

Veritas Storage Foundation™ and High Availability Solutions Release Notes

Linux

5.1 Service Pack 1 Rolling Patch 3

Veritas Storage Foundation and High Availability Solutions Release Notes 5.1 Service Pack 1 Rolling Patch 3

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

Product version: 5.1 SP1 RP3

Document version: 5.1SP1RP3.4

Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction, release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

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

North America and Latin America [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	4	
Chapter 1	About Veritas Storage Foundation and High Availability Solutions	11
	Introduction	11
	About the installrp script	12
	The installrp script options	13
	Overview of the installation and upgrade process	15
	Changes introduced in 5.1 SP1 RP3	16
	Changes related to Veritas Cluster Server	16
	System requirements	16
	Supported Linux operating systems	17
	VMware Environment	18
	Database requirements	18
	Recommended memory and swap space	19
	List of products	19
	Fixed issues	19
	Veritas Volume Manager fixed issues	20
	Veritas File System fixed issues	35
	Veritas Storage Foundation Cluster File System fixed issues	42
	Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues	45
	Veritas Storage Foundation for Oracle RAC fixed issues	45
	Veritas Cluster Server fixed issues	47
	Veritas Storage Foundation for Databases (SFDB) tools fixed issues	55
	Known issues	56
	Issues related to installation	56
	Veritas Dynamic Multi-pathing known issues	66
	Veritas Storage Foundation known issues	71
	Veritas Cluster Server known issues	97
	Veritas Storage Foundation Cluster File System known issues	114
	Veritas Storage Foundation for Oracle RAC known issues	115

	Software limitations	123
	Veritas Volume Manager software limitations	123
	Veritas Cluster Server software limitations	124
	List of RPMs	124
	Downloading the 5.1 SP1 RP3 archive	128
Chapter 2	Installing the products for the first time	129
	Installing the Veritas software using the script-based installer	129
	Installing Veritas software using the Web-based installer	130
	Starting the Veritas Web-based installer	131
	Obtaining a security exception on Mozilla Firefox	131
	Installing 5.1 SP1 RP3 with the Veritas Web-based installer	131
Chapter 3	Upgrading to 5.1 SP1 RP3	133
	Prerequisites for upgrading to 5.1 SP1 RP3	133
	Downloading required software to upgrade to 5.1 SP1 RP3	133
	Supported upgrade paths	134
	Upgrading to 5.1 SP1 RP3	134
	Performing a full upgrade to 5.1 SP1 RP3 on a cluster	135
	Upgrading to 5.1 SP1 RP3 on a standalone system	144
	Upgrading to 5.1 SP1 RP3 on a system that has encapsulated boot disk	146
	Performing a rolling upgrade using the installer	146
	Performing a phased upgrade to SFCFSHA and SFRAC	152
	Upgrading the operating system	173
	Verifying software versions	173
Chapter 4	Uninstalling version 5.1 SP1 RP3	175
	About removing Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP3	175
	Uninstalling Veritas Storage Foundation Cluster File System 5.1 SP1 RP3	175
	Preparing to uninstall Veritas Storage Foundation Cluster File System	176
	Uninstalling Veritas Storage Foundation Cluster File System	177
	Uninstalling Veritas Storage Foundation for Oracle RAC	178
	Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP3	180

Appendix A	DMP as a multi-pathing solution in an Oracle VM Server for SPARC environment	183
	DMP as a multi-pathing solution in an Oracle VM Server for SPARC environment	183
	Enabling DMP in the control and alternate I/O domains	184
	Enabling DMP path failover in the guest domain	185

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [Introduction](#)
- [About the installrp script](#)
- [Overview of the installation and upgrade process](#)
- [Changes introduced in 5.1 SP1 RP3](#)
- [System requirements](#)
- [List of products](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [List of RPMs](#)
- [Downloading the 5.1 SP1 RP3 archive](#)

Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Rolling Patch 3 release.

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

<http://www.symantec.com/docs/TECH75506>

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

<http://www.symantec.com/docs/TECH74012>

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

This rolling patch applies to the following releases of Storage Foundation and High Availability products:

- Storage Foundation and High Availability Solutions 5.1 SP1
- Storage Foundation and High Availability Solutions 5.1 SP1 RP1
- Storage Foundation and High Availability Solutions 5.1 SP1 RP2
- Storage Foundation and High Availability Solutions 5.1 SP1 RP2
- VirtualStore 5.1 SP1 PR3

This rolling patch is available as 5.1 SP1 RP3.

Given that this rolling patch applies to the previously released 5.1 SP1 platform RP releases, Symantec does not plan on the following releases:

- 5.1 SP1 PR2 RP1
- 5.1 SP1 PR3 RP1

About the installrp script

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP3 provides an upgrade script.

See “[Supported upgrade paths](#)” on page 134.

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

The installrp script options

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

Command Line Option	Function
[<i>system1 system2...</i>]	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.
[<i>-precheck</i>]	Use the <i>-precheck</i> option to confirm that systems meet the products' installation requirements before the installation.
[<i>-postcheck</i>]	Use the <i>-postcheck</i> option after an installation or upgrade to help you determine installation-related problems and provide troubleshooting information.
[<i>-logpath log_path</i>]	Use the <i>-logpath</i> option to select a directory other than <i>/opt/VRTS/install/logs</i> as the location where the <i>installrp</i> log files, summary file, and response file are saved.
[<i>-responsefile response_file</i>]	Use the <i>-responsefile</i> option to perform automated installations or uninstalls using information stored in a file rather than prompting for information. <i>response_file</i> is the full path of the file that contains configuration definitions.
[<i>-tmppath tmp_path</i>]	Use the <i>-tmppath</i> option to select a directory other than <i>/var/tmp</i> as the working directory for <i>installrp</i> . This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
[<i>-hostfile hostfile_path</i>]	Use the <i>-hostfile</i> option to specify the location of a file containing the system names for installer.
[<i>-keyfile ssh_key_file</i>]	Use the <i>-keyfile</i> option to specify a key file for SSH. When you use this option the <i>-i ssh_key_file</i> is passed to every SSH invocation.

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

Command Line Option	Function
[<code>-patchpath patch_path</code>]	Use the <code>-patchpath</code> option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by <code>installrp</code> .
[<code>-rsh</code> <code>-redirect</code> <code>-listpatches</code> <code>-makeresponsefile</code> <code>-pkginfo</code> <code>-serial</code> <code>-version</code>]	<p>Use the <code>-rsh</code> option when <code>rsh</code> and <code>rcp</code> are to be forced for communication though <code>ssh</code> and <code>scp</code> is also setup between the systems.</p> <p>Use the <code>-redirect</code> option to display progress details without showing the progress bar.</p> <p>Use the <code>-listpatches</code> option to display product patches in the correct installation order.</p> <p>Use the <code>-makeresponsefile</code> option to generate a response file without doing an actual installation. The text that displays <code>install</code>, <code>uninstall</code>, <code>start</code>, and <code>stop</code> actions are simulations. These actions are not performed on the system.</p> <p>Use the <code>-pkginfo</code> option is used to display the correct installation order of packages and patches. This option is available with or without one of following options: <code>-allpkgs</code>, <code>-minpkgs</code>, and <code>-recpkgs</code>.</p> <p>Use the <code>-serial</code> option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion.</p> <p>Use the <code>-version</code> option to have the installer check and report the installed products and their versions. Identifies the installed and missing RPM and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPM and patches where applicable.</p>

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

Command Line Option	Function
<code>[-upgrade_kernelpkgs -upgrade_nonkernelpkgs]</code>	<p>Use the <code>-upgrade_kernelpkgs</code> option for the rolling upgrade's upgrade of kernel packages to the latest version.</p> <p>Use the <code>-upgrade_nonkernelpkgs</code> option for the rolling upgrade's upgrade of non-kernel packages. In this phase, VCS packages and other agent packages are upgraded to the latest versions. Product kernel drivers are upgraded to the latest protocol version.</p>

Overview of the installation and upgrade process

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

To install the Veritas software for the first time

- Skip this step if you are upgrading to 5.1 SP1 RP3. If you are installing 5.1 SP1 RP3 for the first time:

 - Download Storage Foundation and High Availability Solutions 5.1 SP1 from <http://fileConnect.symantec.com>.
 - Extract the tar ball into a directory called `/tmp/sfha51sp1`.
 - Check <http://sort.symantec.com/patches> to see if there are any patches available for the 5.1 SP1 Installer. Download applicable P-patches and extract them to the `/tmp` directory.
 - Change the directory to `/tmp/sfha51sp1`:


```
# cd /tmp/sfha51sp1
```
 - Install the 5.1 SP1 software. Follow the instructions in the Installation Guide.


```
# ./installer -require complete_path_to_SP1_installer_patch
```
- Download SFHA 5.1 SP1 RP3 from <http://sort.symantec.com/patches> and extract it to a directory called `/tmp/sfha51sp1rp3`.

3 Check <http://sort.symantec.com/patches> to see if there are patches available for the 5.1 SP1 RP3 installer. Download applicable P-patches and extract them to the `/tmp` directory.

4 Change the directory to `/tmp/sfha51sp1rp3`:

```
# cd /tmp/sfha51sp1rp3
```

5 Install 5.1 SP1 RP3:

```
# ./installrp -require complete_path_to_SP1RP3_installer_patch
```

Changes introduced in 5.1 SP1 RP3

This section lists the changes in 5.1 SP1 RP3.

Changes related to Veritas Cluster Server

Veritas Cluster Server includes the following changes in 5.1 SP1 RP3:

Sybase and SybaseBk agents support Intelligent Monitoring Framework (IMF)

Symantec has updated the Sybase and SybaseBk agents to provide online monitoring (PRON) IMF support.

You should set the IMF Mode attribute to 2 (PRON support) only. IMFRegList is added to the Sybase type definition. During the enabling of IMF, make sure that the type definition is updated and IMFRegList is correctly set.

For in-depth monitoring, the DetailMonitor attribute should be used. The LevelTwoMonitorFrequency attribute is not used in this release. As a result, when IMF is enabled and DetailMonitor is also enabled for Sybase resource, the actual in-depth monitor will happen at a frequency of the DetailMonitor attribute value * the IMF MonitorFreq attribute value

System requirements

This section describes the system requirements for this release

Supported Linux operating systems

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

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

<http://www.symantec.com/docs/TECH75506>

The Veritas Storage Foundation (SF) 5.1 SP1 RP3 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) till Red Hat Enterprise Linux 5 (RHEL 5) with Update 9 (2.6.18-348.el5) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Red Hat Enterprise Linux 6 (RHEL 6) with Update 1 (2.6.32-131.0.15.el6 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Red Hat Enterprise Linux 6 (RHEL 6) with Update 2 (2.6.32-220.el6 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Red Hat Enterprise Linux 6 (RHEL 6) with Update 3 (2.6.32-279.el6 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel), SP3 (2.6.16.60-0.54.5) or SP4 (2.6.16.60-0.85.1), on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) with SP1 (2.6.32.12-0.7 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- OEL supported but not OEL UBK
- Oracle Linux 5 (OL 5) with Update 3 (2.6.18-128.el5 kernel) till Oracle Linux 5 (OL 5) with Update 8 (2.6.18-308.el5 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Linux (OL) 6.1 (2.6.32-131.0.15.el6 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Linux (OL) 6.2 (2.6.32-220.el6 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Linux (OL) 6.3 (2.6.32-279.el6 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)

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

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

Symantec supports only Oracle Enterprise Linux, Red Hat Enterprise Linux (RHEL Compatible Mode only), and SUSE Linux Enterprise Server distributed kernel binaries.

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

Information about the latest supported Red Hat erratas and updates and SUSE service packs is available in the following TechNote. Read this TechNote before you install Symantec products.

<http://www.symantec.com/docs/TECH75506>

VMware Environment

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

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

Database requirements

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

<http://www.symantec.com/docs/TECH74389>

Note: Veritas Storage Foundation (SF) and Veritas Storage Foundation Cluster File System (SFCFS) do not support running SFDB tools with DB2 and Sybase, but they support running Oracle, DB2, and Sybase on VxFS and VxVM.

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

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

Recommended memory and swap space

Symantec recommends the following memory and swap space sizes:

- On the system where you run the installation, use the following guidelines for memory minimums when you install on:
 - One to eight nodes, use 1 GB of memory
 - More than eight nodes, use 2 GB of memory or more
- On the system where you run the installation, use the following guidelines for swap space when you install on:
 - One to eight nodes, use *(number of nodes + 1)* x 128 MB of free swap space
 - For a minimum of 256 MB for 1 node and a maximum of 1 GB of swap space for 8 or more nodes

List of products

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

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation and High Availability (SFHA)
- Veritas Storage Foundation Cluster File System (SFCFS)
- Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
- Veritas Volume Manager (VM)
- Veritas File System (FS)
- Veritas Cluster Server (VCS)
- Veritas Dynamic Multi-Pathing (DMP)
- Symantec VirtualStore (SVS)

Fixed issues

This section describes the issues fixed in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

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

Veritas Volume Manager fixed issues

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP3.

[Table 1-2](#) describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP3.

Table 1-2 Veritas Volume Manager 5.1 SP1 RP3 fixed issues

Fixed issues	Description
925653	Node join fails for higher <code>CVMTIMEOUT</code> value.
2886402	When re-configuring devices, <code>vxconfigd</code> hang is observed.
2868790	In RHEL 6, there are some changes in the <code>sysfs</code> tree layout preventing <code>vxesd</code> to collect the HBA topology information through <code>sysfs</code> .
2858853	After master switch, <code>vxconfigd</code> dumps core on old master.
2838059	VVR Secondary panic in <code>vol_rv_update_expected_pos</code> .
2836798	In VxVM, resizing simple EFI disk fails and causes system panic/hang.
2826125	VxVM script daemon is terminated abnormally on its invocation.
2818840	Enhance the <code>vxdumpasm</code> utility to support various permissions and "root:non-system" ownership can be set persistently.
2815517	<code>vxvg adddisk</code> allows mixing of clone & non-clone disks in a <code>DiskGroup</code> .
2801962	Growing a volume takes significantly large time when the volume has version 20 DCO attached to it.
2783293	After upgrade to RHEL5.8(2.6.18-308), all paths get disabled when <code>deport/import</code> operations are invoked on shared dgs with SCSI-3 mode.
2775960	In secondary CVR case, IO hang is seen on a DG during SRL disable activity on other DG.

