) preinstall scripts.
2923867	Internal test hits an assert <code>f:xted_set_msg_pri:1</code> .
2841059	The file system gets marked for a full fsck operation and the attribute inode is marked as 'bad ondisk'.
2916691	Customer experiencing hangs when doing dedups.
2977697	<code>vx_idetach</code> generated kernel core dump while running filestore replication.
2906018	<code>vx_ireaderrors</code> after successful log replay and mount of the file system.

Veritas Cluster Server 6.0.3 fixed issues

[Table 1-10](#) lists the Veritas Cluster Server issues fixed in this release.

Table 1-10 Veritas Cluster Server fixed issues

Incident	Description
2736627	The remote cluster remains in "INIT" state and the ICMP heartbeat status is "UNKNOWN".
2737653	If you override the OnlineTimeout attribute value for the RVGPrimary resource, the agent does not consider it.
2848009	Asynchronous Monitoring Framework (AMF) panics the system when an agent exits.
2861253	In the vxfs driver debug log message, the jeopardy membership status is printed as garbage.
2937673	The AMF driver panics the system when the amfstat utility is executed.

Table 1-10 Veritas Cluster Server fixed issues (*continued*)

Incident	Description
2941155	Veritas Cluster Server (VCS) does not report the group as offline on a failed cluster when the cluster failure is declared in the global cluster configuration.
2964772	If the NFSRestart resource is taken offline, the NFSRestart agent may unexpectedly stop the NFS processes in a local container.
3013940	<p>In non-MPP mode, when no virtual host is configured in the db2nodes.cfg file, the db2start command fails with the following message:</p> <pre>A communication error occurred during START or STOP DATABASE MANAGER processing. SQL1032N No start database manager command was issued. SQLSTATE=57019.</pre>
3013962	In DB2 10.1, the monitor fails to detect the online DB2 instance.

Veritas Storage Foundation for Oracle RAC 6.0.3 fixed issues

There are no issues fixed in SF Oracle RAC 6.0.3.

Software limitations in this release

This section describes the software limitations in this release.

- Veritas Storage Foundation
 See [“Veritas Storage Foundation 6.0.3 software limitations”](#) on page 23.
- Veritas Storage Foundation for Databases (SFDB) Tools
 See [“Veritas Storage Foundation for Databases tools 6.0.3 software limitations”](#) on page 26.
- Veritas Cluster Server
 See [“Veritas Cluster Server 6.0.3 software limitations”](#) on page 27.
- Veritas Storage Foundation Cluster File System High Availability
 See [“Veritas Storage Foundation Cluster File System High Availability 6.0.3 software limitations”](#) on page 26.
- Veritas Storage Foundation for Oracle RAC
 See [“Veritas Storage Foundation for Oracle RAC 6.0.3 software limitations”](#) on page 31.

Veritas Storage Foundation 6.0.3 software limitations

There are no software limitations in this release.

The software limitations for previous releases are as follows:

6.0.1 See [“Veritas Storage Foundation 6.0.1 software limitations”](#) on page 23.

Veritas Storage Foundation 6.0.1 software limitations

This section lists Veritas Storage Foundation software limitations for 6.0.1 release.

Dynamic Multi-Pathing 6.0.1 software limitations

This section lists Dynamic Multi-Pathing software limitations for 6.0.1 release.

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 Volume Manager (VxVM) 6.0.1 software limitations

This section lists VxVM software limitations for this release.

SF Oracle RAC 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.

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.

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    thinclm
xiv0_613    19313     2108            dg1    thinclm
xiv0_614    19313     35              dg1    thinclm
xiv0_615    19313     32              dg1    thinclm
xiv0_616    19313     31              dg1    thinclm
xiv0_617    19313     31              dg1    thinclm
xiv0_618    19313     31              dg1    thinclm
```

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 `lfailed`, `lmissing` or `LDISABLED` are introduced when I/O shipping is active because of storage disconnectivity.

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.

Veritas FileSystem 6.0.1 software limitations

This section lists Veritas FileSystem software limitations for this release.

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.

After uninstalling Veritas File System 6.0, a file system with disk layout Version 7 or later cannot be mounted

If you install Veritas File System (VxFS) 5.0 or later, create a file system with disk layout Version 7 or later, and then uninstall VxFS, you are left with the base VxFS release of 4.1. VxFS 4.1 does not recognize disk layout Version 7 or later, and thus you cannot mount the file system.

Workaround: You must reinstall VxFS 5.0 or later to mount a file system that has disk layout Version 7, VxFS 5.1 SP1 or later to mount a file system that has disk layout Version 8, or VxFS 6.0 to mount a file system that has disk layout Version 9.

Data deduplication is not supported on PA architecture

The data deduplication feature is not supported on PA architecture.

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 tools 6.0.3 software limitations

There are no software limitations in this release.

The software limitations for previous releases are as follows:

6.0.1

See "[Veritas Storage Foundation for Databases \(SFDB\) tools 6.0.1 software limitations](#)" on page 26.

Veritas Storage Foundation for Databases (SFDB) tools 6.0.1 software limitations

This section lists Veritas Storage Foundation for Databases (SFDB) tools software limitations for this release.

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

Only one instance of the `vxsfdm` command can be run at a time. Running multiple instances of `vxsfdm` 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 Storage Foundation Cluster File System High Availability 6.0.3 software limitations

There are no software limitations for this release.

The software limitations for the previous releases are as follows:

6.0.1

See ["Veritas Storage Foundation Cluster File System High Availability \(SFCFSHA\) 6.0.1 software limitations"](#) on page 27.

Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) 6.0.1 software limitations

This section lists Veritas Storage Foundation Cluster File System High Availability software limitations for this release.

VMware vSphere extension for VirtualStore limitations

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

F5 usage is not supported for wizard refreshing (2362940)

F5 usage is not supported for wizard refreshing.

Workaround

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

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

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

Veritas Cluster Server 6.0.3 software limitations

There are no software limitations in this release.

For VCS 6.0.1 software limitations, see [Veritas Cluster Server 6.0.1 software limitations](#).

Veritas Cluster Server 6.0.1 software limitations

This section describes the software limitations in the VCS 6.0.1 release.

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.

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/VRTSvc/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/VRTSvc`:

```
# hatype -modify <ResourceType> AgentFile
/opt/VRTSvc/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

Systems in a cluster must have same system locale setting

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

Limitations with DiskGroupSnap agent [1919329]

The DiskGroupSnap agent has the following limitations:

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

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

Cluster Manager (Java console) limitations

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

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

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

See the *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.

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 package, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm package is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

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 for Oracle RAC 6.0.3 software limitations

There are no software limitations in this release.

The software limitations from previous releases are as follows:

- 6.0.1 See "[Veritas Storage Foundation for Oracle RAC 6.0.1 software limitations](#)" on page 31.

Veritas Storage Foundation for Oracle RAC 6.0.1 software limitations

The software limitations in SF Oracle RAC 6.0.1 are as follows.

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.

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.

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

When all nodes present in the SF Oracle RAC 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 `vxsfenclearpre` utility.

For more information on the `vxfenclearpre` utility, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

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.

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 Oracle RAC environments

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

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.

Known issues in this release

This section describes the known issues in this release.

- Issues related to installation
See [“Issues related to installation 6.0.3”](#) on page 33.
- Veritas Storage Foundation
See [“Veritas Storage Foundation 6.0.3 known issues”](#) on page 38.
- Veritas Storage Foundation for Databases (SFDB) Tools
See [“Veritas Storage Foundation for Databases \(SFDB\) tools 6.0.3 known issues”](#) on page 55.
- Veritas Cluster Server
See [“Veritas Cluster Server 6.0.3 known issues”](#) on page 64.
- Veritas Storage Foundation Cluster File System High Availability
See [“Veritas Storage Foundation Cluster File System High Availability 6.0.3 known issues”](#) on page 61.
- Veritas Storage Foundation for Oracle RAC
See [“Veritas Storage Foundation for Oracle RAC known issues”](#) on page 90.

Issues related to installation 6.0.3

This section lists issues related to installation in this release.

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

```
# hagr -list Frozen=1
```

- 2 Unfreeze all the frozen service groups:

```
# haconf -makerw  
# hagr -unfreeze service_group -persistent  
# haconf -dump -makero
```

Warning messages may be seen during script-based installation (2615447)

When you install SF Oracle RAC using the script-based installer, you may see the following warning message:

```
interpreter "/opt/VRTSperl/bin/perl" not found
```

Workaround: You must install perl to resolve the issue.

To install perl

- 1 Exit the installer.
- 2 Install the VRTSperl package from the product media manually:

```
# cd /dvd_path/depot  
# /usr/sbin/swinstall -x enforce_dependencies=false  
-x autoreboot=false -s `pwd` VRTSperl
```

- 3 Start the installer again.

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

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

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the `/usr/opensv/netbackup/bin/version` file and `/usr/opensv/netbackup/version` file to the `/tmp` directory. If you have clustered

NetBackup installed, you must also copy the

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

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

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

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

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

The VRTSacclib package is deprecated (2032052)

The VRTSacclib package is deprecated. For installation, uninstallation, and upgrades, note the following:

- Fresh installs: Do not install VRTSacclib.
- Upgrade: Ignore VRTSacclib.
- Uninstall: Ignore VRTSacclib.

Errors recorded in the swremove logs of VRTSgab during VCS upgrade from 4.1 to 5.0.1 (1719136)

When VCS is upgraded from 4.1 to 5.0.1 on HP-UX 11i v3 using the Veritas product installer, the installer reports errors for GAB and errors are recorded in the swremove logs related to VRTSgab.

Workaround: You can safely ignore these error messages.

VCS agents dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the update-ux command (1630968)

On PA-RISC architecture, the VCS agents (Oracle, Netlsnr, Sybase, SybaseBk, MultiNICB, and so on) may dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the `update-ux` command.

This is because on HP-UX PA-RISC systems, the default thread stack size is limited to 64k. When the agent requires more than 64k stack memory, it may dump core due to SIGBUS error.

Workaround: Before running the `update-ux` command, edit the `/opt/VRTSvcs/bin/vcsenv` file to append following lines to it:

```
PLATFORM=`uname -s`  
ARCHITECTURE=`uname -m`  
if [ "${PLATFORM}" = "HP-UX" ] && [ "${ARCHITECTURE}" = "9000/800" ];  
then  
    PTHREAD_DEFAULT_STACK_SIZE=524288  
    export PTHREAD_DEFAULT_STACK_SIZE  
fi
```

After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon after 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: Refer to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for the section, "vxconfigd daemon recovery."