Table 1-2 Veritas Volume Manager 5.1 SP1 RP3 fixed issues (*continued*)

Fixed issues	Description
2774406	System may panic while accessing data change map volume.
2763206	The <code>vxdisk rm</code> command dumps core when disk name of very large length is given.
2760181	Panic hit on secondary slave during logowner operation.
2756059	System may panic when large <code>cross-dg</code> mirrored volume is started at boot.
2754819	Diskgroup rebuild through <code>vxmake -d</code> loops infinitely if the diskgroup configuration has multiple objects on a single cache object.
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.
2739709	Disk group rebuild fails as the links between volume and vset were missing from the <code>vxprint -D -</code> output.
2739601	VVR: VRAS: repstatus output occasionally reports abnormal timestamp.
2735951	Uncorrectable write error is seen on subdisk when SCSI device/bus reset occurs.
2729911	IO errors are seen during controller reboot or array port disable/enable.
2726148	System becomes unbootable after setting <code>dmp_native_support</code> tunable on and reboot is done.
2715129	<code>Vxconfigd</code> hangs during Master takeover in a CVM (Clustered Volume Manager) environment.
2711312	New symbolic link is created in root directory when a FC channel is pulled out.
2689845	Data disk can go in the error state when data at the end of the first sector of the disk is same as MBR signature.
2688308	When re-import of disk group fails during master takeover, other shared disk groups should not be disabled.
2680343	Manual disable/enable of paths to an enclosure leads to system panic.
2664825	DiskGroup import fails when disk contains no valid UDID tag on config copy and config copy is disabled.

Table 1-2 Veritas Volume Manager 5.1 SP1 RP3 fixed issues (*continued*)

Fixed issues	Description
2657797	Starting a RAID5 volume fails, when one of the sub-disks in the RAID5 column starts at an offset greater than 1TB.
2656803	Race between vxnetd start and stop operations causes panic.
2648176	Performance difference on Master vs Slave during recovery via DCO.
2647975	Customer ran <code>hastop -local</code> and shared dg had splitbrain.
2637217	Document new storage allocation attribute support in vradmin man page for <code>resizevol/resizeurl</code> .
2627126	IO hang is seen due to IOs stuck at DMP level.
2627056	<code>vxmake -g DGNAME -d desc-file</code> fails with very large configuration due to memory leaks.
2626741	Using <code>vxassist -o ordered</code> and <code>mediatype:hdd</code> options together do not work as expected.
2626199	<code>vxddmpadm list dmpnode</code> prints incorrect path-type.
2620556	IO hung after SRL overflow.
2620555	IO hang due to SRL overflow and CVM reconfig.
2612301	Upgrading kernel on encapsulated root disk does not work properly.
2606978	Private region I/O errors due to <code>DMP_PATH_FAILED</code> do not trigger path failover in linux vxio.
2606709	IO hang when SRL overflow and reboot one node.
2606695	Machine panics in CVR (Clustered Volume Replicator) environment while performing I/O Operations.
2599526	IO hang is seen when DCM is zero.
2578336	Failed to online the cdsdisk.
2576602	<code>vxdbg listtag</code> should give error message and display correct usage when executed with wrong syntax.
2575172	I/Os hung on master node after reboot the slave node.
2567618	VRTSexplorer coredumps in <code>checkhaapi/print_target_map_entry</code> .

Table 1-2 Veritas Volume Manager 5.1 SP1 RP3 fixed issues (*continued*)

Fixed issues	Description
2566174	Null pointer dereference in <code>volcvm_msg_rel_gslock()</code> results in panic.
2561012	The offset of private(and or public) region of disks are shown incorrect in the <code>vxdisk</code> list output which could lead to DG import problem as well as I/O errors and system hang reported by VxFS or other applications.
2560843	In VVR (Veritas Volume Replicator) setup I/Os can hang in slave nodes after one of the slave node is rebooted.
2560835	I/Os and <code>vxconfigd</code> hung on master node after slave is rebooted under heavy I/O load.
2558261	VxVM unable to setup/un-setup powerpath devices.
2556781	In cluster environment, import attempt of imported disk group may return wrong error.
2556467	Disabling all paths and reboot of the host causes losing of <code>/etc/vx/.vxdmprawdev</code> records.
2526623	Memory leak detected in CVM code.
2516584	Startup scripts use 'quit' instead of 'exit', causing empty directories in <code>/tmp</code> .
2513101	User data corrupted with disk label information.
2495332	<code>vxcdsconvert</code> fails if the private region of the disk to be converted is less than 1 MB.
2486301	During SFHA CPI installation, "vxfs" package installation fails on system having good amount of luns coming from A/P-F array.
2483265	VxVM <code>vxdmp V-5-0-0</code> I/O error occurred (<code>errno=0x205</code>).
2441937	<code>vxconfigrestore precommit</code> fails with awk errors.
2425259	<code>vxdg join</code> operation fails with <code>VE_DDL_PROPERTY: Property not found in the list</code> .
2413763	Uninitialized memory read results in a <code>vxconfigd</code> core dump.
2389554	The <code>vx dg listssbin</code> output is incorrect.
2348199	<code>vxconfig</code> dumps core while importing a Disk Group.
2333540	<code>vxdisk resize</code> may incorrectly reduce the size of a VxVM disk.

Table 1-2 Veritas Volume Manager 5.1 SP1 RP3 fixed issues (*continued*)

Fixed issues	Description
2277558	<code>vxassist</code> outputs a misleading error message during snapshot related operations.
2257850	<code>vxdiskadm</code> leaks memory while performing operations related to enclosures.
2252680	<code>vxtask abort</code> does not appropriately cleanup the tasks.
2227678	Second rlink goes into DETACHED STALE state in multiple secondaries environment when SRL has overflowed for multiple rlinks.
2216951	<code>vxconfigd</code> dumps core because <code>chosen_rlist_delete()</code> hits NULL pointer in linked list of clone disks.
2149922	Record the diskgroup import and deport events in syslog.
2104887	<code>vx dg import</code> error message needs improvement for cloned diskgroup import failure.
2088426	Re-onlining of disks in DG during DG deport or destroy.
2000585	<code>vxrecover</code> doesn't start remaining volumes if one of the volumes is removed during the <code>vxrecover</code> command run.
1903700	Removing mirror using <code>vxassist</code> does not work.
1765916	VxVM socket files don't have proper write protection.
1675482	The <code>vx dg list dgname</code> command shows configuration copy in new failed state.
1533134	Warnings regarding depreciated SCSI ioctl appear in syslog.
1431223	The <code>vradmin syncvol</code> and <code>vradmin syncrvg</code> commands do not work if the remote diskgroup and vset names are specified when synchronizing vsets.
1291519	After multiple VVR migrate operations, <code>vrstat</code> fails to output statistics.

This section describes Veritas Volume Manager fixed issues in 51 SP1 RP2 P3.

Table 1-3 Veritas Volume Manager 51 SP1 RP2 P3 fixed issues

Fixed issues	Description
2771452	IO hung because of hung port deletion.
2751102	VxVM(Veritas Volume Manager) buffer spillover leads to VxFS (Veritas File System) inode corruption which further leads to panic in SF (Storage Foundation) suite.
2741240	Invoking "vxdg join" operation during heavy IO load results in a transaction failure and leaves disks in an intermediate state.
2729501	vxdmpadm exclude vxvm path=<> results in excluding unexpected set of paths.
2722850	DMP fail over hangs when the primary controller is disabled while I/O activity is ongoing.
2717362	/etc/vx/vxpath_links file is absent.
2710579	Do not write backup labels for CDS disk - irrespective of disk size.
2701654	Phantom DMP disk partition causes panic.
2700792	The VxVM volume configuration daemon may dump a core during the Cluster Volume Manager(CVM) startup.
2700486	vradmind core dumps when Primary and Secondary have the same hostname and an active Stats session exists on Primary.
2700086	EMC BCV (NR) established devices are resulting in multiple dmp events messages (paths being disabled/enabled).
2698860	vxassist mirror failed for thin LUN because statvfs failed.
2688747	Logowner local sequential I/Os starved with heavy I/O load on logclient.
2675538	vxdisk resize may cause data corruption.
2674465	Adding/removing new LUNs causes data corruption.
2666163	A small portion of possible memory leak incase of mix (clone and non-cloned) diskgroup import.
2660151	vxconfigd is generating a series of LVM header messages for devices (CLONES/replicated devices)Secondary EMC MirrorView LUNS in an error state.
2644248	vxunroot fails as root partition "logvol" mounted on /var/log.

Table 1-3 Veritas Volume Manager 51 SP1 RP2 P3 fixed issues (*continued*)

Fixed issues	Description
2643634	Message enhancement for a mixed(non-cloned and cloned) dg import.
2636005	vxddmpadm device exclusion is not working as expected.
2635476	Volume Manager does not recover a failed path.
2621465	When detached disk after connectivity restoration is tried to reattach gives Tagid conflict error.
2608849	VVR Logowner: local I/O starved with heavy I/O load from Logclient.
2600939	Linux: Make the vol_use_rq tunable accessible in non-debug builds.
2595557	multiple execution of "sysctl -a" caused OS panic.
2557984	VxDMP size in /proc/partitions table doesn't match with the vxdisk list <dmpnode>.
2553729	Disk groups do not get imported and 'clone_disk' flag is seen on non-clone disks after upgrade of VxVM.
2527289	Site consistency: Both sites become detached after data/dco plex failue at each site, leading to I/O cluster wide outage.
2509291	The vxconfigd daemon hangs if host side i/o paths are failing.
2495186	With TCP protocol used for replication, I/O throttling happens due to memory flow control.
2423701	Upgrade of VxVM caused change in permissions.
2419948	Race between the SRL flush due to SRL overflow and the kernel logging code, leads to a panic.
2417546	Raw devices are lost after reboot and cause permissions problem.
2390998	System panicked during SAN reconfiguration because of the inconsistency in dmp device open count.
2365486	In 2-nodes SFRAC configuration, after enabling ports systems panics due to improper order of acquire and release of locks.
2253970	Support per-disk maxiosize for private region I/Os.
2234731	Encapsulation of root disk fails and throws an error while copying the temp boot files.

Table 1-3 Veritas Volume Manager 51 SP1 RP2 P3 fixed issues (*continued*)

Fixed issues	Description
2061082	The <code>vxddladm -c assign names</code> command should work for devices with native support not enabled (VxVM labeled or TPD).

This section describes Veritas Volume Manager fixed issues in 51 SP1 RP2 P2.

Table 1-4 Veritas Volume Manager 51 SP1 RP2 P2 fixed issues

Fixed issues	Description
2185069	panic in <code>vol_rv_mdship_srv_start()</code> .

This section describes Veritas Volume Manager fixed issues in 51 SP1 RP2 P1.

Table 1-5 Veritas Volume Manager 51 SP1 RP2 P1 fixed issues

Fixed issues	Description
2536667	Slave node panics when private region I/O and dg deport operation are executed simulatenously.
2530279	<code>vxesd</code> has been built without any thread locking mechanism.
2524936	DG disabled after <code>vold</code> found the process file table is full.
2510523	The <code>ls -l</code> command hang during RMAN backup on VVR/RAC cluster.
2489350	Memory leak in VVR.
2484334	Panic in <code>dmp_stats_is_matching_group()</code> .
2483053	Primary Slave node runs out of memory, system hang on <code>VRTSvxvm</code> .
2438426	VxVM is failing to correctly discover ZFS LUNs presented via PP after excluding/including <code>libvxpp.so</code> .
2432006	Pending read count with <code>kio</code> cache is not decremented when read object is locked in transaction.
2431448	CVR:I/O hang while transitioning to DCM mode.
2428170	IO hung on Mirror volume and return error on DMP disk, but <code>phydisk(/dev/sdbw)</code> is OK.

Table 1-5 Veritas Volume Manager 51 SP1 RP2 P1 fixed issues (*continued*)

Fixed issues	Description
2420386	Data corruption creating data in a vxfs filesystem, while being grown with vxresize on efi thinrclm disks.
2419803	Secondary Site panics in VVR (Veritas Volume Replicator).
2419486	Data corruption occurs on changing the naming scheme.
2390431	VVR: system crash during DCM flush not finding the parent_sio volsiodone+.
2344186	Volume recovery is not clearing the need sync flag from volumes with DCO in BADLOG state. Thus nodes are unable to join the cluster.
2235382	IO hung in DMP while restoring a path in presence of pending IOs on local A/P class LUN.
2169726	CLONE : Disk group is imported using a Non-cloned and cloned disks, it can lead to data corruption.
2148851	<code>vxdisk resize</code> failed to resize the disk which is expanded physically from array console.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP2

[Table 1-6](#) describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP2.

Table 1-6 Veritas Volume Manager 5.1 SP1 RP2 fixed issues

Fixed issues	Description
2484685	Race between two vol_subdisk sios while doing done processing which causes one thread to free sio_fsm_priv before other thread accesses it
2480600	I/O permanent hung on master node when IO size larger than 512K, and 32+ threads write in parallel
2440349	DCO volume may grow into any 'site' even when 'alloc=site:xxxx' is specified by a list of 'site' to be limited
2431470	vxpfto uses DM name when calling vxdisk, but vxdisk will match DA name first and thus cause corruption
2431423	CVR: Panic in vol_mv_commit_check after I/O error on DCM

Table 1-6 Veritas Volume Manager 5.1 SP1 RP2 fixed issues (*continued*)

Fixed issues	Description
2428875	I/O on both nodes (wait for the DCM flush started), and crash the slave node, lead to the master reconfiguration hang
2428631	Allow same fence key to be used for all Disk groups
2425722	vxsd move operation failed for disk size greater than or equal to 2TB
2425551	IO hung for 6 mintues when reboot the slave node, if there is I/O on both master and slave
2424833	while autosync_deport#2 primary logowner hits ted assert nmcom_send_msg_tcp
2422058	The VxVM diskgroup can NOT import with I/O fencing enabled of both dmp and raw mode
2421067	Vxconfigd hung in both nodes of primary
2419348	DMP panic: race between dmp reconfig and dmp pass through ioctl
2413904	Multiple issues are seen while performing Dynamic LUN reconfiguration
2411698	VVR:iohang: On I/O to both master and slave
2410845	Lots of 'reservation conflict' messages seen on 51SP1RP1P1 clusters with XIV arrays
2408771	vxconfigd does not scan and discover all the storage device; some storage devices are skipped
2407192	Application I/O hangs because of race between CVM reconfiguration and Log-owner change protocol
2406292	Panic in vol_subdisksio_delete()
2400654	Stale arrayinfo file can cause vxdmpadm commands to hang
2400076	vxconfigd produced kernel panic when you run "vxinstall" command
2396293	I/Os loaded, sanboot failed with vxconfigd core dump
2388725	Panic in dmp_get_dmpsymbols when attempting to load an APM
2387993	While testing including/excluding libvxppso vxconfigd goes into disabled mode

Table 1-6 Veritas Volume Manager 5.1 SP1 RP2 fixed issues (*continued*)

Fixed issues	Description
2386120	Enhancement request to add diagnostic logging to help triage a CVM master takeover failure situation
2385694	IO hung if the slave node rebooted
2385680	vol_rv_async_childdone+1147
2383158	VVR: vxio panic in vol_rv_mdship_srv_done+680
2379029	Changing of enclosure name is not working for all devices in enclosure
2369786	VVR:A deadlock about NM_ERR_HEADR_IO
2369177	DDL: do_diskio function should be able to handle offset greater than 2 TB
2365951	Growto failing with error V-5-1-10128 Unexpected kernel error in configuration update
2364253	VVR: Kernel memory is leaked on VVR secondary while using SO snapshots
2359814	vxconfigbackup doesn't handle errors well
2358321	Remove usage of __invalidate_device() from VxVM Symbol is no longer in kABI whitelist
2357798	CVR:Memory leak due to unfreed vol_ru_update structure
2357507	In presence of large number of NR (Not-Ready) devices, server panics due to NMI triggered and when DMP continuously generates large no of path disable/enable events
2356744	VxVM script daemons should not allow its duplication instance in itself
2349352	During LUN provisioning in single path IO mode environment a data corruption is observed
2346470	Excluding and including a LUN in a loop triggers a huge memory leak
2337694	TP "vxdisk -o thin list" showing size 0 for over 2TB LUNs on RHEL5
2337353	vxdmpadm include vxvm dmpnodename= <i>emcpower</i> # includes all excluded dmpnodes along with the requested one
2334534	In CVM environment, vxconfigd level join is hung when Master returns error "VE_NO_JOINERS" to a joining node and cluster nidmap is changed in new reconfiguration

Table 1-6 Veritas Volume Manager 5.1 SP1 RP2 fixed issues (*continued*)

Fixed issues	Description
2323925	If rootdisk is encapsulated and if install-db is present, clear warning should be displayed on system boot
2322752	Duplicate DA records seen for NR devices upon restart of vxconfigd
2320917	vxconfigd core dump and lost dg config after removing volume and disk on thin reclaim LUN
2317703	Vxesd/Vxconfigd leaks file descriptors
2316297	After applying 51SP1RP1 error message "Device is in use" appears during boot time
2299670	Disk Groups created on EFI LUNs do not auto import at boot time using VxVM version 51SP1 and later
2286559	kernel heap corruption detected panic after array controller reboot
2263317	CLONE: Diskgroup import with dgid needs to be clearly documented in manual for the case in which original dg was destroyed and cloned disks are present
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated
2255182	Handling misconfiguration of CLARiiON array reporting one failovermode value through one HBA and different from other HBA
2253970	Support per-disk maxiosize for private region I/Os
2253552	Leak in vxsfdefault_parsey at function vxsf_getdefault (*val)
2249113	vol_ru_recover_primlog_done return the same start address to be read from SRL, if the dummy update is greater than MAX_WRITE
2248730	vxdg import command hangs as vxrecover daemon (spawned by vxdg) doesn't close standard error stream
2242268	panic in voldr1_unlog
2240056	vxdg move' transaction not completing and backups fail
2237089	vxrecover might start the recovery of data volumes before the recovery of the associated cache volume is recovered
2234821	etrack 1946267 - DMP can't detect the re-enabled os device status does not work on RHEL5

Table 1-6 Veritas Volume Manager 5.1 SP1 RP2 fixed issues (*continued*)

Fixed issues	Description
2232411	supporting NetApp Metro Cluster
2228531	cvm master vxconfigd process hung in vol_klog_lock()
2218470	Some of the VxVM init scripts need to be compliant to the Linux Standard Base
2205108	SVS 51SP1: vxconfigd clubbing all luns in a single dmpnode
2204752	Multiple VM commands succeed but throw "GPT entries checksum mismatch" error message for hpdisk format
2200670	vxattachd does not recover disks if disk group is not imported
2197254	While creating volumes on thinrlm disks, the option "logtype=none" does not work with vxassist command
2196918	Snapshot creation with cachesize fails, as it doesn't take into account diskgroup alignment
2196480	The disk initialization failed due to wrong number of cylinders reported in devintf_disk_geom_raw() gotten from raw geometry
2194685	vxconfigd daemon core dump during array side switch ports disable and re-enable
2193429	IO policy not getting preserved when vold is restarted and migration from one devlist to other is taking place
2190020	SUSE complains dmp_deamon applying 1m continuous memory paging is too large
2179259	DMP SCSI bypass needs to be enhanced to handle I/O greater than 2 TB
2165394	CLONE: dg imported by selecting wrong disks After destroying original dg, when try to import clone devices without useclonedev option with dgname, then it import dg with original disks
2154287	Improve handling of Not-Ready(NR)devices which are triggering "VxVM vxdmp V-5-3-1062 dmp_restore_node: Unstable path" messages
2152830	In multilevel clone disks environment, regular DG import should be handled properly and in case of DG import failure, it should report correct error message

Table 1-6 Veritas Volume Manager 5.1 SP1 RP2 fixed issues (*continued*)

Fixed issues	Description
2144775	Failoverpolicy "local" is not getting preserved after upgrade from 51RP1/Sles10Sp2 to 51Sp1/Sles10Sp3
2139179	SSB check invalid when lun copy
2094672	CVR: vxconfigd on master hangs while reconfig is running in cvr stress with 8 users
2033909	In SF-RAC configuration, IO hung after disable secondary path of A/PG array Fujitsu ETERNUS3000
1791397	VVR:RU thread keeps spinning sending START_UPDATE message repeatedly to the secondary
1675599	Memory leaks in DDL and ASLs

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Volume Manager in this release.

Table 1-7 Veritas Volume Manager 5.1 SP1 RP1 fixed issues

Fixed issues	Description
1426480	The <code>volcvm_clear_pr()</code> ioctl now propagates the error returned by DMP to the caller.
1829285	<code>vxconfigd</code> no longer dumps core while assigning a unique native name to a disk.
1869002	Introduced a Circular buffer at the vold level for master-slave communication.
1940052	<code>vxconfigd</code> no longer hangs on the master after removing the HBA alias from the zone and node leave followed by join
1959513	The <code>-o noreonline</code> option of a diskgroup import is now propagated to slave nodes.
1970560	<code>vxconfigd</code> no longer dumps core on the master node when <code>vxconfigd</code> on a passive slave dies and command shipping is in progress.

Table 1-7 Veritas Volume Manager 5.1 SP1 RP1 fixed issues (*continued*)

Fixed issues	Description
2015467	Improved performance for NetBackup 6.5.5 on Veritas Storage Foundation 5.1 VxVM mapping provider.
2038928	Added support for creating and using pre-5.1 SP1 release diskgroups on CDS-initialized disks.
2062190	vxrootadm : split/join operation fails when there is a rvg present in the rootdg/backupdg
2080730	The vxvm exclude file and vxdmp exclude file contents are now consistent after updating the files using the vxdiskadm command and vxdmpadm command.
2082450	vxdisk resize should output more meaningful error message
2088007	possibility of reviving only secondary paths in dmp_revive_paths()
2105547	tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations
2125306	Fixed a few issues related to loading the HBA API library and the vxinstall script.
2129477	vxdisk reclaim no longer fails after resize.
2129989	EVA ASL should report an error message if pref_bit is not set for a LUN
2133503	Renaming enclosure results in dmpevents.log reporting 'mode for Enclosure has changed from Private to Private'
2148682	while shipping a command node hangs in master selection on slave nodes and master update on master node
2152830	In a multi-level clone disks environment, a regular diskgroup import is now handled properly, and in the case of a diskgroup import failure, the correct error message is now displayed.
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core.
2160199	An upcoming master can now import a shared diskgroup, which allows the master takeover to succeed.
2166682	checks needed to make sure that a plex is active before reading from it during fsmirror read interface

Table 1-7 Veritas Volume Manager 5.1 SP1 RP1 fixed issues (*continued*)

Fixed issues	Description
2172488	FMR2 restore doesn't sync the existing snapshot mirrors
2179479	The flags on a disk are no longer incorrectly set as "error" even after running the <code>vxdisk scandisks</code> command after creating a PV and volume group.
2181631	striped-mirror volume cannot be grown across sites with <code>-oallowspansites w/ DRL</code>
2183984	system panic in <code>dmp_update_stats()</code> routine
2188590	an <code>ilock</code> acquired by a SLAVE node for a read on a DCL object can lead to IO hang when the node becomes MASTER before reading is done
2191693	<code>vxmpadm native list</code> command now displays output and error messages.
2194492	VxVM-ASM co-existence enablement
2199496	Fixed a data corruption issue with the "site mirror" Campus Cluster feature.
2199836	The system with the root volume group and DMP native support enabled now successfully boots and mounts.
2200670	The <code>vxattachd</code> command can now recover disks if even if the disk group is not imported.
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault
2215216	<code>vxkprint</code> now reports TP-related values.
2220926	The <code>vxprivutil -D set attr</code> command no longer causes the <code>vxprivutil</code> command to hang.
2226813	Rlinks no longer remain disconnected with the UDP protocol if data ports are specified.
2227923	Renaming an enclosure is now persistent.
2234844	An <code>asm2vxfs</code> conversion with Linux partitions no longer fails.

Veritas File System fixed issues

This section describes Veritas File System fixed issues in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas File System: Issues fixed in 5.1 SP1 RP3

This section describes Veritas File System fixed issues in 5.1 SP1 RP3.

[Table 1-8](#) describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP3.

Table 1-8 Veritas File System 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2884144	Changes required in Linux kernel for RHEL6.3 support.
2841059	full fsck fails to clear the corruption in attribute inode 15.
2839871	process hung in vx_extentalloc_delicache.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2597347	fsck segmentation fault bc_rgetblk ().
2560244	LM/CFS stress test hit kernel panic.

Table 1-9 Veritas File System 5.1 SP1 RP2 P2 fixed issues

Fixed issues	Description
2821152	Internal test hit an assert "f:vx_dio_physio:4,1" via "vx_dio_rdwri" on SLES11_SP1.
2753944	VxFS hang in vx_pd_create.
2745357	The fix is such, that the piggy back data would not be ignored if it's of type VX_IOT_ATTR, in vx_populate_bpdata, to improve performance.
2715028	fsadm -d hang during vx_dircompact.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2670022	Duplicate file names can be seen in a directory.
2666249	File system hangs during Filestore operations.
2651922	Performance degradation of 'll' and high SYS% CPU in vx_ireuse().
2566875	A write(2) operation exceeding the quota limit fails with an EDQUOT error.
2086902	System crash when spinlock was held too long.

This section describes Veritas File System fixed issues in 5.1 SP1 RP2 P1.

Table 1-10 Veritas File System 5.1 SP1 RP2 P1 fixed issues

Fixed issues	Description
2648078	Manual upgrade for VxFS will fail in when ODM is running.
2631276	QIO does not work in a partitioned directory.
2624262	Dedup:fsdedup.bin hit oops at vx_bc_do_brelse.
2611279	Filesystem with shared extents may panic.
2609010	FCL seek error is causing replication failure.
2607631	vxportal may fail to load on boot.
2605089	SVS: clone and golden image corruption with and without customization.
2591702	Improper messages are seen in engine.log on 5.1SP1RP2 with RHEL6 env.
2588593	usage of volume in the output of df command do not back to beginning after created files and deleted files.
2576794	/bin/sh: bad interpreter: Permission denied.
2561334	using flockfile() instead of adding new code to take lock on ._fspadm_enforcesq file descriptor before writing into it.
2529201	fscdsconv limits are wrong in cdslimittab.
2528819	VxFS thread create warning messages.
2515459	mount command still hanged even with the fix of e1466351.
2515380	ff_vxfs ERROR: V-3-24347: program limit of 30701385 exceeded .
2429281	quota file becomes huge with large user ID .
2350956	fsck fails with the following message ASSERT(devid == 0 (start == VX_HOLE && devid == VX_DEVID_HOLE)) failed.
2342067	While mounting FS, in error case we may see panic in vx_kill_sb.
2332314	Internal Test with odm hit an assert fdd_odm_aiodone:3.
2314212	DB2 9.5 onwards can cause contention of the mmap_sem.
2275679	On SLES10 machine with high I/O activity some writes may appear to be stalled.
2271797	Internal hit an assert f:vx_getblk:1a.

Table 1-10 Veritas File System 5.1 SP1 RP2 P1 fixed issues (*continued*)

Fixed issues	Description
2246127	Mount should perform read ahead on IAUs.
2212686	The read /write performance may degrade over a period of time due to memory fragmentation.
2191031	Performance may be impacted due to the frequent triggering of swapping daemon (kswapd) .
1590963	Requirement for vx_maxlink tunable on Linux.
1590324	Umount can hang if linux is using inotify.

Veritas File System: Issues fixed in 5.1 SP1 RP2

[Table 1-11](#) describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP2.