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

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

If the CVM protocol version does 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
```

Web installer does not ask for authentication after the first session if the browser is still open (2509330)

If you install or configure SF Oracle RAC 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.

Error message seen in swagent.log after removing the 6.0.1 VRTS packages (2324553)

After removing the 6.0.1 VRTS packages and before rebooting the system, you sometimes see the following message in the `swagent.log` file:

```
vxfs mount: V-3-21272: mount option(s) incompatible with file system  
/dev/vg00/lvol1
```

This message appears because the VRTS packages are removed and the kernel is not yet loaded.

Workaround: Reboot the system.

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.

Some unused packages are not removed after upgrade SFORA from 5.0.1 to SFHA 6.0.1 (2821560)

If you upgrade from 5.0.1 or a previous release, the `VRTSobc33`, `VRTSspbx`, and `VRTSiccco` packages are not uninstalled even if no other package depends on them. You can safely ignore these packages. When you uninstall the product, the installer uninstalls these packages.

Workaround: There is no workaround for this issue. This issue is harmless.

Erroneous resstatechange trigger warning

You may encounter the following warning when you restart resources:

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

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

Issues related to any OS or supported technology

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.

Veritas Storage Foundation 6.0.3 known issues

The known issues for Veritas Storage Foundation for previous releases are as follows:

6.0.1

See "[Veritas Storage Foundation 6.0.1 known issues](#)" on page 39.

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

Veritas Volume Manager 6.0.3 known issues

This section lists Veritas Volume Manager known issues for this release.

On HP-UX 11.31, vxdiskadm option 22-2 Dynamic Reconfiguration (DR) operation 'Remove Luns' might fail with error (2957766)

When the user tries to remove LUNs from the system using `vxdiskadm option 22-2 Dynamic Reconfiguration operation 'Remove Luns'`, the device removal operation fails and reports the following error message:

```
ERROR: Please make sure to remove Luns from Array
```

This is due to the Dynamic Reconfiguration Tool not being able to find devices that are not part of the legacy HP-UX I/O device tree but are seen only in the agile I/O device tree.

Workaround:

Perform the following steps:

- 1 Remove the device with no hardware (NO_HW in output of 'ioscan -fNC disk') using `rmsf (1M)`.
- 2 Run `ioscan (1M)`.
- 3 Run `vxdisk scandisks`.

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 Storage Foundation 6.0.1 known issues

This section lists Veritas Storage Foundation known issues in 6.0.1 release.

Dynamic Multi-Pathing (DMP) 6.0.1 known issues

This section lists Dynamic Multi-Pathing (DMP) known issues in 6.0.1 release.

Path name character limit when converting LVM volumes over DMP to VxVM volumes over DMP (2035399)

The HP-UX `lvdisplay` utility truncates physical volume path names to 22 characters. If a path name is truncated, utilities such as `vxvmconvert` or `vxautoconvert` that depend on the `lvdisplay` output may not function properly. If you intend to use the `vxvmconvert` utility or the `vxautoconvert` utility to convert LVM over DMP to VxVM over DMP, Symantec recommends that you reduce the length of the enclosure name to at most 8 characters before enabling native stack support.

DMP path discovery behavior when a device is removed from PowerPath control (2144891)

To remove a device from PowerPath control, you use the `powermt unmanage` command. When you remove a device from PowerPath control, DMP requires two device discovery cycles to discover the attributes of the paths of the device correctly.

Workaround:

Issue the following command to start the device discovery:

```
# vxdisk scandisks
```

After the discovery completes, issue the command again to start a second device discovery cycle.

I/O fails on some paths after array connectivity is restored, due to high restore daemon interval (2091619)

If a path loses connectivity to the array, the path is marked as suspected to fail and hence is not used for I/O. After the connectivity is restored, the restore daemon detects that the path is restored when the restore daemon probes the paths. The restore daemon makes the path available for I/O. The restore daemon probes the paths at the interval set with the tunable parameter `dmp_restore_interval`. If you set the `dmp_restore_interval` parameter to a high value, the paths are not available for I/O until the next interval.

Changes in enclosure attributes are not persistent after an upgrade to VxVM 6.0.3 (2082414)

The Veritas Volume Manager (VxVM) 6.0.3 includes several array names that differ from the array names in releases prior to release 5.1SP1. Therefore, if you upgrade from a previous release to VxVM 6.0.3, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0.3. Manually reconfigure the enclosure attributes to resolve the issue.

[Table 1-11](#) shows the Hitachi arrays that have new array names.

Table 1-11 Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC
TagmaStoreUSPV	Hitachi_USP-V
TagmaStoreUSPVM	Hitachi_USP-VM

Table 1-11 Hitachi arrays with new array names (*continued*)

Previous name	New name
<New Addition>	Hitachi_R700
Hitachi AMS2300 Series arrays	New array names are based on the Model Number 8x. For example, AMS_100, AMS_2100, AMS_2300, AMS_2500, etc.

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0.3. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

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.

Enclosure name limitation when using HP-UX LVM `vgcreate` command on DMP device (3439724)

For HP-UX LVM on a DMP device, you cannot use the `vgcreate` command if the enclosure-based name of the DMP device contains the 's' character. This is a limitation of the `vgcreate` utility on HP-UX LVM.

Workaround:

Rename the enclosure to replace the 's' with some other character in the name of the enclosure before you run the `vgcreate` command. To rename the enclosure, use the following command:

```
# vxdlmpadm setattr enclosure
enclr_name name=new_enclr_name
```

After disconnecting and reconnecting the Fibre Channel, DMP is unable to present the device tree (2509636)

On some HP-UX 11i version 3 systems, after disconnecting and reconnecting the Fibre Channel, DMP is unable to present the device tree.

Workaround:

Restart the `vxconfigd` daemon with the following command:

```
# vxconfigd -k
```

Failback to primary paths does not occur if the node that initiated the failover leaves the cluster (1856723)

When CVM is configured on non-A/A storage, if a node loses access to the storage through all the primary paths, then all the nodes in the cluster switches to the secondary paths. If the node which raised the protocol leaves the cluster and if all the rest of the nodes in the cluster are seeing the primary paths as healthy, then failback to primary paths never happens.

The pvcreate command failed with DMP devices on IA machines (2482178)

When `dmp_native_support` is enabled on HP-UX 11i version 3 systems, you must run the `pvcreate` command on `/dev/disk/<disk#>` before creating a LVM volume group on the corresponding DMP device.

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

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

Workaround:**To configure path-level attributes**

- 1 Remove the path entries from the `/etc/vx/dmppolicy.info` file.
- 2 Reset the path attributes.

Issues if the storage connectivity to data disks is lost on a CVM slave node while vxconfigd was not running on the node (2562889)

If storage connectivity to data disks is lost on a CVM slave node while `vxconfigd` was not running on the node, this may result in following issues when `vxconfigd` comes up on this node:

- The shared disk groups on the disconnected storage are marked as dgdisabled on the slave node only.
- The shared disk groups are available to rest of the cluster nodes but no transactions, such as VxVM configuration changes, are possible on any shared disk group.
- Attempts to deport such shared disk groups will fail.

Work-arounds:

Use one of the following work-arounds:

- Remove the faulty slave node out of CVM cluster, restore storage connectivity, and rejoin the node to the cluster.
- Restart `vxconfigd` on the CVM master node.

After LUNs remapped using different target ID, DMP reports error with device discovery (2526605)

After LUNs are re-mapped using different target IDs, device discovery fails with the following error message:

```
VxVM vxdisk ERROR V-5-1-16007 Data Corruption Protection Activated -  
User Corrective Action Needed To recover, first ensure that the OS  
device tree is up to date (requires OS specific commands).
```

Workaround:**To recover from this issue**

- 1 Use Operating System (OS) commands to ensure that the OS device tree is up to date.
- 2 Remove the specified devices from VxVM control:

```
# vxdisk rm devicename
```

- 3 Restart device discovery.

```
# vxdisk scandisks
```

DMP native support is not persistent after upgrade to 6.0 (2526709)

The DMP tunable parameter `dmp_native_support` is not persistent after upgrade to DMP 6.0. After you upgrade, set the tunable parameter using the following command:

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

Veritas File System 6.0.1 known issues

This section lists Veritas File System known issues in 6.0.1 release.

Veritas File System known issues

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

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

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

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

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

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

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

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

Deleting a large number of files at the same time drastically increases CPU usage (2129455)

When you delete a large number of files at the same time, the CPU usage drastically increases beyond what you should expect.

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.

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.

Debug kernel panics with spin_deadlock_failure panic string while enabling auxiliary swap space (2521695)

The debug kernel panics with a `spin_deadlock_failure` panic string while enabling auxiliary swap space. The following example is of the relevant part of the stack trace:

```
spinlock+0x50  
vx_inactive+0x140  
vx_vn_inactive+0x30  
vn_rele_inactive+0x1e0  
vx_dnlc_getpathname+0x12b0
```

Not all partitioned directory entries display after exporting a VxFS file system over an HP-UX NFS server (2623412)

After you export a VxFS file system over an HP-UX NFS server, the file system might not list all of the entries in partitioned directories if accessed by NFS clients. This issue is specific to HP-UX NFS servers and VxFS disk layout Version 8 and later.

Workaround: There is no workaround for this issue.

Possible assertion failure in vx_freeze_block_threads_all() (2244932)

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

Workaround: There is no workaround for this issue.

Performance degradation for buffered writes with delayed allocation turned on (2646933)

With the delayed allocation feature turned on, you might observe a performance degradation for buffered writes.

Workaround: Turn off delayed allocation.

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

Veritas Volume Manager 6.0.1 known issues

Veritas Volume Manager known issues

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

In some cases with large LUN setup, the storage disappears after DMP device scan (2828328)

This issue is typically seen on a large LUN setup. In some cases, the storage disappears after the DMP device scan. The DMP device scan is generated with the `vxdisk scandisks` command or the `vxctl enable` command. Even if the OS command `ioscan` can discover devices, VxVM/DMP cannot.

Workaround:

Restarting the `vxconfigd` daemon on the affected node may resolve the issue. If that does not work, you must reboot the system.

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. It may lead to corruption. The recommended format is the Cross-platform Data Sharing (CDS) disk format.

Workaround:

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

System may not boot from a VxVM root disk on a thin LUN (2753626)

The system may fail to boot from a VxVM root disk on a thin LUN. This is an intermittent issue seen only with thin LUNs. The boot process aborts with the following error:

```
System Console is on the Built-In Serial Interface
AF_INET socket/streams output daemon running, pid 52
afinet_prelink: module installed
Starting the STREAMS daemons-phase 1
NOTICE: reading the krs value is failed rc 2
Swap device table: (start & size given in 512-byte blocks)
entry 0 - major is 2, minor is 0x1; start = 0, size = 6242304
Starting vxconfigd in boot mode (pre_init_rc).
pre_init_rc[86]: 81 Illegal instruction
Error returned from vxconfigd -m boot, halting
ERROR: The configuration could not be locked. It may be in use by
another process.
Calling function e000000001a98660 for Shutdown State 1 type 0x1
```

Workaround:

In most cases, rebooting the system resolves the issue.

The vxdmp and other drivers have the incorrect release version (2878024)

The `vxdmp` and other drivers have the incorrect release version. The version displays as 50.0, as shown in the following output:

```
kcmodule -v vxdmp
Module          vxdmp (50.0)
Description     VxVM DMP Subsystem
Timestamp       Wed Aug  1 10:17:12 2012 [50195688]
State           static (best state)
State at Next Boot static (best state)
Capable         static unused
Depends On      interface HPUX_11_31_PERF:1.0
```

vxdg split or join operations can fail for disks with a disk media name greater than or equal to 27 characters (2063387)

If a disk's media name is greater than or equal to 27 characters, certain operations, such as diskgroup split or join, can fail with the following error:

```
VxVM vxdg ERROR : vxdg move/join dg1
                    dg2 failed subdisk_name : Record
already exists in disk group
```

VxVM uses disk media names to create subdisk names. If multiple subdisks are under the same disk, then the serial number, starting from 1, is generated and appended to the subdisk name so as to identify the given subdisk under the physical disk. The maximum length of the subdisk name is 31 characters. If the disk media name is long, then the name is truncated to make room for serial numbers. Therefore, two diskgroups can end up having same subdisk names due to this truncation logic, despite having unique disk media names across diskgroups. In such scenarios, the diskgroup split or join operation fails.

Workaround:

To avoid such problems, Symantec recommends that disk media name length should be less than 27 characters.

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
```

After initializing a disk for native LVM, the first instance of vxdisk list fails with a 'get_contents' error and errant flags are displayed (2074640)

After you initialize a disk that is under the operating system's native LVM control and not under Veritas Volume Manager (VxVM) control by using the `pvcreate path_to_physical_disk` command, the first time that you run the `vxdisk list disk_name` command results in a VxVM error message related to `get_contents`, and the `flags` field is incorrectly populated. However, in the next instantiation of the same command, VxVM does not produce an error and the flags are correctly populated with the LVM tag.

Workaround:

Issue the `vxdisk list disk_name` command a second time.

vxconfigd fails to allocate memory until the daemon is restarted (2112448)

Veritas Volume Manager (VxVM) utilities may fail with the following error message:

```
Memory allocation failure
```

This error implies that there is insufficient memory for the `vxconfigd` daemon. A program's data segment size is enforced by the operating system tunable `maxdsiz`. The default value of `maxdsiz` is 1 GB. With this default `maxdsiz` value, the `vxconfigd` daemon can allocate a maximum of 1 GB of memory.

Workaround:

You might need to increase the operating system `maxdsiz` tunable's value appropriately to increase the data storage segment for the programs.

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

After increasing the value, you must stop and restart the `vxconfigd` daemon. Depending on the `maxdsiz` tunable value, `vxconfigd` can allocate a maximum up to 2 GB of memory on PA machines, and 4 GB of memory on IA machines.

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.

vxdisksetup fails on a LUN that is larger than 1 TB and has the cdsdisk format if the system is using Tachyon HBAs (2146340)

The `vxdisksetup` command fails to initialize a LUN that is larger than 1 TB and has the `cdsdisk` format if the system is using Tachyon HBAs. The `vxdisksetup` command displays the following error:

```
VxVM vxdisk ERROR V-5-1-5433 Device disk_name: init failed:  
Disk is not useable, bad format
```

Workaround:

There is no workaround for this issue.

vxdisk -f init can overwrite some of the public region contents (1190117)

If a disk was initialized by a previous VxVM version or defined with a smaller private region than the new default of 32 MB, then the public region data will be overridden.

Workaround:

Specify explicitly the length of `privoffset`, `puboffset`, `publen`, and `privlen` while initializing the disk.

vxsnap addmir command sometimes fails under heavy I/O load (2441283)

The `vxsnap addmir` command sometimes fails under heavy I/O load and produces multiple errors.

Workaround: Rerun the `vxsnap addmir` command.

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 all dgs list" followed by "vxdg listclone" to get all the disks containing "clone_disk" or "udid_mismatch" flag on respective host.

Known Issue related to EFI disk initialization (2585433)

For disks initialized with EFI format using `idisk`, DA record becomes invisible from "vxdisk list" output after executing "vxdisk scandisks".

Workaround:

For devices to be correctly seen with slices in "vxdisk list" output, VxVM needs to flush the cached open and reopen the disk device. Further, VxVM needs to search for this new EFI format on the disk and generate new DA record.

To recover from this issue

- ◆ To achieve this functionality run following VxVM commands:

```
# vxdisk rm <DANAME>
# vxdctl cacheflush
# vxdisk scandisks
```

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 Oracle RAC 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.

Re-enabling connectivity if the disks are in local failed (lfailed) state (2425977)

In a Cluster Volume Manager (CVM) cluster, you can disable connectivity to the disks at the controller or enclosure level with the `vxddmpadm disable` command. In this case, CVM may place the disks into the `lfailed` state. When you restore connectivity with the `vxddmpadm enable` command, CVM may not automatically clear the `lfailed` state. After enabling the controller or enclosure, you must run disk discovery to clear the locally failed state.

To run disk discovery

- ◆ Run the following command:

```
# vxdisk scandisks
```

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

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

Workaround:

To work around this issue

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

During online migration from LVM to VxVM volumes, LVM sometimes incorrectly reports the remapped LVM device paths as valid LVM volumes

Problem: In a migrated or committed configuration, only the renamed LVM names of the form `<lvolname>_vxlv` are valid LVM volumes. The original LVM names, in turn, point to target VxVM volumes. However, LVM sometimes incorrectly reports these original LVM device paths pointing to VxVM volumes, as valid LVM volumes.

Do not assume these as LVM volumes or do any operations on them, as it would disrupt the application's access to the target VxVM volumes.

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

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

Workaround: The following procedure resolves this issue.

To resolve this issue

- 1 Set the `iotimeout` tunable to 600:

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

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

```
# vxddctl enable
```

The `vxrecover` command fails with 'Cannot execute /etc/vx/type/static/vxassist: No such file or directory' (2857827)

In some cases, the `vxrecover` command fails with the following error:

```
Cannot execute /etc/vx/type/static/vxassist: No such file or directory.
```

This case typically happens if linked volume grow or shrink recovery is triggered as part of the `vxrecover` operation.

Workaround:

To resolve this issue, copy the file from `/usr/sbin/vxassist` to `/etc/vx/type/static/vxassist`.

Diskgroup import of BCV luns using `-o updateid` and `-ouseclonedev` 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 `-ouseclonedev`, 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.

Cluster Volume Manager issues

The following are Cluster Volume Manager issues.

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.

Issue with a configuration with large number of disks when the joining node is missing disks (2869514)

In a configuration with large number of disks (more than 500) where the joining node is missing a few disks (for example. 100 disks), the node join time takes a long time. The joining node attempts to online all the disks as it searches for the missing disks on the node. When the disks are not found the REMOTE LMISSING disks are created on the joining node to complete the join process. This process is found to take time and in such cases the VCS resource online process can timeout.

Workaround:

- Connect the missing disks on the joining node.
- If the intention is to join with missing disks, the VCS timeout needs to be increased.

A master node is not capable of doing recovery if it cannot access the disks belonging to any of the plexes of a volume (2764153)

A master node with missing disks is not capable of doing recovery, as it does not have access to the disks belonging to any of the plexes of a volume.

Workaround:

If other nodes have access to the storage, they can do the recovery. Switch the master role to some other node with better storage connectivity.

Cascaded failure of nodes with ioship enabled may cause the vxconfigd daemon to hang (2865771)

In a shared disk group environment with ioship enabled, the `vxconfigd` daemon may hang in certain cases. When the I/O is initiated from the slave node that has lost connectivity to the disks locally, the I/O is shipped to other nodes. If the node

processing the shipped I/O also leaves the cluster shortly after the first node, and tries to rejoin the cluster as a slave, the cascaded failures may cause the `vxconfigd` daemon to hang.

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.

Server panic after losing connectivity to the voting disk (2787766)

This issue occurs on A/P arrays. If the voting disk loses connectivity to the primary paths, DMP takes some time to analyze the error and fail over the paths. During this time, the `cssd` reports a timeout and panics. When using Oracle ASM over DMP devices, set the `disktimeout` parameter to an appropriate value. This parameter indicates the maximum time allowed for a voting file I/O to complete. If this time is exceeded, the voting disk is marked as offline.

The default of `disktimeout` is 200. If the value of the tunable is less than this value, reset the value to the default value.

Workaround:

To set the `disktimeout` to 200:

```
$CRS_HOME/bin/crsctl set css disktimeout 200 [-force] test
```

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.

Veritas Storage Foundation for Databases (SFDB) tools 6.0.3 known issues

There are no known issues for Veritas Storage Foundation for Databases (SFDB) tools in this release.

The known issues for Veritas Storage Foundation for Databases (SFDB) tools for previous releases are as follows:

6.0.1

See [“Veritas Storage Foundation for Databases \(SFDB\) tools 6.0.1 known issues”](#) on page 55.

Veritas Storage Foundation for Databases (SFDB) tools 6.0.1 known issues

The following are known issues in this release of Veritas Storage Foundation for Databases (SFDB) tools.