Table 1-11 Veritas File System fixed issues

Fixed issues	Description
2340953	cfs.stress.enterprise hit an assert f:vx_iget:1a.
2481984	file system will hang if a customer creates 400 shares
2247387	LM stress.S3 test hit an assert "vx_ino_update:2"
2486589	threads blocked behind vx_ireuse_steal
2440584	node panic in vx_sync() during shutdown
2424240	Dedup ioctl sharing extents incorrectly under certain scenarios
2431674	panic in vx_common_msgprint() via vx_inactive()
2480935	V-3-26626: File Change Log IOTEMP and ACCESTEMP index creation failure for /vx/fsvm with message Argument list too long
1892045	Improve the memory allocation for per-cpu data.
2413172	There is a priority 1 issue reported by AXA Rosenberg for Filestore replication and issue seems related to VxFS
2399228	TRuncate up size updates can be missed

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

Fixed issues	Description
2412604	It does not work when set homedir user softlimit numspace quota after generate data
2242630	Remove limits on inode and buffer cache sizes
2422574	Reboot one node and the node can't mount file system , after turn on the homedir quota on
2403126	cfs recovery didn't finished timely in the primary node after one slave left.
2283893	Add functionality of free space defragmentation through fsadm.
2372093	new fsadm -C hung
2387609	User quota corruption
2371710	user quota information corrupts on 5.1SP1
2371903	newline in vx_osdep.c: snprintf(cmp->cm_name, sizeof(cmp->cm_name), "vxclonefs-%d" breaks native LVM(pvs)
2384831	vxfs panic in iput() from vx_softcnt_flush() ,after filesystem full fsck,and run reboot
2399178	fsck : pass2c needs performance enhancements
2374887	Accessing FS hung. FS marked full fsck after reboot of node.
2374887	Accessing FS hung. FS marked full fsck after reboot of node.
2283315	cfs-stress_S5 hits assert of "f:vx_reorg_emap:10 via vx_extmap_reorg"
2368737	RCQ processing code should set FULLFSCK flag if it finds a corrupt indirect block.
1956458	fsckpt_fbmap for changed blocks failed with ENXIO due to inode mapped to hole in ILIST of down stream checkpoint
2337470	In the process of shrink fs, the fs out of inodes, fs version is 5.0MP4HF*
2332460	vxquota slow on some systems
2300682	Question about IOTemp on fsppadm query
2316793	After removing files df command takes 10 seconds to complete
2302426	Unaligned Reference Fault in vx_copy_getemap_structs

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

Fixed issues	Description
2272072	Threads stuck in vx_rwsleep_rec_lock_em
2290800	investigation on ilit HOLE
2192895	Panic while set/get acls - possible race condition
2059611	Panic in vx_unlockmap() due to NULL ml_tranp
2282201	vxdump core dumped whilst backing up layout 7 local VxFS file system
2337737	killing IOs to a CFS and ls command to the same CFS is hanging.
2316094	There was discrepancy between vxi_bcache_maxkbyte and vx_bc_bufhwm.
2253938	EAU delegation timeouts
2419991	ncheck: no way to limit output to specific filesets, as with limiting output to specific inodes
2419989	ncheck -i does not limit output to the specified inodes when using -o device/block/sector
2074806	dm_punch_hole request does not invalidate pages
2296107	Operation not applicable appear on fspadm query result
2246579	Panic at getblk() when growing a full filesystem with fsadm
2061177	fsadm -de' command erroring with 'bad file number' on filesystem(s) on 5.0MP3RP1
1475345	write() system call hangs for over 10 seconds on VxFS 3.5 on 11.23
2251223	df -h after removing files takes 10 seconds
2253617	LM stress aborted due to "run_fsck : Failed to full fsck cleanly".
2220300	vx_sched' is hogging CPU resources.
2161379	repeated hangs in vx_event_wait()
1949445	hang due to large number of files in a directory
1296491	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2239412	system panics while writing to cfs share exported as NFS to ESX server4.1.

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

Fixed issues	Description
2169324	Test LM-stress_S5 hits an assert of "f:vx_idelxwri_off:5a vai vx_trunc_tran"

Veritas File System: Issues fixed in 5.1 SP1 RP1

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

Table 1-12 Veritas File System fixed issues

Fixed issues	Description
1296491	Fixed issues seen during a force unmount of a parent cluster file system while a child was being mounted or unmounted.
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.
2030119	fspadm core dumps when analysing a badly formatted XML file, is resolved
2032525	Fixed the cause of an NFS stale file handle.
2061554	Sequential extents are now collated.
2111921	Improved the performance of VxFS file systems with concurrent I/O or direct I/O enabled.
2149659	Fixed the cause of an error that resulted during the truncate operation of a file with a shared extent in the presence of a Storage Checkpoint containing FileSnaps.
2162822	During online migration from ufs to vxfs, df command returns a non-zero return value.
2163084	The <code>listxattr()</code> call now uses <code>rwlock</code> .
2169273	During online migration, nfs export of the migrating file system leads to system panic
2177253	A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for <code>vxupgrade</code> purposes
2178147	Linking a IFSOC file now properly calls <code>vx_dotdot_op()</code> , which fixes the cause of a corrupted inode.

Table 1-12 Veritas File System fixed issues (*continued*)

Fixed issues	Description
2181833	The <code>vxfilesnap</code> command no longer gives an incorrect error message on a Storage Checkpoint file system.
2184528	<code>fsck</code> no longer fails to repair corrupt directory blocks that have duplicate directory entries.
2178147	Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system
2198553	A forced unmount now properly clears the <code>bd_super</code> structure member.
2221623	Fixed a performance loss due to a <code>delxwri_ilst</code> spin lock with the default values for <code>vx_idelxwri_timelag</code> .
2226257	Fixed the cause of a system panic in the <code>in_ilst_add()</code> call, which led to corruption in the <code>vx_ftenter()</code> codepath when using named data streams.

Veritas Storage Foundation Cluster File System fixed issues

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP3.

[Table 1-13](#) describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP3.

Table 1-13 Veritas Storage Foundation Cluster File System 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2925918	mount point getting hanged after starting async conversion of a ckpt to 'nodata'.
2781444	<code>cfs->stress->reconfig</code> test hit asserts "f:vx_inactive_remove:1d" and "f:xted_validate_pnqtran:5" on RHEL6_U3.

Table 1-13 Veritas Storage Foundation Cluster File System 5.1 SP1 RP3 fixed issues (*continued*)

Fixed issues	Description
2392571	hacf-verify/etc/VRTSvcS/conf/config/ gives error after phased upgrade on 2nd subcluster which results in stack startup issues.
1852788	cfs.stress.reconfig test exited with error message run_fsck: Failed to complete extended operations and hit an assert f:vx_msgprint:ndebug via vx_recover on RHEL6_U3.

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP2 P2.

Table 1-14 Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 P2 fixed issues

Fixed issues	Description
2824895	vcscvmqa "cfsumount" test getting fail.
2684573	Enhancement request for force option of the cfsumount command.
2674639	VxFS returning error 61493 (VX_EFCLNOSPC) on CFS.
2669724	CFSMountAgent core dump due to assertion failure in VCSAgThreadTbl::add().

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP2 P1.

Table 1-15 Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 P1 fixed issues

Fixed issues	Description
2565400	Poor read performance with DSMC (TSM) backup on CFS filesystems.
2530747	Threads doing rename, create and remove can wait indefinitely for the exclusive cluster wide dirlock of the file-system.
2528888	CFS mount fails after recovery from I/O path failure.
2326037	Write operation on a Cluster mounted filesystem may fails with ENOENT.

Table 1-15 Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 P1 fixed issues (*continued*)

Fixed issues	Description
2433934	Performance discrepancy between CFS and standalone VxFS using NFS.
2196308	Performance degradation on CFS with 5.1 RP2.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP2

[Table 1-16](#) describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP2.

Table 1-16 Veritas Storage Foundation Cluster File System fixed issues

Fixed issues	Description
2146573	qdetails performance downgraded

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP1

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

Table 1-17 Veritas Storage Foundation Cluster File System fixed issues

Fixed issues	Description
1296491	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2169538	The cfsmntadm add command fails, if one host name is a substring of another host name in the list
2180905	fsadm -S shared mountpoint gives wrong error message when mount points of veritas filesystem version is other than 8.
2181833	"vxfilesnap" gives wrong error message on checkpoint filesystem on cluster
2184114	In a large filesystem, stat operation on cluster mount leaves the file system frozen for too long leading to CVMVoldg and CFSMount agent timeouts.

Table 1-17 Veritas Storage Foundation Cluster File System fixed issues
(continued)

Fixed issues	Description
2203917	ODM I/O performance bottleneck due to threads waiting in odm_rwsleep_lock() to get lock on odm_iop_table is resolved
2180476	System panic in vx_iupdat_clustblks() due to an unprotected inode getting corrupted.

Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues

This section describes Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation Cluster File System for Oracle RAC: Issues fixed in 5.1 SP1 RP3

There are no fixed issues for Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC) in 5.1 SP1 RP3. The [Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP3](#) are applicable to SFCFSRAC.

Veritas Storage Foundation Cluster File System for Oracle RAC: Issues fixed in 5.1 SP1 RP2

There are no Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues in this release.

Veritas Storage Foundation Cluster File System for Oracle RAC: Issues fixed in 5.1 SP1 RP1

There are no Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues in this release.

Veritas Storage Foundation for Oracle RAC fixed issues

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP3.

Table 1-18 Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2853860	PrivNIC resource state goes to FAULTED state after setting UseVirtualIP to 1.
2853857	when usevirtualIP=1, it overwrites the IP on virtual interface of some nodes of the cluster.
2740150	SFRAC CPI does not set OfflineWaitLimit attribute for CSSD agent resource.
2900969	System panics while accessing ODM device.

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP2 P1.

Table 1-19 Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 P1 fixed issues

Fixed issues	Description
2603511	Database operations can fail on nodes running Oracle RAC 11.2.0.3 and later. The following message is reported in the system logs: ODM ERROR V-41-4-1-105-22 Invalid argument

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP2

[Table 1-20](#) describes the incidents that are fixed in Veritas Storage Foundation for Oracle RAC in 5.1 SP1 RP2.

Table 1-20 Veritas Storage Foundation for Oracle RAC fixed issues

Fixed issues	Description
2429449	The cssd agent explicitly uses hard-coded string "cssd" as resource name.

Table 1-20 Veritas Storage Foundation for Oracle RAC fixed issues (*continued*)

Fixed issues	Description
2390892	Starting the VCSMM driver on two or more nodes in the cluster causes a memory leak in the <code>vcsmm_set_cluster_proto</code> function during memory allocation
2374987	Failed to remove original IP address by PrivNIC and MultiPrivNIC agents during failover/failback operation
2374970	CRSResource agent support for 11gR2

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP1

There are no fixed issues in this release.

Veritas Cluster Server fixed issues

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP3.

Table 1-21 Veritas Cluster Server fixed issues

Fixed issues	Description
2907684	AMF causes process inquiring programs like <code>ps</code> to hang. VCS agents also hang.
2896402	The resource unregistration gets executed with wrong state when running the <code>hagrp -online/-offline</code> or <code>hares -online/-offline</code> command.
2855755	VxFEN might fail to start or online coordination point replacement (OCPR) might fail if a CP server used as a coordination point for the first time and not reachable that time.
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility <code>hagrp</code> gives incorrect output with the <code>"-clear"</code> , <code>"-flush"</code> , <code>"-state"</code> options.

Table 1-21 Veritas Cluster Server fixed issues (*continued*)

Fixed issues	
2831283	System got panic on GAB with below: panic string: BAD TRAP: type=31 rp=2a10d4cf530 addr=28 mmu_fsr=0 occurred in module "gab" due to a NULL pointer dereference.
2818567	LLT ARP flood issue.
2813773	AMF driver panics the box with the message AMF ASSERT panic: FAILED (dev_name).
2811578	Tracking incident for qualification of the VCS Oracle agent (ver. 6.0SP1 and 5.1SP1RP2) on RHEL6 and OEL6.
2804891	lltconfig on boot up core dump and unable to send packets using sendto().
2788059	System did not panic when "PanicSystemOnDGLoss" is set.
2746816	Remove the syslog() call from SIGALRM handler.
2746802	VCS engine should not clear the MigrateQ and TargetCount when failover service group is probed on a system.
2741299	CmdSlave dumped core with SIGSEGV.
2735410	The High Availability Daemon (HAD) core dumps and gets restarted.
2732228	VCS is unable to shut down with the init script.
2731133	When NFSRestart resource is brought offline, it forcefully stops automountd process.
2729867	Global group did not failover to remote site after HAD gets killed and the primary site node crashed.
2729816	Service group failover failure caused by ToQ not getting cleared when OnlineRetryLimit larger than 0.
2728802	Apache agent should work correctly even if Mountpoint for httpd directory is not present on the failover node.
2710892	Node is unable to join fencing cluster after reboot, due to snapshot mismatch.
2692173	The Child service group can be online on the same node with parent group when -nopr is used for an online remote firm dependency.
2684818	If a pure local attribute like PreOnline is specified before SystemList in main.cf then it gets rejected when HAD is started.

Table 1-21 Veritas Cluster Server fixed issues (*continued*)

Fixed issues	
2679251	System panic at dr_mod_cleanup in diskres_udev.
2660011	Restart of an agent moves a critical resource to FAULTED state and hence the group, even if value of ManageFaults attribute is set to NONE at service group level.
2647049	On SLES11, VCS logs do not reflect time zone changes.
2636874	AMF calls VxFS API with spinlock held.
2593173	DiskGroup agent do not detect serial split-brain situation.
2561722	The imf_register entry point failure count gets incremented even when we imf_unregister entry point fails.
2558988	CurrentLimits not getting updated when a node faults.
2531558	graceful shutdown of node should not trigger race condition on peer.
2292294	pfiles, lsof, procfiles, or truss commands hang in the gablogd daemon.
1919382	Mount agent fails to detect the mounted file system with trailing "/".

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP2 P1.

Table 1-22 Veritas Cluster Server 5.1 SP1 RP2 P1 fixed issues

Fixed issues	Description
2708619	SCSI READ BUFFER IOCTL's to device partition fail, vxfen is unable to configure fencing.

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP2 P2.

Table 1-23 Veritas Cluster Server 5.1 SP1 RP2 P2 fixed issues

Fixed issues	Description
2768871	vxfentsthdw should pick up the dmp device path containing whole disk name instead of a partition/slice.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP2

Table 1-24 describes the incidents that are fixed in Veritas Cluster Server in 5.1 SP1 RP2.

Table 1-24 Veritas Cluster Server 5.1 SP1 RP2 fixed issues

Fixed Issues	Description
2382583	Fixed issue with CP Agent where it does not show coordination point information in engine log when CP server is not accessible.
2416842	_had consuming over 99% CPU time. Multiple ha commands are hung in pollsys()
2411653	Add check for MAX message size in GAB
2407755	Application and Netlsnr Agents failing
2407653	In case of forceful unload of AMF module, Module reference count of 'vxfs'/'ext3' should be handled correctly.
2406748	We are able to register already online process for offline monitor with AMF.
2405780	Cable pull test fails when Mii is set to 0
2405391	LLT: The arp ack packet should include the nodename of the node.
2403851	AMF status is showing Module loaded but not configured.
2403782	Sybase agent scripts are setting incorrect path for cat command on linux.
2403633	ContainerInfo attribute should be allowed to be updated even when Group is not completely offline
2400485	Once vxfenconfig -c with mode A has returned EFAULT ("1036 Unable to configure..."), all subsequent runs of vxfenconfig -c with mode B fail with error EBADMSG ("1050 Mismatched modes...").
2400330	whyonlining does not behave as advertised in VCS 5.1SP1.
2399898	hagrp -switch of child group fails in 5.0MP3RP2 and later if 2 or more parent groups online on alternate.
2398807	VCS should be setting a soft limit for file descriptors in /opt/VRTSvcs/bin/vcsenv.
2394176	vxfenswap process hangs, "ps -ef" shows "vxfenconfig -o modify" on one node but not on other. "vxfenswap -a cancel" kills the stuck operation.

Table 1-24 Veritas Cluster Server 5.1 SP1 RP2 fixed issues (*continued*)

Fixed Issues	Description
2386326	cannot configure fencing, vxfenadm prints same Serial Number for all LUNs which have more than 96 bytes of SCSI Inquiry data in page 0x83
2382592	Issue with displaying "ResourceInfo" Attribute of SRDF Resource using hares -display
2382493	Parent service group does not failover in case of online local firm dependency with child service group.
2382463	Had weight(1) is not added if we reach the boundary condition(10000) in System policy in CPS preferred fencing.
2382335	vxfststhdw fails to choose the same fencing disk on two nodes.
2381083	Broadcast address 0.0.0.0 was set by IP-Agent
2372483	SambaServerAgent generated core dumps on FileStore 5.7.
2372072	User core for "hacf"
2366201	Enhanced Fencing to start when a majority of the coordination points are available.
2354932	hacli -cmd' triggers had coredump on 5.1SP1RP1 system
2330980	No notifications about resources should be sent to agents running on nodes already existing in SystemList of Group, when a node is added / deleted to SystemList.
2330045	RemoteGroup resource does not go offline when network fails.
2330041	VCS group dependencies do not online parallel parent group after upgrading SF 5.0MP3 RP2 to SF5.1SP1.
2318334	Oracle needs its database's \$Oracle_home/lib library to be first in LD_LIBRARY_PATH before /usr/lib.
2301731	Panic in amf_lock() due to bad mutex during system shutdown.
2287061	When enabling the amf, cfsmount agent cannot start normally. The basic event registration with AMF driver is failing.
2276622	Cannot configure SCSI-3 fencing using RamSan DMP devices.
2271882	MonitorMethod attribute does not reflect IMF value without setting Listener attribute on Netlsnr Resource.

Table 1-24 Veritas Cluster Server 5.1 SP1 RP2 fixed issues (*continued*)

Fixed Issues	Description
2253441	VCS should setup the right default netmask when NetMask attribute is not set
2528475	support IPMultiNIC/ IPMultiNICB type in preonline_ipc for VCS5.x.
2509515	The resource fails to go offline when Options attribute and class B address is used.
2483964	Monitor for Process resource faults right after online, though the process appears to be running correctly.
2483314	Oracle agent core dumps when large number of oracle instances are running. (Around 50)
2483044	had' crashed with SIGSEGV when asserting against gp->activecount()->gets32GL(nodename) == 0\, in "Resource.C" in check_failover function
2477372	LLT: reduce "lld" CPU consumption by reducing the wakeup calls
2477296	Application service group did not fail over on node panic
2477280	Application resource is not failover when system reboot after Concurrency Violation
2439772	wac resource offline failed after network interruption on SFHA5.1RP2, Solaris 10
2438261	Failed to perform online migration from scsi raw to scsi dmp policy.
2426663	On OCPR from customized mode to scsi3 mode, vxfsend does not terminate
2426572	Persistent resource is reported OFFLINE (not FAULTED) when system is added to group using hagr -modify command
2423990	Application Agent is not working properly when nonexistent user is configured.
2382559	Online Migration fails with the message pl/O fencing does not appear to be configured on nodeq
2382460	Configuring fencing is successful with 3 disks even when single_cp=1 and formatting of warning messages required in vxfsend_A.logo
2382452	Syntax errors while unconfiguring CP server using configure_cps.pl scripto.

Table 1-24 Veritas Cluster Server 5.1 SP1 RP2 fixed issues (*continued*)

Fixed Issues	Description
2367721	Enable selinux permissive / enforcing for Virtual Fire Drill by modifying owner.vfd.
2366701	Query regarding usage of variable in VCS attributes
2366201	Allow Fencing to start when a majority of the coordination points are available.
2364875	Enhancing the Bundled agents to support the RHEL 6 environment.
2330047	VCS share agent hostname comparison is case sensitive.
2511385	Sybase online script marks the database as online before Database has recovered
2439695	VXFEN module gets loaded even though user chooses not to enable VXFEN.
2426572	Persistent resource is reported OFFLINE (not FAULTED) when system is added to group using hagrps -modify command
2411860	Various VCS service group switch failures
2407755	Application and Netlsnr Agents failing
2405514	Panic in amf_lock() due to bad mutex during system shutdown.
2400330	whyonlining does not behave as advertised in VCS 5.1SP1
2382452	Syntax errors while unconfiguring CP server using configure_cps.pl script
2372072	User core for "hacf"
2296172	Failover Policy does not take into consideration AutoFailover = 2 and SystemZones when the nodes within a SystemZone are brought down/rebooted.
2393939	Enhanced Apache agent version parsing to accommodate IBM HTTP server 7.0.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Cluster Server in 5.1SP1RP1 release.

Table 1-25 Veritas Cluster Server 5.1 SP1 RP1 fixed issues

Fixed Issues	Description
1949294	fdsetup can now correctly parse disk names containing characters such as "-".
1949303	fdsetup no longer allows volume that are not part of the RVG, which fixes a possible cause of the RVGSnapshot agent failing.
2011536	Added IMF support for the db2udb agent.
2159991	Fixed an issue with messages in the engine_A.log file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2172181	Fixed an issue with AMF-related messages for the CAVF agent in the engine_A.log file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2179652	The monitor script of the db2udb agent can now handle empty attribute values.
2184205	Fixed an issue with HAD in which the parent service group did not fail over if the parent service group had an online local firm dependency with a child service group.
2194473	HAD no longer dumps core while overriding the static attribute to the resource level.
2205556	Fixed an issue with the offline EP of the DNS agent, which did not remove all A/AAAA records if OffDelRR=1 for multi-home records.
2205563	A clean EP now properly removes resource records when OffDelRR=1.
2205567	Fixed an issue in which having an attribute set to master.vfd caused the DNS agent to fail to query the DNS server.
2208675	There is now a return value check for broadcast ping in NIC/MultiNICA monitor, which fixes one possible cause of the MultiNic resource is going into the FAULTED state in IPv6 with the Link option configuration.
2208901	Fixed an issue with the RVGSnapshot agent.
2209337	Fixed an issue with VCSAPI where the RemoteGroup agent crashed if the VCSAPI log level was set to a non-zero value.
2214539	Fixed an issue in which rebooting a node sometimes set the intentonline of a group to 2, even if the group was online somewhere else. This caused the group to use the autostartlist and not perform a failover.

Table 1-25 Veritas Cluster Server 5.1 SP1 RP1 fixed issues (*continued*)

Fixed Issues	Description
2217446	Fixed an issue that caused the installation of <code>VRTSvcsag</code> to fail.
2218556	Fixed an issue in the <code>cpsadm</code> command in which it sometimes failed if LLT was not installed or configured on a single node cluster.
2218561	Fixed an issue in which <code>MonitorTimeStats</code> incorrectly showed 303 seconds intermittently.
2219955	Fixed an issue in which a split-brain condition occurred even when using VCS Steward.
2220749	Fixed an issue in which the Cluster Manager (Java Console) was not encrypting the <code>DBAPword</code> attribute of the Oracle Agent.
2241419	Fixed an issue in which <code>halogin</code> did not work in a secure environment where the root broker was not a VCS node.

Veritas Storage Foundation for Databases (SFDB) tools fixed issues

This section describes Veritas Storage Foundation for Databases (SFDB) tools fixed issues in 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Storage Foundation for Databases (SFDB) tools fixed issues in 5.1 SP1 RP3.

[Table 1-26](#) describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 5.1 SP1 RP3.

Table 1-26 Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2848204	<code>vxdbd sigsec</code> in <code>stcopy</code> in <code>dbed_ul_print_valist</code> .
2848193	<code>vxdbd coredumps</code> in <code>build_function_header</code> on <code>malloc</code> failure.
2848176	<code>vxdbd</code> memory leak in <code>build_function_header</code> .
1957142	<code>reverse_resync_abort</code> and <code>reverse_resync_commit</code> fails.

Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP2

There are no SFDB fixed issues in 5.1 SP1 RP2.

Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP1

[Table 1-27](#) describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in this release.

Table 1-27 Storage Foundation for Databases fixed issues

Incident	Description
2203917	Process table has been changed to use per-hash-bucket locks, and the number of buckets has been increased from 32 to 256.
2237709	The <code>dbdst_preset_policy</code> command no longer aborts when you specify the volume class as MEDIUM.

Known issues

This section covers the known issues in 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

- [Issues related to installation](#)
- [Veritas Dynamic Multi-pathing known issues](#)
- [Veritas Storage Foundation known issues](#)
- [Veritas Cluster Server known issues](#)
- [Veritas Storage Foundation Cluster File System known issues](#)
- [Veritas Storage Foundation for Oracle RAC known issues](#)

Issues related to installation

This section describes the known issues during installation and upgrade in 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

Incorrect version listed after upgrading (2121881)

When you upgrade from SFCFS 5.1 RP2 to SFCFS 5.1 SP1, the previous SFCFS version is incorrectly listed as 5.1.1.0.

This affects the following products:

- SFCFS
- SFCFSRAC
- SFRAC

Incorrect error messages: error: failed to stat (2120567)

During installation, you may receive errors such as, "error: failed to stat /net: No such file or directory." Ignore this message. You are most likely to see this message on a node that has a mount record of /net/x.x.x.x. However, the /net directory is unavailable at the time of installation.

During product migration the installer overestimates disk space use (2088827)

The installer displays the space that all the product RPM and patches needs. During migration some RPM are already installed and during migration some RPM are removed. This releases disk space. The installer then claims more space than it actually needs.

Workaround: If the disk space is less than that installer claims, but more than actually required. Run the following command:

```
# installer -nospacecheck
```

Error messages in syslog (1630188)

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

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

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

```
Jul  6 10:58:59 swlx62 setroubleshoot: SELinux is preventing the
restorecon from using potentially mislabeled files
```

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

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

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

```
The partitioning on disk /dev/sda is not readable by
The partitioning tool parted, which is used to change the
partition table.
```

```
You can use the partitions on disk /dev/sda as they are.
You can format them and assign mount points to them, but you
cannot add, edit, resize, or remove partitions from that
disk with this tool.
```

Or

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

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

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

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

SFCFS upgrade shows partial upgrade warning(2166810)

When you install 5.1 SFCFS and try to upgrade to SFCFS 5.1 SP1 using the `./installsfcfs` command, you may receive a partial upgrade error message.

Workaround: Use the `./installer -upgrade` command instead of the `./installsfcfs` command.

installrp fails to install 5.1 SP1 RP3 when the root user shell is set to csh (2523643)

The VCS installation fails, if superuser (root) login is using C shell (csh). Currently the installer does not support C shell (`/usr/bin/csh`).

Workaround: Change your superuser (root) shell to `/usr/bin/sh` and retry the installation.

Installation precheck can cause the installer to throw a license package warning (2320279)

If the installation precheck is attempted after another task completes (for example checking the description or requirements) the installer throws the license package warning. The warning reads:

```
VRTSvlic not installed on system_name
```

Workaround:

The warning is due to a software error and can be safely ignored.

While configuring authentication passwords through the Veritas product installer, the double quote character is not accepted (1245237)

The Veritas product installer prompts you to configure authentication passwords when you configure Veritas Cluster Server (VCS) as a secure cluster, or when you configure Symantec Product Authentication Service (AT) in authentication broker (AB) mode. If you use the Veritas product installer to configure authentication passwords, the double quote character (") is not accepted. Even though this special character is accepted by authentication, the installer does not correctly pass the characters through to the nodes.

Workaround: There is no workaround for this issue. When entering authentication passwords, do not use the double quote character (").

EULA changes (2161557)

The locations for all EULAs have changed.

The English EULAs now appear in `/product_dir/EULA/en/product_eula.pdf`

The EULAs for Japanese and Chinese now appear in those language in the following locations:

The Japanese EULAs appear in `/product_dir/EULA/ja/product_eula.pdf`

The Chinese EULAs appear in `/product_dir/EULA/zh/product_eula.pdf`

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

NetBackup 6.5 or older version is installed on a VxFS file system. Before upgrading to Veritas Storage Foundation (SF) 5.1, the user umounts all VxFS file systems including the one which hosts NetBackup binaries (`/usr/opensv`). While upgrading SF 5.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure packages `VRTSspb`, `VRTSat`, and `VRTSicisco`, which causes NetBackup to stop working.

Workaround: Before you umount the VxFS file system which hosts NetBackup, copy the two files `/usr/opensv/netbackup/bin/version` and `/usr/opensv/netbackup/version` to the `/tmp` directory. After you umount the NetBackup file system, manually copy these two version files from the `/tmp` directory to their original path. If the path does not exist, make the same directory path with the command: `mkdir -p /usr/opensv/netbackup/bin` and `mkdir -p /usr/opensv/netbackup/bin`. Run the installer to finish the upgrade process. After upgrade process is done, remove the two version files and their directory paths.

How to recover from systems that are already affected by this issue: Manually install `VRTSspb`, `VRTSat`, and `VRTSicisco` RPM after the upgrade process is done.

The VRTSaclib is deprecated (2032052)

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

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

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

Splitting a cluster that uses server-based fencing is currently not supported. [2110148]

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

However, if you use the installer to reconfigure the Storage Foundation HA, the installer retains the same cluster UUID of `clus1` in both `clus1A` and `clus1B`. If both `clus1A` and `clus1B` use the same CP servers for I/O fencing, then the CP server

allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: None.