Relinking ODM after upgrading from 5.0.x

The `VRTSodm` library path has changed from `/opt/VRTSodm/lib/libodm.sl` to `/opt/VRTSodm/lib/libodm.so`.

After upgrading to from 5.0.x you must update the ODM link for your database to the new `VRTSodm` library path `/opt/VRTSodm/lib/libodm.so`.

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SF Oracle RAC. There is no workaround at this point of time.

Database Storage Checkpoint unmount may fail with device busy (2591463)

In some cases, when a database that is cloned using a Database Storage Checkpoint is shut down, an error similar to the following may occur:

```
SFAE Error:0457: Failed to unmount device
/dev/vx/dsk/datadg/datavol:Ckpt_1317707593_rw_1317708154.
Reason: VxFS returned error : umount: /tmp/clonedb/data: device is
busy
```

Workaround

As an Oracle user, force shut down the clone database if it is up and then retry the unmount operation.

Attempt to use SmartTier commands fails (2332973)

The attempts to run SmartTier commands such as `dbdst_preset_policy` or `dbdst_file_move` fail with the following error:

```
fsppadm: ERROR: V-3-26551: VxFS failure on low level mechanism  
with message - Device or resource busy
```

This error occurs if a sub-file SmartTier command such as `dbdst_obj_move` has been previously run on the file system.

There is no workaround for this issue. You cannot use file-based SmartTier and sub-file SmartTier simultaneously.

Attempt to use certain names for tiers results in error (2581390)

If you attempt to use certain names for tiers, the following error message is displayed:

```
SFORA dbdst_classify ERROR V-81-6107 Invalid Classname BALANCE
```

This error occurs because the following names are reserved and are not permitted as tier names for SmartTier:

- BALANCE
- CHECKPOINT
- METADATA

Workaround

Use a name for SmartTier classes that is not a reserved name.

Clone operation failure might leave clone database in unexpected state (2512664)

If the clone operation fails, it may leave the clone database in an unexpected state. Retrying the clone operation might not work.

Workaround

If retrying does not work, perform one of the following actions depending on the point-in-time copy method you are using:

- For FlashSnap, resync the snapshot and try the clone operation again.
- For FileSnap and Database Storage Checkpoints, destroy the clone and create the clone again.
- For space-optimized snapshots, destroy the snapshot and create a new snapshot.

Contact Symantec support if retrying using the workaround does not succeed.

FlashSnap resync fails if there is an existing space-optimized snapshot (2479901)

If you try a FlashSnap resync operation when there is an existing space-optimized snapshot, the resync operation fails with the following error:

```
Error: VxVM vxdg ERROR V-5-1-4597 vxdg join FS_oradg oradg failed
datavol_snp : Record already exists in disk group
archvol_snp : Record already exists in disk group
```

Workaround

Destroy the space-optimized snapshot first and then perform the FlashSnap resync operation.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0MP2 to 6.0.3 (2003131)

While upgrading from 5.0 MP2 to 6.0.3 the following error message could be seen when running `sfua_rept_migrate`:

```
# /opt/VRTSdbed/migrate/sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
SFORA sfua_rept_migrate ERROR V-81-8903 Could not start repository
database.
/usr/lib/dld.sl: Can't find path for shared library: libcur_colr.1
/usr/lib/dld.sl: No such file or directory
sh: 3845 Abort(coredump)
Symantec DBMS 3.0.85.0 vxdbms_start_db utility
ASA failed. Sybase ASA error code: [134].
Sybase ASA Error text: {{{}}
```

```
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround

To upgrade without an existing SFDB repository set up

- 1 Verify X/Open curses is installed on the system.
- 2 Create the following link:

```
ln -s /usr/lib/libxcurses.1  
/usr/lib/libcur_colr.1
```
- 3 Run:

```
# sfua_rept_migrate
```

Upgrading in an HP Serviceguard environment (2116452)

When upgrading SFDB to 5.1SP1 from the previous release in an HP Serviceguard environment, first verify that the `cmviewcl` command can be executed by a non-root user. This permission change must be done before executing SFDB upgrade commands.

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`.

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.

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.

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'
```

Swverify error related to VRTSdbed observed after a Phase 2 rolling upgrade of SFRAC 6.0.1 on HP-UX 11.31 (2869263)

Upgrade of the SF or SFRAC stack from 5.x to 6.0.1 could display an swverify warning, as follows:

```
WARNING: Directory "/var/vx/vxdba/locks" should have mode "755" but the  
actual mode is "1755".  
WARNING: Directory "/var/vx/vxdba/logs" should have mode "755" but the  
actual mode is "1755".  
WARNING: Fileset "VRTSdbed.DBED,1=/,r=6.0.100.000" had file warnings.
```

Workaround: Ignore the warning, or change the directory permissions to 755 for both `/var/vx/vxdba/locks` and `/var/vx/vxdba/logs`.

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.

Veritas Storage Foundation Cluster File System High Availability 6.0.3 known issues

There are no known issues for Veritas Storage Foundation Cluster File System High Availability in this release.

The known issues for Veritas Storage Foundation Cluster File System High Availability for previous releases are as follows:

6.0.1

See [“Veritas Storage Foundation Cluster File System High Availability \(SFCFSHA\) 6.0.1 known issues”](#) on page 61.

Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) 6.0.1 known issues

This section lists Veritas Storage Foundation Cluster File System High Availability known issues in 6.0.1 release.

Veritas Storage Foundation Cluster File System High Availability known issues

This section describes the known issues in this release of Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

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.

An ENOSPC error may return to the cluster file system application (2867282)

In some cases, when a large number of exclusion zones are set by commands such as `fsadm`, an ENOSPC error may return to the cluster file system application when delegations with free extents are not available.

Workaround: There is no workaround for this issue.

CFS commands might hang when run by non-root (2403263, 3038283)

The CFS commands might hang when run by non-root.

Workaround

To resolve this issue

- ◆ Use `halogin` command to save the authentication information before running any CFS commands on a non-root sessions.

When you run the `halogin` command, VCS stores encrypted authentication information in the user's home directory.

Miscalculated file set usage (2123429)

When file set quotas are enabled, it may be possible for VxFS to get into a state where it thinks a very large number of blocks are allocated to Storage Checkpoints. This issue can be seen using the `fscckptadm` command:

```
# fscckptadm 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:

```
# fscckptadm quotaoff /mnt1
# fscckptadm quotaon /mnt1
# fscckptadm getquotalimit /mnt1
Filesystem  hardlimit  softlimit  usage  action_flag
/mnt1      10000     10000     99
```

The `cfsmntadm add` command may fail with no errors (2169538)

The `cfsmntadm add` command fails, if one host name is a substring of another host name in the list.

Note: VOM is affected by this issue when adding a CFS mount to a cluster that has systems with host names that are substrings of each other.

Workaround

Run the `cfsmntadm` command with the `"all="` option on one of the nodes in the CFS cluster to add the cfsmounts to all nodes.

Multiple CFSmount resources are in a single service group they may not all come online after a reboot (2164670)

In some cases when multiple CFSmount resources are in a single service group, they all may not come online after a reboot. You will need to manually bring them online after a reboot.

Workaround

Create a resource dependency between the various CFSmount resources.

Panic due to null pointer de-reference in vx_bmap_lookup() (2582232, 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.

MountAgent process can get stuck from repeatedly switching a service group from one node to another (2170318)

The `MountAgent` process can get stuck from repeatedly switching a service group from one node to another. This occurs because the `MountAgent` process is waiting for notification, but the notification is unregistered.

Workaround: There is no workaround for this issue.

Veritas Cluster Server 6.0.3 known issues

This section lists Veritas Cluster Server known issues for this release.

For VCS 6.0.1 known issues, see [Veritas Cluster Server 6.0.1 known issues](#).

HPVirtualMachine resource may fail to failover VM after VM is migrated (3207190)

If a Virtual Machine (IVM) being monitored through `HPVirtualMachine` resource is migrated from source to target, then migration process sets NR (Not Runnable) flag for the VM on source. Due to this flag, after a successful migration to target, VCS is not able to bring the `HPVirtualMachine` resource online on source under the following conditions:

- If the VM faults on target subsequently after migration and VCS tries to bring it online on source
- If the VM is brought offline on target through VCS and then brought online through VCS on source
- If the service group containing the `HPVirtualMachine` resource is tried to switch from target to source through `hagrps -switch`
- If the VM is brought offline on target through VCS and brought online through `hagrps -online -any` VM service group may try to come online on source first and go into FAULTED state

Workaround: To bring VM online on source again, execute the following command on the source VMHost:

```
#hpvmmodify -P <VMGuestName> -x modify_status=enabled -x  
runnable_status=enabled
```

Note: Migration of VM from target to source requires `runnable_status` set to disabled value on source VMHost. Hence, after the `runnable_status` is set to enabled value (for bringing VM online), migration from target to source fails.

Veritas Cluster Server 6.0.1 known issues

This section describes the known issues in the VCS 6.0.1 release.

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.

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.

During online migration from LVM to VxVM volumes, LVM sometimes incorrectly reports the remapped LVM device paths as valid LVM volumes

Problem: In a migrated or committed configuration, only the renamed LVM names of the form `<lvolname>_xlv` are valid LVM volumes. The original LVM names, in turn, point to target VxVM volumes. However, LVM sometimes incorrectly reports these original LVM device paths pointing to VxVM volumes, as valid LVM volumes.

Do not assume these as LVM volumes or do any operations on them, as it would disrupt the application's access to the target VxVM volumes.

Issues related to installing and upgrading VCS

Manual upgrade of VRTSvlic package loses keyless product levels [2737124]

If you upgrade the `VRTSvlic` depot manually, the product levels that were set using `vxkeyless` may be lost. The output of the `vxkeyless display` command will not display correctly. To prevent this, perform the following steps while manually upgrading the `VRTSvlic` depot.

1. Note down the list of products configured on the node for keyless licensing.

```
# vxkeyless display
```

2. Set the product level to `NONE`.

```
# vxkeyless set NONE
```

3. Upgrade the `VRTSvlic` package

```
# swremove VRTSvlic
```

This step may report a dependency, which can be safely overridden.