Errors observed during partial upgrade of SFHA

While upgrading the VCS packages during an SFHA upgrade from 5.0 MP3 RP2 to 5.1SP1, CPI failed to uninstall the I/O fencing packages (VRTSvxfen, VRTSllt, and VRTSgab). [1779129]

Workaround

Before upgrading SFHA from 5.0 MP3 RP2 to 5.1SP1, you must apply the I/O fencing hotfix 5.0MP3RP2HF2.

Issue with soft links getting deleted in a manual upgrade

While performing a manual upgrade (from 5.1 to 5.1SP1) of the VRTSvlic package, some of the soft links created during your previous installation are deleted. As a result, `vxkeyless` binary is not found in its specified path.

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

For example: `rpm -Uvh --nopreun VRTSvlic-3.02.51.010-0.x86_64.rpm`

Manual upgrade of VRTSvlic package loses keyless product levels (2115662)

If you upgrade the `VRTSvlic` package 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` package.

To manually upgrade the `VRTSvlic` package

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.

```
# rpm -Uvh --nopreun VRTSvlic-3.02.51.010-0.x86_64.rpm
```

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

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

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

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

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

Table 1-28 Whether attributes are configured and required actions that you need to perform during upgrade

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
Configured	May or may not be configured	May or may not be configured	In this case the <code>ifconfig</code> command is used. If RouteOptions is set, attribute value is used to add/delete routes using command <code>route</code> . As the Options attribute is configured, IPv4RouteOptions values are ignored.	No need to configure IPv4RouteOptions .

Table 1-28 Whether attributes are configured and required actions that you need to perform during upgrade (*continued*)

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
Not configured	May or may not be configured	Must be configured	In this case the <code>ip</code> command is used. <code>IPv4RouteOptions</code> must be configured and are used to add/delete routes using the <code>ip route</code> command. As <code>Options</code> attribute is not configured, <code>RouteOptions</code> value is ignored.	Configure <code>IPv4RouteOptions</code> and set the IP of default gateway. The value of this attribute typically resembles: <code>IPv4RouteOptions = "default via gateway_ip"</code> For example: <code>IPv4RouteOptions = "default via 192.168.1.1"</code>

Issues with keyless licensing reminders after upgrading VRTSvlic (2141446)

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

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.
2. Set the product level to `NONE` with the command:

```
# vxkeyless set NONE
```

3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic`:
 - Verify if the key has `VXKEYLESS` feature Enabled using the following command:

```
# vxlicrep -k <license_key> | grep VXKEYLESS
```
 - Delete the key if and only if `VXKEYLESS` feature is Enabled.

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

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

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

SELinux error during installation of VRTSvcsag package

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

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

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

Workaround: Reinstall the SELinux package, if you observe this issue.

5.1 SP1 RP3 configuration title shows as "Veritas Storage Foundation for Oracle RAC 5.1 SP1 PR2 Configure Program" (2908221)

The installer scripts under `/opt/VRTS/install` is still using base version and the configuration title still shows base version after you install 5.1 SP1 RPx patches.

Workaround

There is no workaround. The installer won't be updated when installing RP patches.

CPI failed to install VRTSvxfs rpms on cluster nodes on RHEL 6.3 (2909884)

To install 5.1 SP1 RP3 on RHEL 6.3, you need to install 5.1 SP1 PR2 and then use the `installrp` command to upgrade to 5.1SP1RP3. But because RHEL 6.3 is not supported platform in 5.1 SP1 PR2, CPI failed to install VRTSvxfs rpms on the cluster nodes.

Workaround:

There is no workaround. This issue has no functional impact.

On RHEL6 or later systems with encapsulated root disks, boot failure may occur when you upgrade VxVM 5.1 SP1 RP3 to a later release (2750782)

On RHEL6 or later systems with encapsulated root disks, boot failure may occur when you upgrade VxVM 5.1 SP1 RP3 to a later release. The 5.1 SP1 RP3 uninstall script causes the system to fail to boot up.

Workaround:

To resolve this issue:

- 1 Unroot the encapsulated root disk.

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

- 2 Reboot the system.

- 3 Uninstall VxVM 5.1 SP1 RP3:

- To list the VxVM RPM installed on your system, enter the following command:

```
# rpm -qa | egrep VRTSvxvm
```

- Symantec recommends using the `rpm` command to uninstall the VxVM RPM:

```
# rpm -e --nodeps rpm-version.arch.rpm
```

For example:

```
# rpm -e --nodeps VRTSvxvm-5.1.133.000.x86_64.rpm
```

Or you can use the product installer for uninstallation.

See the *Veritas Storage Foundation and High Availability Installation Guide* on how to uninstall.

4 Reinstall the VxVM RPM you are upgrading to:

Symantec recommends using the `rpm` command to install the VxVM RPM:

```
# rpm -ivh rpm-version.arch.rpm
```

For example:

```
# rpm -ivh VRTSvxvm-5.1.134.000-SP1RP4_RHEL6.x86_64.rpm
```

Or you can use the product installer for installation.

See the *Veritas Storage Foundation and High Availability Installation Guide* on how to install.

Veritas Dynamic Multi-pathing known issues

This section describes the Veritas Dynamic Multi-pathing known issues in 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

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 with the `NODE_SUSPECT` flag. After the connectivity is restored, the restore daemon detects that the path is restored when the restore daemon probes the paths. The restore daemon clears the `NODE_SUSPECT` flag and 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.

Node is not able to join the cluster with high I/O load on the array with Veritas Cluster Server (2124595)

When the array has a high I/O load, the DMP database exchange between master node and joining node takes a longer time. This situation results in VCS resource online timeout, and then VCS stops the join operation.

Workaround:

Increase the online timeout value for the HA resource to 600 seconds. The default value is 300 seconds.

To set the OnlineTimeout attribute for the HA resource type CVMCluster

- 1 Make the VCS configuration to be read/write:

```
# haconf -makerw
```

- 2 Change the OnlineTimeout attribute value of CVMCluster:

```
# hatype -modify CVMCluster OnlineTimeout 600
```

- 3 Display the current value of OnlineTimeout attribute of CVMCluster:

```
# hatype -display CVMCluster -attribute OnlineTimeout
```

- 4 Save and close the VCS configuration:

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

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

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

Work around:

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

Device discovery failure during VxVM startup in SUSE (2122771)

When the system boots up, some devices are not displayed in the `vxdisk list` output. This issue occurs if the `vold` daemon does the device discovery before the operating system (OS) completes its device discovery. Therefore, the `vold` daemon may miss some devices.

Work-around:

Configure the `vxvm-startup` script to wait until the operating system discovery is completed before starting `vold`. In the `/etc/vx/vxvm-startup` file, set `DEV_DISCOVER_DELAY` to the expected device discovery time taken by the OS. By default, `DEV_DISCOVER_DELAY` is set to 0.

You must reboot the system before this configuration applies. To discover the missed devices, run the `vxdisk scandisks` command or the `vxctl enable` command.

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

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

```
# vxddmctl settune dmp_fast_recovery=off
```

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

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

Upgrading the Linux kernel when the root volume is under DMP control

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

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

To update the kernel on a RHEL5 system

- 1 Update kernel with the `rpm` command.

```
# rpm -ivh kernel_rpm
```

- 2 Turn on the `dmp_native_support` tunable:

```
# vxddmctl settune dmp_native_support=on
```

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

- 3 Reboot.

On SLES10 or SLES11

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

To update the kernel on a SLES10 or SLES11 system

- 1 Turn off DMP native support

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

- 2 Reboot the system.

- 3 Upgrade kernel using the `rpm` command

```
# rpm -ivh kernel_rpm
```

- 4 Turn on DMP native support.

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

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

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

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

A system can hang or panic in a SAN boot environment due to udev device removal after loss of connectivity to some paths on SLES 11 (2219626)

The issue may occur with NetApp LUNs in ALUA mode, with a SAN boot configuration. When a device fails with a `dev_loss_tmo` error, the operating system (OS) device files are removed by `udev`. After this removal, a system can hang or panic due to I/O disruption to the boot device. To avoid this issue, use the following workaround.

Workaround

To update the kernel and create the new rules file

- 1 Save the existing rules file.

```
# mkdir /savefiles
# cd /etc/udev/rules.d/
# mv 40-VxVM.rules /savefiles
```

- 2 Download and upgrade to kernel 2.6.27.45-0.1.1 or later from Novell.
- 3 Create the file `/etc/udev/rules.d/40-rport.rules` with the following content line:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports",
ACTION=="add",RUN+="/bin/sh -c 'echo 20 >
/sys/class/fc_remote_ports/%k/fast_io_fail_tmo;echo 864000
>/sys/class/fc_remote_ports/%k/dev_loss_tmo'"
```

- 4 Reboot the system.
- 5 If new LUNs are dynamically assigned to the host, run the following command:

```
# udevadm trigger --action=add --subsystem-match=fc_remote_ports
```

The DMP tunables `partitionsize` and `iopolicy` got overwritten after upgrade from SFRAC 51 SP1 PR2 to SFRAC 5.1 SP1 RP3 (2935183)

After upgrading SF stack from SFRAC 51 SP1 PR2 to SFRAC 5.1 SP1 RP3 using the `responsefile` on RHEL6, the following DMP tunables are reset to default:

- `partitionsize`
- `iopolicy`

Workaround:

After upgrade, use the `vxddmpadm setattr enclosure` CLI to set the tunables back to the desired values.

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

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

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

Workaround:

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

```
obtain_device_list_from_udev = 0
```

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

```
# vxddm padm settune dmp_native_support=on
```

Excluded DMP devices shows up in vxdisk list for 3PAR array (2562817)

Exclude the DMP devices for 3PAR array and disable one of the ports from the switch side and then reboot the system. Once the machine comes up, the excluded DMP devices show up in the `vxdisk list` output.

Workaround:

To exclude the devices,

1. Run the `udevadm trigger` command to generate all the hwwpaths.
2. Run the `vxdisk scandisks` command to rescan all the devices.

Veritas Storage Foundation known issues

This section describes the Veritas Storage Foundation known issues in 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

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

Veritas Storage Foundation known issues

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

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

After you have installed a Storage Foundation product, some VM process may hang during the configuration phase.

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

Oracle 11gR2 is not supported with the managed host client included in this release (2568403)

The Veritas Operations Manager (VOM) managed host client version 3.xxx that is included in this release is not supported with Oracle 11gR2.

Workaround: After installing this release, upgrade the managed host client to version 4.0 RU1 or later.

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

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

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

Or

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

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

db2exp may frequently dump core (1854459)

If a host is configured to an SFM central server with DB2 version 9.x, then the command-line interface db2exp may frequently dump core.

Workaround: There is a hotfix patch available for this issue. Contact Symantec Technical Support for the hotfix patch.

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

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

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

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

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

```
DBI1070I Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

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

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

To communicate in a dual-stack environment

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

```
127.0.0.1 swlx20-v6
```

Or

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

where `127.0.0.1` is the IPv4 loopback address.

where `swlx20-v6` and `swlx20-v6.punipv6.com` is the IPv6 hostname.

Oracle 11gR1 may not work on pure IPv6 environment (1819585)

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

Tools like `dbca` may hang during database creation.

Workaround: There is no workaround for this, as Oracle 11gR1 does not fully support pure IPv6 environment. Oracle 11gR2 release may work on a pure IPv6 environment, but it has not been tested or verified.

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

After upgrading SF stack from 5.0 MP3 SP1 RP2 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the SFM GUI.

Workaround:

To resolve this known issue

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

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

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

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

```
Cannot open file /etc/vx/vxdba/rac11g1/.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'rac11g1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdg ERROR V-5-1-582 Disk group SNAP_rac11dg1: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dg1.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dg1 failed.
```

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

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

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

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

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

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

To resolve this issue

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

A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880622)

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

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

Veritas Volume Manager known issues

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

Node join can lead to hang if an upgrade of the cluster protocol version is in progress (2103567)

If you attempt to join a node to the cluster while Cluster Volume Manager (CVM) is upgrading the cluster protocol version, the system may hang. This issue occurs if the node is attempting to join the cluster after you issue the `vxctl upgrade` command to upgrade the CVM cluster.

Work-around:

Avoid joining a new node to the cluster until the CVM cluster upgrade is completed.

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

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

```
VxVM vxencap INFO V-5-2-5327 Missing file: /boot/grub/device.map
```

Workaround:

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

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

Post encapsulation of the root disk, system comes back up after first reboot unencapsulated (2119038)

In some cases, after encapsulating the root disk and rebooting the system, it may come up without completing the encapsulation. This happens because the vxvm-reconfig startup script is unable to complete the encapsulation process.

Workaround

Reboot the system or run the following command.

```
# service vxvm-reconfig start
```

This will reboot the system and complete the remaining stages of encapsulation.

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

To support iSCSI devices, Veritas Volume Manager (VxVM) does not start non-root devices until runlevel2. The boot process expects all local (non-NFS) mount points in the `/etc/fstab` file to be present at boot time. To avoid boot failures, all VxVM entries in the `/etc/fstab` file must have the `_netdev` attribute, and must not have the `fsck` required flag set. These attributes enable VxVM to defer mounting of VxVM devices until after VxVM has started.

vxrestored daemon fails to restore disabled paths (1663167)

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

Workaround:

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

To enable the `mpt_disable_hotplug_remove` tunable

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

```
options mptsas mpt_disable_hotplug_remove=0
```

- 2 Rebuild the `initrd` image:

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

- 3 Reboot the system.

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

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

Workaround:

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

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

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

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

Workaround:

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

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

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

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

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

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

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

Workaround:

To recover from this situation

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

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

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

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

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

Root disk encapsulation issue (1603309)

Encapsulation of root disk will fail if it has been assigned a customized name with vxdkmpadm(1M) command. If you wish to encapsulate the root disk, make sure that you have not assigned a customized name to its corresponding DMP node. See vxdkmpadm(1M) and the section "Setting customized names for DMP nodes" on page 173 for details.

VxVM starts before OS device scan is done (1635274)

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

Workaround:

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

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.

The layout operation fails when there are too many disks in the disk group. (2015135)

The attempted layout operation on a disk group containing approximately more than 300 LUNs or disks may fail with the following error:

```
Cannot setup space
```

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

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

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

Workaround: There is no workaround for this issue.