```
swinstall -s 'pwd' VRTSvlic
```

4. Restore the list of products that you noted in step 1.

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

Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

After upgrading from 5.0.1 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]
```

Installer does not detect the duplicate cluster ID in an already configured SF Oracle RAC cluster [2368898]

When you run the installer using `installsfrac -configure` command and if you choose to check the cluster ID, the installer correctly checks if the cluster ID is in use by any other setup. However, if you perform the same check on an already configured SF Oracle RAC cluster, it is unable to detect it.

Workaround: No workaround.

VxSS may go to a faulted state After stack and OS upgrade [2564568]

During the upgrade using the installer, if cluster is in secure mode and the upgrade is from 1123 to 1131, the installer may send a the following warning message.

```
Warning: /opt/VRTSsat/bin/vxatd is not running  
on <system name>. Will be unable to setup trust with shared  
broker, however secure upgrade can still proceed.
```

Workaround: You can ignore this warning and proceed with the upgrade.

Manual install of VRTSvcs depot using /usr/sbin/swinstall -s 'pwd' VRTSvcs may fail [2399744]

Manual installation of VRTSvcs package using `/usr/sbin/swinstall -s `pwd` VRTSvcs` might fail on freshly installed HP-UX machine.

AVXFS is an HP-owned integration product and has dependency on VRTSvlic 3.02.24.0. This causes `/usr/sbin/swinstall` to not select VRTSvcs dependent products from the depot.

Workaround: The workaround for this issue is to remove the AONLINEJFS, OnlineJFS01, and AVXFS depots manually before stack installation. Moreover, before installing VCS on freshly installed HP-UX machine, uninstall older VRTSvlic 3.02.24.0 depot (if installed).

During rolling upgrade swverify command displays errors [2439492]

While performing Rolling upgrade from VCS 5.1SP1 to VCS 6.0 or later, after phase 1 of the Rolling Upgrade process, `swverify` command throws the following errors for all VCS related man pages:

```
"/opt/VRTS/man/man1m/hastatus.1m" missing.
```

This does not cause any product functionality issues. The man pages reappear after phase 2 of the Rolling Upgrade is complete.

Workaround: Not required, as the man pages reappear after Rolling Upgrade phase 2. In you wish to retain all the man pages even after phase 1, copy the `/opt/VRTS/man` directory in a safe location before starting the Rolling Upgrade procedure.

Errors seen during verification of VRTSamf package (2599242)

If VRTSvcsea package is manually upgraded using `swinstall` command from VCS 5.1SP1RP1 to VCS 6.0, permissions for `/opt/VRTSamf/imf/imf_register` file are affected. Due to this, errors are seen during verification of VRTSamf.

Workaround: First uninstall the VRTSvcsea package from VCS5.1SP1RP1 using `swremove` command and then proceed to install VRTSvcsea package from VCS 6.0 release.

Operational issues for VCS

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

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

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

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the *Veritas Cluster Server Installation Guide* for the list of ports and services used by VCS.

- Configure the firewall policy such that the TCP ports required by VCS are not blocked. Refer to your respective firewall or OS vendor documents for the required configuration.

Issues related to the VCS engine

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

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

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

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

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

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

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

If the top level resource which does not have any parent dependency is disabled then the other resources do not come online and the following message is displayed:

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

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

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

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

Workaround: Terminate the agent process.

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

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

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

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

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

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

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

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

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

Oracle group fails to come online if Fire Drill group is online on secondary cluster [2653695]

If a parallel global service group faults on the local cluster and does not find a failover target in the local cluster, it tries to failover the service group to the remote cluster. However, if the firedrill for the service group is online on a remote cluster, offline local dependency is violated and the global service group is not able to failover to the remote cluster.

Workaround: Offline the Firedrill service group and online the service group on a remote cluster.

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

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

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

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

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

Workaround: No workaround.

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

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

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

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

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

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.

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

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

Workaround

- 1 Delete `/var/VRTSat/profile/<user_name>`.
- 2 Delete `/home/user_name/.VRTSat`.

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

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

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

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

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

or

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

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).

Issues related to the bundled agents

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

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

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

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.

NIC agent may report incorrect interface state due to less traffic [2556355]

When PingOptimize is set to 1 and no NetworkHosts is specified, NIC agent depends on packet count to report the health of the interface. If the traffic on the interface is not sufficient enough, NIC agent may report incorrect state of the interface.

Workaround: Any of the following workaround must resolve the issue:

- Setting PingOptimize = 0. This makes NIC agent ping the broadcast address whenever there is no traffic on the interface.
- Setting valid NetworkHosts value. This makes NIC agent to ping NetworkHosts to check health of status.

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.

Resource gets faulted with zero byte logical volume of LVM agent [2393787]

LVM Agent does not support zero byte logical volume and the resource goes into faulted state.

When you configure the resource and try to bring it online, the resource goes into faulted state.

Workaround: No workaround.

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.

Process resource fails to come online if call to exec are present in the profile of the root user [2611530]

Process agent fails to bring the resource online if there are calls to exec in the shell profile file of the root user.

Workaround: Make sure there no calls to exec in the shell profile of the root user.

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.

```
# hagrps -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 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.

CoordPoint agent reports faulted if agile disk naming scheme is used [2871893]

If LevelTwo monitoring is enabled, i.e., if LevelTwoMonitorFreq attribute value is greater than zero, CoordPoint agent reports faulted if agile disk naming scheme is used on HP-UX platform.

Workaround: Use a disk naming scheme other than agile naming scheme.

Probing LVMLogicalVolume resource makes it OFFLINE when LVMLogicalVolume attribute is "" and volume group is not activated [2858185]

If volume group is not activated and the LogicalVolume attribute is set to "" (empty string), a probe of LVMLogicalVolume resource shows the resource state as OFFLINE instead of UNKNOWN.

Workaround: Activate the volume group before configuring LVMLogicalVolume resource under VCS control.

LVMVolumeGroup resource remains ONLINE if VolumeGroup is deactivated outside VCS [2858165]

If LVMVolumeGroup resource is deactivated outside of VCS, clean entry point cannot clean the resource. As a result, the resource state remains ONLINE and clean is called repeatedly.

Workaround: Do not deactivate LVMVolumeGroup from outside SF Oracle RAC.

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.

Concurrency violation in the service group [2870982]

Concurrency violation and data corruption of a volume resource may occur if storage connectivity is lost or all paths under VxDMP are disabled and PanicSystemOnDGLoss is set to 0 This happens when:

- In a cluster configuration, if cluster-wide UseFence attribute is set to SCSI3 and service group contains Volume resource and DiskGroup resource with the PanicSystemOnDGLoss attribute set to 0 (zero).
- If storage connectivity is lost or all paths under VxDMP are disabled, VCS fails over the service group. If storage connectivity is restored on the node on which the service group was faulted and disk group is not deported manually, then volume may get started if disk group is not deported during the service group failover. Thus, volume resource shows its state as online on both the nodes and thus causes concurrency violation. This may lead to data corruption.

Workaround: Ensure that the disk group is deported soon after storage connectivity is restored. Always configure volume resource whenever disk group resource is configured and set the attribute PanicSystemOnDGLoss to 1 or 2 as required.

Service group with LVMLogicalVolume resource does not failover if the underlying storage is not available (2916108)

If service group with LVMLogicalVolume resource is ONLINE and the underlying storage is disconnected or is unavailable, then it does not failover. This is because

the clean entry point is not able to clean the resource, causing the service group to remain ONLINE.

Workaround: Make sure the underlying storage is always available.

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.

Issues related to the agent framework

Agent may fail to heartbeat under heavy load [2073018]

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

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

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

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

Agent framework cannot handle leading and trailing spaces for the dependent attribute (2027896)

Agent framework does not allow spaces in the target resource attribute name of the dependent resource.

Workaround: Do not provide leading and trailing spaces in the target resource attribute name of the dependent resource.

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 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.

Issues related to global clusters

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.

On reboot of cluster nodes that are connected via a single switch, a race condition may cause one of the llc links to not come up (2848001)

If cluster nodes are connected via a single switch and nodes are rebooted multiple times then sometimes a race condition may cause one of the links to be down. Run the `lltstat -nvv` command to know the link that is down.

Workaround: Restart LLT on the rebooted node.

Cannot use CPI response files to add nodes to a cluster that is using LLT over UDP (2869763)

When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/llttab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files.

Workaround: None

GAB known issues

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

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

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

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinitd on user request
```

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

Workaround: There is no workaround for this issue.

GAB can panic due to had not responding (2166263)

GAB can panic due to `had` not responding. This is caused by threads becoming stuck in the `vx_event_wait()` call and the `vx_rwsleep_rec_lock_em()` call.

Workaround: There is no workaround for this issue.

Cluster panics during reconfiguration (2590413)

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

Workaround: There is no workaround for this issue.

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 SF Oracle RAC 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 SF Oracle RAC, 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.

Fencing does not come up on one of the nodes after a reboot (2573599)

If VxFEN unconfiguration has not finished its processing in the kernel and in the meantime if you attempt to start VxFEN, you may see the following error in the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1007 Vxfen already configured
```

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

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

Workaround: Start VxFEN again after some time.

The vxfenswap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)

The `vxfenswap` utility runs the `vxfenconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxfenswap` command using RSH (with the `-n` option), then RSH does not detect the failure of validation of coordination points on a node. From this point, `vxfenswap` proceeds as if the validation was successful on all the nodes. But, it fails at a later stage when it tries to commit the new coordination points to the VxFEN driver. After the failure, it rolls back the entire operation, and exits cleanly with a non-zero error code. If you run `vxfenswap` using SSH (without the `-n` option), then SSH detects the failure of validation of coordination of points correctly and rolls back the entire operation immediately.

Workaround: Use the `vxfenswap` utility with SSH (without the `-n` option).

In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)

When you start server-based I/O fencing, the node may not join the cluster and prints error messages in logs similar to the following:

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1043
Detected a preexisting split brain. Unable to join cluster.
```

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
operation failed.
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@sys1,
domaintype vx; not allowing action
```

The `vxfend` daemon on the application cluster queries the coordination point server (CP server) to check if the cluster members as seen in the GAB membership are registered with the CP server. If the application cluster fails to contact the CP server due to some reason, then fencing cannot determine the registrations on the CP server and conservatively assumes a pre-existing split-brain.

Workaround: Before you attempt to start VxFEN on the application cluster, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

The `cpsadm` command fails if LLT is not configured on the application cluster (2583685)

The `cpsadm` command fails to communicate with the coordination point server (CP server) if LLT is not configured on the application cluster node where you run the `cpsadm` command. You may see errors similar to the following:

```
# cpsadm -s 10.209.125.200 -a ping_cps
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using
environment variable
CPS_NODEID
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

However, if you run the `cpsadm` command on the CP server, this issue does not arise even if LLT is not configured on the node that hosts CP server. The `cpsadm` command on the CP server node always assumes the LLT node ID as 0 if LLT is not configured.