Removing a volume from a thin LUN in an alternate boot disk group triggers disk reclamation (2080609)

If you remove a volume from an alternate boot disk group on a thin LUN, this operation triggers thin reclamation, which may remove information required for the disk to be bootable. This issue does not affect the current boot disk, since VxVM avoids performing a reclaim on disks under the bootdg.

Workaround: If you remove a volume or plex from an alternate boot disk group with the `vxedit` command, specify the `-n` option to avoid triggering thin reclamation. For example:

```
# vxedit -g diskgroup -rfn rm volumename
```

VxVM vxdg V-5-1-16063 is returned from 'vxdg rmdisk' when attempting to perform storage reclamation on Hitachi AMS 2500 array

See <http://www.symantec.com/docs/TECH162709> for more information.

On 5.1SP1RP2, changing path attributes by using 'vxdatapadm' gives an error message "Failed to save path information persistently" (2433012)

On Linux system with "selinuxenabled" set to 1, setting path attributes using "vxdatapadm setattr path *pathname* pathtype=*attr*", CLI will set the path attribute but these changes may not remain persistent across reboots. Following error displays if CLI fails to save the changes persistently:

```
VxVM vxdatapadm ERROR V-5-1-14526 Failed to save path information  
persistently
```

Workaround: Reset these attributes using vxdatapadm CLI, if system is rebooted.

Node not able to join when recovery is in progress (2165829)

When a node leaves a cluster, there is an associated volume recovery for the leaving node. It is not allowed to join any node to the cluster during the recovery because the nodeid of the leaving node cannot be assigned to the joining node.

Workaround:

Retry node join after the recovery is completed.

Panic or vxconfigd hang after installing vxvm package on RHEL6(2917137)

On RHEL6, because of the change in the interface for inter-module communication, VRTSaslapm whose version is earlier than 5.1.100.500 with VxVM 5.1 SP1 RP3 causes machine to panic or vxconfigd hang.

Workaround:

To upgrade to 5.1 SP1 RP3 on RHEL6

- 1 Uninstall VRTSaslapm if it is already installed.
- 2 Upgrade VxVM to 5.1 SP1 RP3 .
- 3 Install the latest VRTSaslapm package whose version is no ealier than 5.1.100.500.

On SLES11, upgrading kernel on encapsulated bootdisk does not work as documented (2920815)

Kernel upgrade on the encapsulated root disk does not work properly and may cause system boot failure. While upgrading kernel on the encapsulated root disk you may experience any of the following issues:

- Incorrect entry for VxVM(Veritas Volume Manager) root disk in menu.lst
- System boot failure - VxVM_root_backup entry is not removed from menu.lst when un-encapsulating root disk

Workaround:

Perform following steps on the machine with the encapsulated root disk to upgrade kernel:

- 1 Unroot the encapsulated root disk:

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

- 2 Upgrade the kernel:

```
# rpm -Uvh Kernel-upgrade version
```

- 3 Reboot to boot into the upgraded kernel.

- 4 Re-encapsulated the root disk:

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

If vxconfigd is under heavy load, “vxassist settag” may make volume tagging information inconsistent (2484764)

If there are a lot of VxVM operations running, vxconfigd is under heavy load. If you execute the vxassist settag operations when vxconfigd is under stress, these operations will succeed, but the volume tagging information may be inconsistent. In such cases, you will not be able to use the tag for the further operations for that particular volume. And if you run the vxassist listtag operation, it will fail with error:

```
Inconsistent tag information found on disk
```

Workaround:

There is no workaround for this issue.

vxconfigd dumps core on all the nodes in Campus Cluster setup (2937600)

Campus Cluster Scenario (two sites A and B, with 2 nodes in each site):

1. Disabled site A storage from all the four nodes and also shutdown site A nodes.
2. Enabled site A storage and activated site A nodes.
3. Site B nodes panic.

After the reboot of the nodes in site A, when nodes try to join the cluster, vxconfigd dumps core.

Workaround:

There is no workaround for this issue right now.

You are unable to encapsulate or unroot the root disk if the device name changed (2222447)

You are unable to encapsulate the root disk or unroot the root disk if the root disk's device name changed over reboot. The device name changes if you perform disk operations such as adding or removing LUNs and then rebooting the machine. This may lead to boot failure.

Workaround:

There is no workaround for this issue.

Veritas File System known issues

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

VxFS read ahead can cause stalled I/O on all write operations (1965647)

Changing the `read_ahead` parameter can lead to frozen I/O. Under heavy load, the system can take several minutes to recover from this state.

Workaround: There is no workaround for this issue.

VxFS module loading fails when freevxfs module is loaded (1736305)

The following module loading error can occur during RPM installation if the `freevxfs` module is loaded:

```
Error in loading module "vxfs". See documentation.
```

```
ERROR: No appropriate VxFS drivers found that can be loaded.  
See VxFS documentation for the list of supported platforms.
```

Workaround: Ensure that the `freevxfs` module is not loaded before installing the `VRTSvxfs` RPM. The following command shows if the `freevxfs` module is loaded:

```
# lsmod | grep freevxfs
```

If the `freevxfs` module is loaded, unload the module:

```
# rmmod freevxfs
```

A mount can become busy after being used for NFS advisory locking (1508386)

If you export a VxFS file system using NFS and you perform file locking from the NFS client, the file system can become unable to be unmounted. In this case, the `umount` command fails with the `EBUSY` error.

Workaround: Force unmount the file system:

```
# vxumount -o force /mount1
```

Storage Checkpoints can exceed the quota limit (2102201)

Under some circumstances, Storage Checkpoints can exceed the quota limit set by the `fsckptadm setquotalimit` command. This issue can arise if all of the following conditions are met:

- The Storage Checkpoint quota has been enabled.
- The Storage Checkpoint quota is not exceeded.
- A file content modification operation, including removing a file, needs to push some or all blocks of the file to the Storage Checkpoint.
- Number of blocks that need to be pushed to the Storage Checkpoint is enough to exceed Storage Checkpoint quota hard limit.

Workaround: There is no workaround for this issue.

vxfsconvert can only convert file systems that are less than 1 TB (2108929)

The `vxfsconvert` command can only convert file systems that are less than 1 TB. If the file system is greater than 1 TB, the `vxfsconvert` command fails with the "Out of Buffer cache" error.

Possible write performance degradation with VxFS local mounts (1837394)

Some applications that allocate large files without explicit preallocation may exhibit reduced performance with the VxFS 5.1 release and later releases compared to the VxFS 5.0 MP3 release due to a change in the default setting for the tunable `max_seqio_extent_size`. One such application is DB2. Hosting DB2 data on a single file system extent maximizes the potential for sequential pre-fetch processing. When DB2 detects an application performing sequential reads against database data, DB2 begins to read ahead and pre-stage data in cache using efficient sequential physical I/Os. If a file contains many extents, then pre-fetch processing is continually interrupted, nullifying the benefits. A larger `max_seqio_extent_size` value reduces the number of extents for DB2 data when adding a data file into a tablespace without explicit preallocation.

The `max_seqio_extent_size` tunable controls the amount of space that VxFS automatically preallocates to files that are allocated by sequential writes. Prior to the 5.0 MP3 release, the default setting for this tunable was 2048 file system blocks. In the 5.0 MP3 release, the default was changed to the number of file system blocks equaling 1 GB. In the 5.1 release, the default value was restored to the original 2048 blocks.

The default value of `max_seqio_extent_size` was increased in 5.0 MP3 to increase the chance that VxFS will allocate the space for large files contiguously, which tends to reduce fragmentation and increase application performance. There are two separate benefits to having a larger `max_seqio_extent_size` value:

- Initial allocation of the file is faster, since VxFS can allocate the file in larger chunks, which is more efficient.
- Later application access to the file is also faster, since accessing less fragmented files is also more efficient.

In the 5.1 release, the default value was changed back to its earlier setting because the larger 5.0 MP3 value can lead to applications experiencing "no space left on device" (ENOSPC) errors if the file system is close to being full and all remaining space is preallocated to files. VxFS attempts to reclaim any unused preallocated space if the space is needed to satisfy other allocation requests, but the current implementation can fail to reclaim such space in some situations.

Workaround: If your workload has lower performance with the VxFS 5.1 release and you believe that the above change could be the reason, you can use the `vxtunefs` command to increase this tunable to see if performance improves.

To restore the benefits of the higher tunable value

- 1 Increase the tunable back to the 5.0 MP3 value, which is 1 GB divided by the file system block size.

Increasing this tunable also increases the chance that an application may get a spurious ENOSPC error as described above, so change this tunable only for file systems that have plenty of free space.
- 2 Shut down any applications that are accessing any large files that were created using the smaller tunable setting.
- 3 Copy those large files to new files, which will be allocated using the higher tunable setting.
- 4 Rename the new files back to the original names.
- 5 Restart any applications that were shut down earlier.

NFS issues with VxFS checkpoint (2027492)

NFS clients mounting VxFS checkpoints that are NFS-exported by SFCFS cluster nodes using a Virtual IP may receive the following error message upon Virtual IP failover:

```
Stale NFS file handle
```

This is a result of major numbers of VxFS checkpoints not necessarily being the same on all SFCFS cluster nodes.

There is no workaround at this time.

Veritas Volume Replicator known issues

This section describes the known issues in this release of Veritas Volume Replicator (VVR).

vradmin syncvol command compatibility with IPv6 addresses (2075307)

The `vradmin syncvol` command does not work with the compressed form of IPv6 addresses. In IPv6 environments, if you run the `vradmin syncvol` command and identify the target host using compressed form of the IPv6 address, the command fails with following error message:

```
# vradmin -s -full syncvol vol1 fe80::221:5eff:fe49:ad10:dg1:vol1
VxVM VVR vradmin ERROR V-5-52-420 Incorrect format for syncvol.
```

Also, if you run the `vradmin addsec` command and you specify the Secondary host using the compressed IPv6 address, the `vradmin syncvol` command also fails – even if you specify the target as `hostname`.

Workaround: When you use the `vradmin addsec` and `vradmin syncvol` commands, do not specify compressed IPv6 addresses; instead, use hostnames.

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

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

Workaround:

To resolve this issue

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

Interrupting the vradmin syncvol command may leave volumes open (2063307)

Interrupting the `vradmin syncvol` command may leave volumes on the Secondary site in an open state.

Workaround: On the Secondary site, restart the `in.vxrsyncd` daemon. Enter the following:

```
# /etc/init.d/vxrsyncd.sh stop

# /etc/init.d/vxrsyncd.sh start
```

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

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

`ElectPrimary` command to elect the new Primary or if the previous `ElectPrimary` command did not complete successfully.

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround: The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

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

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

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

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

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

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

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

Workaround: The following procedure resolves this issue.

To resolve this issue

- 1 Edit the `/etc/hosts` file to specify the correct IPv6 address.
- 2 Restart the `vradmind` daemon on all VVR hosts:

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

Veritas product 5.0MP3 Rolling Patch 2 required for replication between 5.0 MP3 and 5.1 SP1 (1800600)

In order to replicate between Primary sites running Veritas product 5.0 MP3 and Secondary sites running Veritas product 5.1 SP1, or vice versa, you must install the Veritas product 5.0MP3 Rolling Patch 2 on the nodes using 5.0MP3. This patch resolves several outstanding issues for replicating between versions.

While vradmind changeip is running, vradmind may temporarily lose heart beats (2162625)

This issue occurs when you use the `vradmind changeip` command to change the host name or IP address set in the Primary and Secondary RLINKs. While the `vradmind changeip` command runs, `vradmind` may temporarily lose heart beats, and the command terminates with an error message.

Workaround:**To resolve this issue**

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

- 2 Restart `vradmind` to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following:

vxassist relay layout removes the DCM (2162522)

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

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

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

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

Workaround:**To resize layered volumes that are associated to an RVG**

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

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

- 3 Stop the affected RVG. Enter the following:

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

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

```
# vxvol -g diskgroup dis vol
```
- 5 Resize the volumes. In this example, the volume is increased to 10 GB. Enter the following:

```
# vxassist -g diskgroup growto vol 10G
```
- 6 Associate the data volumes to the RVG. Enter the following:

```
# vxvol -g diskgroup assoc rvg vol
```
- 7 Start the RVG. Enter the following:

```
# vxrvg -g diskgroup start rvg
```
- 8 Resume or start the applications.

Cannot relayout data volumes in an RVG from concat to striped-mirror (2162537)

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

Workaround:

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

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

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

```
# vxrvg -g diskgroup stop rvg
```
- 4 Disassociate the volumes from the RVG. Enter the following:

```
# vxvol -g diskgroup dis vol
```
- 5 Relayout the volumes to striped-mirror. Enter the following:

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon

Issue: After upgrading VVR to an IPv6-only environment in 5.1 release, `vradmin` commands may not work when a colon is specified in the RVG, data volume(s) and/or SRL name. It is also possible that after upgrading VVR to an IPv6-only environment, `vradmin createpri` may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

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

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

When VVR is used for replicating shared disk groups in an SFCFS or SFRAC environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg` commands:

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

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

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

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

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

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

If using VEA to create a replicated data set fails, messages display corrupt strings in the Japanese locale (1726499, 1377599)

When using VEA to create a replicated data set, because the volumes do not have a DCM log on all nodes, the message window displays corrupt strings and unlocalized error messages.

Workaround: There is no workaround for this issue.

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

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

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

Workaround:

To restore `vradmin` functionality after a master switch operation

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

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

- 2 Re-enter the command that failed.

Veritas Storage Foundation for Databases (SFDB) tools known issues

The following are known issues in this release of Veritas Storage Foundation products.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0.x to 5.1SP1 (2184482)

The `sfua_rept_migrate` command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 5.1SP1. The error message is:

When upgrading from Veritas product version 5.0 to Veritas product 5.1SP1 the `S*vxdbms3` startup script is renamed to `NO_S*vxdbms3`. The `S*vxdbms3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbms3` startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround

Before running `sfua_rept_migrate`, rename the startup script `NO_S*vxdbms3` to `S*vxdbms3`.