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily,

or if the node is a single-node VCS configuration where LLT is not configured, then the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

Workaround: Set the value of the `CPS_NODEID` environment variable to 255. The `cpsadm` command reads the `CPS_NODEID` variable and proceeds if the command is unable to get LLT node ID from LLT.

Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)

Even if the cluster nodes have no registration on the CP server and if you provide coordination point server (CP server) information in the `vxfenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

Workaround: Manually add the cluster information to the CP server to resolve this issue. Alternatively, you can use `installer` as the installer adds cluster information to the CP server during configuration.

CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for more details.

ping_cps and server_security fail to communicate with the secure CPS Server if the CPS variables are not exported manually (2791763)

The `cpsadm -a ping_cps` and `cpsadm -a server_security` commands fail to communicate with the secure CPS Server from a client if the CPS variables are not exported manually.

Workaround: Set and export the following variables manually on the client cluster:

```
# CPS_DOMAINTYPE="vx"  
# export CPS_DOMAINTYPE  
# EAT_HOME_DIR="/opt/VRTScps"  
# export EAT_HOME_DIR  
# CPS_HOME="/opt/VRTScps"  
# export CPS_HOME  
# CPS_USERNAME="CPSADM@VCS_SERVICES"  
# export CPS_USERNAME
```

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

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

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

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
```

During a race scenario, the larger subcluster of a cluster can lose to a smaller subcluster, which may cause the large subcluster to panic (2858189)

It may happen that during a split-brain scenario, GAB and vxfen modules may take more time to confirm memberships of nodes on a larger subcluster than the time taken to for the same action on a smaller subcluster. So, GAB and vxfen modules on the larger subcluster may lose the race to confirm new node memberships. Hence, the larger subcluster may panic.

CoordPoint agent goes into faulted state if you change the disk naming scheme

If LevelTwo monitoring is enabled (LevelTwoMonitorFreq attribute value set to a value greater than zero), the coordpoint resource goes into faulted state if you use agile disk naming scheme.

Workaround: Use a disk naming scheme other than the agile naming scheme.

Fencing command, vxfenadm, does not print the registration keys in character format (2760308)

The `vxfenadm` command does print character format of keys with leading NULL bytes. This behavior happens because the `vxfenadm` command prints entire registration key as a string and if there is a leading NULL byte in the string key the character format of the key is not printed.

Workaround: None

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 for Oracle RAC 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.

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.

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.

vxfen module does not come up after phased upgrade from release version 4.1MP1 to 6.0.1 (2846209)

With HP-UX 11iv3, after upgrade, vxfen module does not allow raw disks to be specified as coordinator disks. So, even if you set the `vxfen_disk_policy` attribute to **raw** in the `/etc/vxfenmode` file fencing does not come up.

Workaround: Set the `vxfen_disk_policy` to `dmp` in the `/etc/vxfenmode` file.

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

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

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

Issues related to Intelligent Monitoring Framework (IMF)

Registration error while creating a Firedrill setup [2564350]

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

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

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

Workaround: No workaround.

Perl errors seen while using `haimfconfig` command

Perl errors seen while using `haimfconfig` command:

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

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

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

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

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

It should be replaced as follows in `main.cf`:

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

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

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

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

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

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

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

Error messages displayed during reboot cycles [2847950]

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

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

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

Workaround: No workaround.

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

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

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

Workaround: No Workaround.

Error message seen during system shutdown [2954309]

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

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

The system continues to proceed with the shutdown.

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

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

For example:

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

Stopping the daemon gracefully stops all the child processes spawned by the daemon. However, using `kill -9 pid` to terminate a daemon is not a recommended

option to stop a daemon, and subsequently you must kill other child processes of the daemon manually.

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

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

Workaround: Use the correct commands:

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

VCS 5.0.1 Rolling Patch 1 known issues

The VCS issues in this release are as follows:

- The Oracle agent with 11g Release 2 does not support Health check monitoring using the MonitorOption attribute. If the database is 11g Release 2, the MonitorOption attribute for the Oracle agent should be set to 0. The Oracle agent with 11g Release 2 database does not support the Intentional Offline feature. [1975007]
- The ASMInst agent does not support pfile or spfile for the ASM Instance on the ASM diskgroups in 11g Release 2. Symantec recommends that you store the file on the local file system. [1975010]
- If you try to enable debug logs for the DB2 agent, the logs are not written to the `engine_A.log` file. [1954752]

Workaround: Download and install the GNU Awk software from the GNU Web site. Then, create a soft link to the default awk binary on the cluster nodes as follows:

```
# ln -s /usr/local/bin/gawk /bin/awk
```

- The VRTSperl patch takes more than 10 minutes to install on an HP Integrity system node:
On an HP Integrity system node, installing the VRTSperl patch takes more than 10 minutes and requires that VCS is offline during this period. The installation time may vary based on the configuration of the machine on which the VRTSperl patch is being installed.

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 for Oracle RAC known issues

There are no known issues in this release.

The known issues in SF Oracle RAC 6.0.1 are as follows:

- 6.0.1 See ["Oracle RAC issues"](#) on page 90.
- See ["SF Oracle RAC issues"](#) on page 91.

Oracle RAC issues

This section lists the known issues in Oracle RAC.

During installation or system startup, Oracle Grid Infrastructure may fail to start [1933542]

After successful installation of Oracle RAC 11g Release 2 Grid Infrastructure, while executing the `root.sh` script, `ohasd` may fail to start. Similarly, during system startup, Oracle Grid Infrastructure may fail to start though the VCS engine logs may indicate that the `cssd` resource started Oracle Grid Infrastructure successfully.

The following message may be displayed on running the `strace` command:

```
# /usr/bin/strace -ftt -p pid_of_ohasd.bin  
14:05:33.527288 open("/var/tmp/.oracle/npohasd",  
O_WRONLY <unfinished ...>
```

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

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.

SF Oracle RAC issues

This section lists the known issues in SF Oracle RAC for this release.

Installation known issues

This section describes the known issues during installation and upgrade.

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.

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>

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

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

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

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

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

3. Remove the `vxfscckd-pid` file:

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

4. Bring the `vxfscckd` resource online:

```
# hares -online vxfscckd_resname -sys node_name
```

Startup or shutdown failure messages reported for LLT, GAB, VXFEN, and VCSMM (1666327)

If you need to reboot the system when you install SF Oracle RAC, the init scripts for LLT, GAB, VXFEN, and VCSMM report start or stop failure messages. This is because SF Oracle RAC is not yet configured and the required configuration files are not yet generated for these components. These messages may be ignored.

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.

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.

Downloading the patches

The patches included in Veritas Storage Foundation and High Availability Solutions 6.0.3 are available for download from the Symantec website. After downloading the file, use `gunzip` and `tar` to uncompress and extract.

For the 6.0.3 download archive and instructions, visit:

<http://sort.symantec.com/patch/matrix>

Upgrading to version 6.0.3

This chapter includes the following topics:

- [About the installmr script](#)
- [About using the installer to upgrade when the root disk is encapsulated](#)
- [Performing a full upgrade to 6.0.3 on a cluster](#)
- [Performing a full upgrade to 6.0.3 on a standalone system](#)
- [Performing a rolling upgrade to 6.0.3 on a cluster](#)

About the installmr script

Veritas Storage Foundation and High Availability Solutions 6.0.3 provides an installation script. To install the patches that are included in this release, the recommended method is to use the `installmr` script. The `installmr` script lets you install all the patches that are associated with the packages installed. After using the `installmr` script, you may need to restart systems.

[Table 2-1](#) lists the command line options for the `installmr` script.

Table 2-1 Command line options for the `installmr` script

Command Line Option	Function
<code>installmr [<system1> <system2>...]</code>	Specifies the systems on which to run the installation options. If not specified, the command prompts for a system name.
<code>[-ignorepatchreqs]</code>	The <code>-ignorepatchreqs</code> option is used to allow installation or upgrading even if the pre-requisite packages or patches are missed on the system.

Table 2-1 Command line options for the installmr script (*continued*)

Command Line Option	Function
[-precheck]	The <code>-precheck</code> option is used to confirm that systems meet the products install requirements before installing.
[-postcheck]	The <code>-postcheck</code> option is used to check for any issues after installation or upgrading.
[-logpath <log_path>]	The <code>-logpath</code> option is used to select a directory other than <code>/opt/VRTS/install/logs</code> as the location where <code>installmr</code> log files, summary file, and response file are saved.
[-responsefile <response_file>]	The <code>-responsefile</code> option is used to perform automated installations or uninstalls using information stored in a file rather than prompting for information. <code><response_file></code> is the full path of the file that contains configuration definitions.
[-tmppath <tmp_path>]	The <code>-tmppath</code> option is used 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 <timeout_value>]	The <code>-timeout</code> option is used 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.
[-hostfile <hostfile_path>]	The <code>-hostfile</code> option specifies the location of a file containing the system names for installer.
[-keyfile <ssh_key_file>]	The <code>-keyfile</code> option specifies a key file for SSH. When this option is used, <code>-i <ssh_key_file></code> is passed to every SSH invocation.

Table 2-1 Command line options for the installmr script (*continued*)

Command Line Option	Function
[<code>-patchpath <patch_path></code>]	The <code>-patchpath</code> option is used to define the complete path of a directory available to all install systems (usually NFS mounted) that contains all patches to be installed by <code>installmr</code> .
[<code>-rsh</code>]	The <code>-rsh</code> option is used when <code>rsh</code> and <code>rcp</code> are to be used for communication between systems instead of <code>ssh</code> and <code>scp</code> . When the <code>-rsh</code> option is not used, <code>password -less ssh</code> configuration is not a must. If a <code>password -less ssh</code> configuration does not work, then the installer will try <code>rsh</code> .
[<code>-redirect</code>]	The <code>-redirect</code> option is used to display progress details without showing advanced display functionality so output can be redirected to a file.
[<code>-listpatches</code>]	The <code>-listpatches</code> option is used to display product patches in the correct installation order.
[<code>-makeresponsefile</code>]	The <code>-makeresponsefile</code> option generates a response file without doing an actual installation. The text displaying <code>install</code> , <code>uninstall</code> , <code>start</code> , and <code>stop</code> actions are a part of a simulation. These actions are not actually performed on the system.
[<code>-pkgset</code>]	The <code>-pkgset</code> option is used to discover the package set installed on the systems specified.
[<code>-pkgtable</code>]	The <code>-pkgtable</code> option is used to display product depots in correct installation order.
[<code>-pkginfo</code>]	The <code>-pkginfo</code> option is used to display the correct install 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> .

Table 2-1 Command line options for the installmr script (*continued*)

Command Line Option	Function
[<code>-serial</code>]	The <code>-serial</code> option is used to perform install, uninstall, start, and stop operations, typically performed simultaneously on all systems, in a serial fashion.
[<code>-rolling_upgrade</code>]	The <code>-rolling_upgrade</code> option is used 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.
[<code>-rollingupgrade_phase1</code>]	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade phase 1. During this phase, the product kernel depots will be upgraded to the latest version
[<code>-rollingupgrade_phase2</code>]	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade phase 2. During this phase, VCS and other agent depots will be upgraded to the latest version. During this phase, product kernel drivers will be rolling-upgraded to the latest protocol version.
[<code>-ignite</code>]	The <code>-ignite</code> option is used to generate a product bundle which is used by an HPUX Ignite Server for automated installation of all depots and patches for every product. The <code>-ignite</code> option is supported on HPUX only.
[<code>-version</code>]	The <code>-version</code> option is used to check the status of installed products on the system.
[<code>-comcleanup</code>]	The <code>-comcleanup</code> option is used 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.

Table 2-1 Command line options for the installmr script (*continued*)

Command Line Option	Function
[<code>-nolic</code>]	The <code>-nolic</code> option is used to install product depots on systems without entering product licenses. Configuration, startup, or installation of license based features are not performed when using this option.

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

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

Table 2-2 Upgrading using installer when the root disk is encapsulated

Starting version	Ending version	Action required
4.x or 4.x MPx	6.0.3	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.0 or 5.0 MPx or 5.0.1 or 5.0.1 RPx	6.0.3	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 SP1 or 5.1 SP1 RPx	6.0.3	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

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.

The following are the stages of performing a full upgrade on a cluster:

- Freeze service group operations and stop Veritas Cluster Server (VCS) on the cluster.
- Take the nodes offline and install the software patches.

- Bring the nodes online to restart cluster failover services.

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 for 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 version 6.0.3 on an SF Oracle RAC cluster](#)

Performing a full upgrade to 6.0.3 for Veritas Cluster Server

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

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

To upgrade VCS

- 1 Log in as superuser.
- 2 Upgrade the Operating System and reboot the systems if required.
- 3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 6.0.3 maintenance release binaries, change to the directory that contains the installmr script. Start the pre-upgrade check:

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

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

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

- 6 Restart the nodes:

```
# shutdown -r now
```

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 a SFHA and VCS cluster.

To perform a full upgrade to 6.0.3 on an SFHA cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 3 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

```
# /opt/VRTS/bin/hagrp -offline grp_name -sys node_name
```

If the applications are not under VCS control, use native application commands to stop the application.

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

```
# /usr/sbin/mount -v | grep vxfs
```

- 5 Unmount all Storage Checkpoints and file systems:

```
# /usr/sbin/umount /checkpoint_name
```

```
# /usr/sbin/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 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] node1 node2 ... nodeN
```

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

11 Review the output as the program displays the results of the check and saves the results of the check in a log file.

12 Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved.

13 Start the upgrade.

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

Review the output.

14 Restart the nodes:

```
# shutdown -r now
```

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 8, restart each RVG:

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

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

```
# /usr/sbin/mount /filesystem
```

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

```
# /usr/sbin/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 SFCFS cluster.

To perform a full upgrade to 6.0.3 on an SFCFSHA cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 3 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

```
# /opt/VRTS/bin/hagrp -offline grp_name -sys node_name
```

If the applications are not under VCS control, use native application commands to stop the application.

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

```
# /usr/sbin/mount -v | grep vxfs
```

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

```
# /usr/sbin/umount /checkpoint_name
```

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

```
# /usr/sbin/mount -v | grep vxfs
```

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

```
# /usr/sbin/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 `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.

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.

10 On each node, 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
```

- 11 From the directory that contains the extracted and untarred 6.0.3 maintenance release binaries, change to the directory that contains the installmr script.

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

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

- 12 Restart the nodes:

```
# shutdown -r now
```

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

- 14 Bring the CVM service group online on each node:

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

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

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

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

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

- 17 Remount all VxFS file systems on all nodes:

```
# /usr/sbin/mount /filesystem
```

- 18 Remount all Storage Checkpoints on all nodes:

```
# /usr/sbin/mount /checkpoint_name
```

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

Perform the steps in the following procedure to upgrade to version 6.0.3.

Note: If you are upgrading from SF Oracle RAC versions 3.5, 3.5 Update 3, and 3.5 Update 4, you need to first upgrade to version 4.1, then upgrade to 6.0.1, and finally upgrade to 6.0.3. For instructions on upgrading to versions 4.1 and 6.0.1, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for the corresponding version.

Contact Technical Support to obtain the software media or to download the software for version 4.1.

To perform a full upgrade to version 6.0.3 on an SF Oracle RAC cluster

1 Upgrade to version 6.0.3.

See [“Upgrading SF Oracle RAC using the Veritas script-based installation program”](#) on page 107.

Alternatively, use the Web-based installation program to upgrade SF Oracle RAC.

```
# ./webinstaller start
```

Follow the installation prompts to upgrade SF Oracle RAC.

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

3 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys node_name
```

- If the Oracle database is not managed by VCS:

```
# srvctl start database -d db_name
```

4 Start all applications that are not managed by VCS. Use native application commands to start the applications.

- #### 5
- 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 the HP-UX operating system

If you are on an unsupported version of the operating system, you need to upgrade it to HP-UX 11i Version 3 March 2011 Operating Environments Update Release or later supported OS versions.

Note: If you are upgrading from SF Oracle RAC 5.0 or 5.0 RP1 on HP-UX 11i v3, you need to stop VCS before you upgrade the operating system. To stop VCS on all nodes, run the following command as the superuser:

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

If you are upgrading the operating system from HP-UX 11i v2, make sure that you choose the following depots along with the related OS release depots:

- `Base-VxFS-version`
Where *version* is the base VxFS version bundled with the operating system.
- `Base-VxTools-version`
Where *version* is the base VxTools version bundled with the operating system.
- `Base-VxVM-version`
Where *version* is the base VxVM version bundled with the operating system.

To upgrade the operating system from HP-UX 11i v2, run the `update-ux` command specifying the Veritas depots along with the HP-UX operating system depots:

```
# swinstall -s os_path Update-UX
# update-ux -s os_path HPUX11i-DC-OE \
Base-VxFS-version Base-VxTools-version \
Base-VxVM-version
```

where `os_path` is the full path of the directory containing the operating system depots.

To upgrade the operating system from HP-UX 11i v3, run the `update-ux` command as follows:

```
# update-ux -s os_path HPUX11i-DC-OE
```

where `os_path` is the full path of the directory containing the operating system depots.

For detailed instructions on upgrading the operating system, see the operating system documentation.

Upgrading SF Oracle RAC using the Veritas script-based installation program

The product installer performs the following tasks to upgrade SF Oracle RAC:

- Verifies the compatibility of the systems before the upgrade.
- Stops the SF Oracle RAC processes before the upgrade.
- Uninstalls SF Oracle RAC.
- Installs the SF Oracle RAC packages on the nodes.
- Starts SF Oracle RAC on all the nodes.
- Displays the location of the log files, summary file, and response file.

To upgrade to version 6.0.3 using the script-based program

- 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`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `CRSResource.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
```

- 3 Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

If not, a warning message is displayed.

Warning: Backup `/etc/vx/cbr/bk` directory.

- 4 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 5 For Oracle RAC 9i, log in as Oracle user on each node and stop `gsd`:

```
$ $ORACLE_HOME/bin/psctl stop
```

- 6 For Oracle RAC 10g and Oracle RAC 11g:

Stop all Oracle RAC resources.

- If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrps -offline group_name -any
```

- If the database instances are not managed by VCS, then run the following on one node:

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

- 7 ■ 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
```

```
# hagrps -modify oracle_group AutoStart 0
```

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

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

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

8 Stop VCS on all nodes:

```
# hastop -all
```

One way to check whether or not the configuration is valid is to check the main.cf file as follows:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

However, this method can not verify whether all configurations are valid. If SF Oracle RAC was running properly before the upgrade, the configurations are valid.

9 Unmount the non-system mounts on the VxFS file system, which is not under VCS control.

```
# mount -v |grep vxfs
```

```
# umount /mount_point
```

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

```
# fuser -cu /mount_point
```

10 If you plan to upgrade the operating system, stop all ports.

If you are running version 5.1 and later, stop the ports using the installer:

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

If you are running version 5.0.1x and earlier, stop the ports manually as follows:

```
# vcsmmconfig -U
# /sbin/init.d/odm stop
# /sbin/init.d/vxfen stop
# /sbin/gabconfig -U
# kcmodule vxfen=unused
# kcmodule odm=unused
# kcmodule vcsmm=unused
# kcmodule vxglm=unused
# kcmodule vxgms=unused
# lmxconfig -U
# kcmodule lmx=unused
# kcmodule gab=unused
# lltconfig -U
# kcmodule llt=unused
```

- 11 Change to the directory that contains the `installmrscript`:

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

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

- 12 Restart each node in the cluster.

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

- 13 Complete the remaining tasks to finish the upgrade:

See [“Performing a full upgrade to version 6.0.3 on an SF Oracle RAC cluster”](#) on page 104.

Performing a full upgrade 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 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 If required, apply the OS kernel patches.
- 4 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -sys node_name
```

If the applications are not under VCS control, use native application commands to stop the application.

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

```
# mount -v |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` script. Run the `installmr` script:

```
# ./installmr system
```

- 12 Restart the system.

```
# shutdown -r now
```

Performing a rolling upgrade to 6.0.3 on a cluster

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

- [About rolling upgrades](#)
- [Prerequisites for rolling upgrades](#)
- [Performing a rolling upgrade using the script-based 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 are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

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 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.

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 script-based installer

Navigate to the installer program to start the rolling upgrade. The following procedure assumes four nodes: node1, node2, node3, node4.

To perform the rolling upgrade on kernel packages: phase 1

- 1 Log in as superuser to one of the nodes in the first sub-cluster.
- 2 Back up the configuration files on your system.
- 3 For SF Oracle RAC only: Perform the following steps:
 - Stop the Oracle RAC resources on each node.
 - If the database instances are not managed by VCS, then run the following on one node:
 For Oracle RAC 11.2.0.2 and later versions:


```
$ srvctl stop instance -d db_name \  

                            -n node_name
```

 For Oracle RAC 11.2.0.1 and earlier versions:


```
$ srvctl stop instance -d db_name \  

                            -i instance_name
```
 - 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. Failing to perform this step results in the database attempting to come online after the upgrade; the attempt fails due to the presence of old libraries on the system.

```
# haconf -makerw
# hagrpl -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

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

- Unmount all the non-system CFS mount points which are not under VCS control.

```
# mount -v |grep vxfs | grep cluster
# umount /mount_point
```

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

```
# fuser -cu /mount_point
```

- 4 Stop the applications that are not managed by VCS and that which use VxFS or VxVM disk groups on each node, whether local or CFS.

Use native application commands to stop the application.

- 5 Unmount all the VxFS file systems which is not under VCS control.

```
# mount -v |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 -cu /mount_point
```

- 6 Start the installer.

```
# ./installmr -rollingupgrade_phase1 node1 node2
```

The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes.

- 7 Type **y** to continue.

The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade.

8 Type **y** to continue.

If you choose to specify the nodes, type **n** and enter the names of the nodes.

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. After the installer shows the package list, it detects whether or not 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 installer to automatically switch service groups

The downtime is the time that it normally takes to fail over the service group.

After switching the failover service group, the installer detects if there are online parallel service groups on the node to be upgraded. If there are online parallel service groups, the installer prompts you to use the installer for automatic switching of service groups.

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

The installer performs the upgrade and restarts processes. If some processes fail to start, you may need to reboot the nodes and manually check the status of the cluster.

9 For SF/SFHA/SFCFSHA/SF Oracle RAC: Restart the nodes in the first sub-cluster:

```
# shutdown -r now
```

10 Perform the following steps on the nodes in the first sub-cluster:

- Manually mount the VxFS and CFS file systems that VCS does not manage.
- Start all applications that VCS does not manage. Use native application commands to start the applications.

11 Perform the following steps for SF Oracle RAC:

- Bring the Oracle database service group online.
If VCS manages the Oracle database:

```
# hagrps -online oracle_group -sys node_name
```

If VCS does not manage the Oracle database:

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

- Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 12 Complete step 1 to step 5 on the nodes in the second sub-cluster.
 - 13 Start the installer on the nodes in the second sub-cluster.

```
# ./installmr -rollingupgrade_phase1 node3 node4
```

- 14 For VCS: Repeat step 7 through step 8 and step 10.
 For SF/SFHA/SFCFS: Repeat step 7 through step 10.
 For SF Oracle RAC: Repeat step 7 through step 11.

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

Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. In this phase, the installer installs all non-kernel depots on all the nodes in cluster and restarts the cluster.

- 1 Begin phase 2 of the upgrade by typing **y** to continue.
 The installer determines the remaining packages to upgrade.
- 2 Press **Enter** to continue.
 The installer stops Veritas Cluster Server (VCS) processes but the applications continue to run.
 In case of failure in the startup of some of the processes, you may need to reboot the nodes and manually check the cluster's status.
- 3 Type **y** to continue.
 The installer performs prechecks, uninstalls old packages, and installs the new packages.
 It performs post-installation tasks and completes the configuration for the upgrade.
- 4 Type **y** to help Symantec improve the automated installation. If you have network connection to the Internet, the installer checks for updates.
 A message prompting review of the summary file appears.

5 Type **y** to read the installation summary file.

6 Verify the cluster's status:

```
# hastatus -sum
```

Uninstalling version 6.0.3

This chapter includes the following topics:

- [About uninstalling Veritas Storage Foundation and High Availability Solutions 6.0.3](#)
- [About the `uninstallmr` script](#)
- [Rolling back using the `uninstallmr` script](#)
- [Uninstalling 6.0.3 with the Web-based installer](#)

About uninstalling Veritas Storage Foundation and High Availability Solutions 6.0.3

This section describes how to roll back either by using the `uninstallmr` script or the Web-based installer.

Roll back of version 6.0.3 to the 6.0.1 release is supported for the following products:

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)
- Veritas Storage Foundation for Oracle RAC (SF for Oracle RAC)
- Veritas Cluster Server (VCS)
- Dynamic Multi-Pathing (DMP)

Symantec recommends that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the previous release.

For extensive uninstallation and reinstallation procedures, refer to the appropriate product's Installation Guide.

About the `uninstallmr` script

Veritas Storage Foundation and High Availability Solutions 6.0.3 provides a script that you can use to roll back to the 6.0.1 release. To uninstall the patches that are included in this release, the recommended method is to use the `uninstallmr` script.

[Table 3-1](#) lists the command line options for the `uninstallmr` script.

Table 3-1 Command line options for the `uninstallmr` script

Command Line Option	Function
<code>uninstallmr [<system1> <system2>...]</code>	Specifies the systems on which to run the <code>uninstallmr</code> script. If not specified, the command prompts for a system name.
<code>[-logpath <log_path>]</code>	The <code>-logpath</code> option is used 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>[-responsefile <response_file>]</code>	The <code>-responsefile</code> option is used to perform automated installations or uninstalls using information stored in a file rather than prompting for information. <code><response_file></code> is the full path of the file that contains configuration definitions.
<code>[-tmppath <tmp_path>]</code>	The <code>-tmppath</code> option is used 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>[-hostfile <hostfile_path>]</code>	The <code>-hostfile</code> option specifies the location of a file containing the system names for <code>uninstallmr</code> .
<code>[-keyfile <ssh_key_file>]</code>	The <code>-keyfile</code> option specifies a key file for SSH. When this option is used, <code>-i <ssh_key_file></code> is passed to every SSH invocation.
<code>[-rsh]</code>	The <code>-rsh</code> option is used when <code>rsh</code> and <code>rcp</code> are to be used for communication between systems instead of <code>ssh</code> and <code>scp</code> .

Table 3-1 Command line options for the `uninstallmr` script (*continued*)

Command Line Option	Function
[<code>-redirect</code>]	The <code>-redirect</code> option is used to display progress details without showing advanced display functionality so that output can be redirected to a file.
[<code>-makeresponsefile</code>]	The <code>-makeresponsefile</code> option generates a response file without doing an actual installation. The text displaying <code>install</code> , <code>uninstall</code> , <code>start</code> , and <code>stop</code> actions are a part of a simulation. These actions are not actually performed on the system.
[<code>-serial</code>]	The <code>-serial</code> option is used to perform <code>install</code> , <code>uninstall</code> , <code>start</code> , and <code>stop</code> operations, typically performed simultaneously on all systems, in a serial fashion.
[<code>-version</code>]	The <code>-version</code> option is used to check the status of installed products on the system.
[<code>-listpatches</code>]	The <code>-listpatches</code> option is used to display product patches in correct installation order.
[<code>-comcleanup</code>]	The <code>-comcleanup</code> option is used 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.
[<code>-timeout <timeout_value></code>]	The <code>-timeout</code> option is used 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>-ignorepatchreqs</code>]	The <code>-ignorepatchreqs</code> option is used to allow installation or upgrading even if the pre-requisite packages or patches are missed on the system.

Rolling back using the `uninstallmr` script

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

To roll back on a standalone system

- 1 Browse to the directory that contains the `uninstallmr` script.
- 2 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -sys node_name
```

If the applications are not under VCS control, use native application commands to stop the application.

- 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 -v |grep vxfs
```

- 4 Stop all VxVM volumes. For each disk group enter:

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

Verify that no volumes remain open.

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

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

```
# ./uninstallmr
```

- 6 Restart the system:

```
# shutdown -r now
```

The `uninstallmr` script removes 6.0.3 patches. After maintenance rollback completes, modules are loaded and processes are restarted. `uninstallmr` will also report any warning happened during uninstallation.

To roll back in a cluster setup

- 1 Log in as the super user on one of the nodes in the cluster.
- 2 On each node, take the Oracle resources in the VCS configuration file (`main.cf`) offline.

```
# hagrpl -offline oracle_group -sys node_name
```

For example:

```
# hagrpl -offline oral -sys sys1
```

```
# hagrpl -offline oral -sys sys2
```

These commands stop the Oracle resources under VCS control.

If the database is not managed by VCS, stop the Oracle database as follows:

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

- 3 Stop VCS:

```
# hastop -all
```

- 4 Use native application commands to stop the applications that use VxFS or VxVM disk groups on each node and that are not under VCS control, whether local or CFS.
- 5 Unmount all the non-system VxFS file system which is not under VCS control.

```
# mount -v |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 -cu /mount_point
```

- 6 Run the `uninstallmr` command. On each node in the cluster, type:

```
# ./uninstallmr
```

To roll back on all the cluster nodes in one go, type:

```
# ./uninstallmr system1 system2 systemn
```

- 7 Restart the nodes:

```
# shutdown -r now
```

- 8 Manually mount the VxFS and CFS file systems that VCS does not manage.
- 9 Start all applications that VCS does not manage. Use native application commands to start the applications.

Uninstalling 6.0.3 with the Web-based installer

This section describes how to uninstall this release with the Web-based installer.

To uninstall 6.0.3

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Start the Web-based installer.
- 4 On the **Select a task and a product** page, select **Uninstall 6.0.3** from the **Task** drop-down list and click **Next**.
- 5 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Validate**.
- 6 After the validation completes successfully, click **Next** to uninstall the patches on the selected system.
- 7 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 8 After the installer stops the processes, the installer removes the patches from the specified system. Click **Next**.

- 9 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.
- 10 Click **Finish**.
The Web-based installer prompts you for another task.