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

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

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

```
dbed_vmclonedb started at 2009-08-26 11:32:16
Editing remote_login_passwordfile in initclone2.ora.
Altering instance_name parameter in initclone2.ora.
Altering instance_number parameter in initclone2.ora.
Altering thread parameter in initclone2.ora.
SFORA dbed_vmclonedb ERROR V-81-4918 Database clone2 has not been
correctly recovered.
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.20569/recover.log.
```

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

Workaround

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

To resolve this known issue

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

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

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

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

4 Start the database.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0.x to 5.1SP1

The `sfua_rept_migrate` command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 5.1SP1. The error message is:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

When using SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 5.1SP1 the `S*vxdbms3` startup script is renamed to `NO_S*vxdbms3`. The `S*vxdbms3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbms3` startup script and gives the above error message.

Workaround:

Before running `sfua_rept_migrate`, rename the startup script `NO_S*vxdbms3` to `S*vxdbms3`.

Database fails over during Flashsnap operations (1469310)

In an Veritas product environment, if the database fails over during Flashsnap operations such as the `dbed_vmsnap -o resync` command and various error messages appear. This issue occurs because Flashsnap commands do not create a VCS resource for the SNAP disk group. As such, when the database fails over, only the primary disk group is moved to another node.

Workaround:

There is no workaround for this issue.

The error messages depend on the timing of the database failover. To fix the problem, you need to bring the FlashSnap state to `SNAP_READY`. Depending on the failure, you may have to use base VxVM commands to reattach mirrors. After mirrors are attached, you need to wait until the mirrors are in `SNAPDONE` state. Re-validate the snapplan again.

Reattach command failure in a multiple disk group environment (1840672)

In a multiple disk group environment, if the snapshot operation fails then `dbed_vmsnap` fails to reattach all the volumes. This operation must be performed as root user.

Workaround

In case the reattach operation fails, use the following steps to reattach the volumes.

To reattach volumes in a multiple disk group environment if the snapshot operation fails

- 1 Join the snapshot disk groups to primary diskgroups. The snapshot disk group name is a concatenation of “SNAPSHOT_DG_PREFIX” parameter value in `snapplan` and primary disk group name. Use the following command to join the disk groups:

```
# vxvg join snapshot_disk_group_name
           primary_disk_group_name
```

- 2 Start all the volumes in primary disk group.

```
# vxvol -g primary_disk_group_name startall
```

- 3 Reattach the snapshot volumes with primary volumes. The snapshot volume names is a concatenation of “SNAPSHOT_VOL_PREFIX” parameter value in `snapplan` and primary volume name. Use the following command to reattach the volumes.

```
# vxsnap -g primary_disk_group_name reattach snapshot_volume_name
source=primary_volume_name
```

Repeat this step for all the volumes.

Clone command fails if archive entry is spread on multiple lines (1764885)

If you have a `log_archive_dest_1` in single line in the `init.ora` file, then `dbed_vmclonedb` will work but `dbed_vmcloneb` will fail if you put in multiple lines for `log_archive_dest_1`

Workaround:

There is no workaround for this issue.

The `/opt/VRTSdbed/bin/dbdst_obj_move` command may fail with error messages on 10gRAC env (2927308)

The `dbdst_obj_move` command may fail with FSPPADM error:

```
/opt/VRTS/bin/dbdst_obj_move -S $ORACLE_SID -H $ORACLE_HOME \  
-v -t tab_part4 -s 0 -e 10 -c SLOW  
FSPPADM err : Not enough space
```

```
SFORA dst_obj_adm ERROR V-81-6414 Internal Error at fsppadm_err
```

This error can be caused by the old filesystem layout version. To use the `dbdst_obj_move` command, you need filesystem layout 8 or higher.

Workaround:

Upgrade the filesystem layout to version 8.

To upgrade the filesystem layout to version 8:

- 1 Use the following command to check the filesystem layout version:

```
# /opt/VRTS/bin/fstyp -v /dev/vx/dsk/oradatadg/oradatavol \  
| grep version
```

- 2 Use the following command to upgrade the filesystem layout to version 8:

```
# /opt/VRTS/bin/vxupgrade -n 8 /oradata
```

sometimes clone fails with error on sles11 with oracle 11.2.0.3 (2907542)

On SLES 11 with oracle 11.2.0.3, the clone operation fails sometimes with error:

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

This is because of the following oracle bugs which are currently open:

Bug 14102418: DATABASE DOESNT START DUE TO ORA-1513

Bug 14036835: SEEING ORA-01513 INTERMITTENTLY

This problem occurs on the RHEL and SLES platforms. It affects Oracle version 11.2.0.2 and version 11.2.0.3. There is currently no solution for this issue.

Workaround

Retry the cloning operation a few times until it succeeds.

The `dbed_vmclonedb -o recoverdb` command may fail with error messages (2928666)

The `dbed_vmclonedb -o recoverdb` command may fail with following error messages:

```
SFORA dbed_vmclonedb ERROR V-81-4882 An error occurred while reconfiguring  
Oracle instance 'ckptc1n'.
```



```
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is  
at /tmp/dbed_vmclonedb.50528402/startup.log.
```

Also check the `startup.log` file if it contains the following information:

```
./home/oracle>cat /tmp/dbed_vmclonedb.47251692/startup.log  
ORA-16019: cannot use LOG_ARCHIVE_DEST_1 with LOG_ARCHIVE_DEST or  
LOG_ARCHIVE_DUPLEX_DEST
```

The error occurs if the log archive destination parameter of Oracle RAC instances is configured individually using a similar command such as :

```
SQL> alter system set log_archive_dest_1 = 'location=/arch MANDATORY' SID='RACID';
```

Workaround:

Use the following command to set `log_archive_dest_1` for all of the instances at once:

```
SQL> alter system set log_archive_dest_1 = 'location=/arch MANDATORY' SID=*;
```

Veritas Cluster Server known issues

This section describes the Veritas Cluster Server known issues in 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

- [Operational issues for VCS](#)
- [Issues related to the VCS engine](#)
- [Issues related to the bundled agents](#)
- [Issues related to the VCS database agents](#)
- [Issues related to the agent framework](#)
- [Issues related to VCS in Japanese locales](#)
- [Issues related to global clusters](#)
- [Issues related to LLT](#)
- [Issues related to GAB](#)
- [Issues related to I/O fencing](#)
- [Issues related to Symantec Product Authentication Service with VCS](#)
- [Issues related to Veritas Cluster Server agents for Veritas Volume Replicator](#)
- [Issues related to IMF](#)

NFS cluster I/O fails when storage is disabled

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.

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

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

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

ha command does not work when VCS_DOMAIN or VCS_DOMAINTYPE is set with remote broker (2272352)

When VCS_DOMAIN or VCS_DOMAINTYPE is set with remote broker, ha command does not work.

Workaround:

- 1 Set VCS_REMOTE_BROKER to the remote AB:

```
# export VCS_REMOTE_BROKER=remote_broker
```

- 2 Set VCS_DOMAIN and VCS_DOMAINTYPE:

```
# export VCS_DOMAINTYPE=ldap
# export VCS_DOMAIN=ldap_domain_name
```

- 3 Run halogin:

```
# halogin ldap_user
```

Provide password when prompted.

- 4 Unset VCS_DOMAIN and VCS_DOMAINTYPE:

```
# unset VCS_DOMAINTYPE
# unset VCS_DOMAIN
```

- 5 Run any ha command. The command should run fine if the ldap_user has the correct privileges

hacmd -display G or A or O or E or S or C dumps core (2415578)

VCS ha commands do not work when object names are single character long.

VCS agent for Oracle: 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.

had stopped before vxfen startup (2429272)

VCS stops with error VCS CRITICAL V-16-1-10031 VxFEN driver not configured. VCS Stopping. Manually restart VCS after configuring fencing.

If UseFense is set to SCSI3, upon starting VCS checks if the VxFen driver is configured. If VxFen driver is configured to use CPS or if there is a pre-existing split brain, the driver takes a long time to complete configuration. VCS periodically queries the driver until the driver is configured or exits after 90 seconds.

Workaround

VCS must be restarted manually after the VxFen driver is configured.

Application Agent does not handle a case when user is root, envfile is set and shell is csh. (2490299)

The Application Agent uses the `system` command to execute the Start/Stop/Monitor/Clean Programs for 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.

Forcefully un-configuring AMF does not change the monitor method of agent to TRADITIONAL (2564376)

If you forcefully unconfigure AMF when a script based agent is registered with AMF, the monitor method of the agent is not changed to `TRADITIONAL`. It remains `IMF`.

Workaround: Restarting the agent will resolve the issue.

Forcefully un-configuring AMF causes the engine log to be flooded with error messages (2535690)

If you forcefully unconfigure AMF when a script based agent is registered with AMF, the `getnotification` thread continuously polls and displays error messages in the engine log.

Workaround: Restarting the agent will resolve the issue.

NFS resource goes offline on its own and errors out when restarted (2490415)

If multiple agent processes are running because an agent process is restarted multiple times by `_had`, then only one of the agent process survives and other agent processes go offline on its own. Even though the agent process is running, `_had` does not recognize it and hence does not perform any resource operations.

Workaround: Kill the agent process to recover from this situation. Refer to the engine log for further actions (if required) to restart the agent.

HAD dumps core when `hagrp -clear` is executed on a group in OFFLINE|FAULTED state and a resource in the fault path is waiting to go online (2536404)

This issue occurs if you have a resource dependency, such as `r1 -> r2 -> r3`. While resources `r2` and `r3` are online and you initiate bringing resource `r1` online, before the `OnlineTimeout` occurs, resources `r2` and `r3` suffer a fault. Resource `r2` faults first, and then `r3` faults. After the fault of both resources is detected, the group is becomes in an `OFFLINE|FAULTED` state and resource `r1` is stuck waiting to become online. If you execute the `hagrp -clear` command to clear the fault, then HAD dumps core on all nodes due to assertion failure.

Workaround: Flush the pending online operation using the `hagrp -clear` command before clearing the fault.

The `hares -display` command fails if the resource is part of a global service group (2358600)

The `hares -display` command incorrectly processes the response received from the `had` process. Due to the processing error, `hares -display` does not show the resource details.

Workaround: Use the `-localclus` or `-clus` option with `hares -display`.

Excessive delay messages in one-node configurations of Storage Foundation High Availability (SFHA) (2486415)

The GAB error, "Excessive delay between successive calls to GAB heartbeat" appear in the engine log while running a single node or standalone VCS cluster where GAB is disabled.

GAB heartbeat log messages are logged as informational when there is delay between heartbeats (HAD being stuck). When HAD runs in -onenode, GAB does not need to be enabled. When HAD is running in -onenode, for self-checking purposes, HAD simulates heartbeat with an internal component of HAD itself. These log messages are logged because of delay in simulated heartbeats.

Workaround: Log messages are for informational purpose only. When HAD is running in -onenode, no action is needed on excessive delay between heartbeats.

Agent does not retry resource registration with IMF to the value specified in RegisterRetryLimit key of IMF attributes (2411882)

Agent maintains internal count for each resource to track the number of times a resource registration is attempted with IMF. This count gets incremented if any attempts to register a resource with IMF fails. At present, any failure in un-registration of resource with IMF, also increments this internal count. This means that if resource un-registration fails, next time agent may not try as many resource registrations with IMF as specified in RegisterRetryLimit. This issue is also observed for multiple resources.

Workaround: Increase the value of RegisterRetryLimit if un-registration and registration of one or more resources is failing frequently. After increasing the RegisterRetryLimit, if resource registration continues to fail, report this problem to Symantec.

DBED commands such as sfua_db_config and dbed_ckptcreate fail

DBED commands such as `sfua_db_config` and `dbed_ckptcreate` fail with the following error message:

```
VXDBA_PRODUCT exec_remote ERROR V-81-7700 Can not connect to the vxdbd. It might be down. Check the status and restart it if it is not up.
```

However, the output of the `ps` command shows that `vxdbd` is running:

```
# ps -ef | grep vxdbd  
root 14572 1 0 05:07:35 - 0:13 /opt/VRTSdbed/common/bin/vxdbd -d
```

Workaround: Set the environment variable EAT_HOME_DIR to the value /opt/VRTSdbed/eat before running the command:

```
# EAT_HOME_DIR=/opt/VRTSdbed/eat  
  
# export EAT_HOME_DIR
```

installation of Storage Foundation 5.1SP1RP2 fails on SLES10 SP4

Installation of Storage Foundation 5.1SP1RP2 fails on SLES10 SP4 with the error:

The following required OS rpms were not found on host:

```
compat-libstdc++-5.0.7-22.2.x86_64
```

Workaround: Install the compat-libstdc++-32bit RPM package. It is available on the SLES10 SP3 installation media.

Oracle agent incorrectly reports the global resource as online when the resource inside the local zone is online and the Sid's are same (2561563)

Oracle agent incorrectly reports the resource configured for Oracle instance running in global container as online, if the resource configured for Oracle instance running in local container also has same value for Sid attribute and the instance in local container is online.

The above issue is also applicable for ASMInst and Netlsnr agents.

For Netlsnr agent the above issue appears when the Home and listener attributes of the resources running in global and local container are same.

The issue does not appear for Oracle and ASMInst agents when multiple local containers have resources configured with the same value of Sid attribute.

The issue does not appear for Netlsnr agent when multiple local containers have resources configured with the same value of Home and listener attributes.

NFS mount does not work whenever the NFS daemon is started on port other than 2049 (2477799)

On RHEL 6.0 and 6.1, if the paths are exported by a NFS service (nfsd) which is running on any port other than the default (2049), the mount command fails to

mount the exported paths. If NFS service is running on the default port (2049), then you are able to mount the exported paths successfully.

This issue occurs only on RHEL 6.0 and 6.1. It does not exist on RHEL 5.6, 5.7 and SLES 10, 11.

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 may not all 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.

Operational issues for VCS

Issues with configuration of resource values

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 with bunker replay

When ClusterFailoverPolicy is set to Auto and the AppGroup is configured only on some nodes of the primary cluster, global cluster immediately detects any system fault at the primary site and quickly fails over the AppGroup to the remote site. VVR might take longer to detect the fault at the primary site and to complete its configuration changes to reflect the fault.

This causes the RVGPrimary online at the failover site to fail and the following message is displayed:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname
could not be imported on bunker host hostname. Operation
failed with error 256 and message VxVM
VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote server
```

```
unreachable...
```

```
Timestamp VCS ERROR V-16-2-13066 (hostname) Agent is calling  
clean for resource(RVGPrimary) because the resource  
is not up even after online completed.
```

Resolution: To ensure that global clustering successfully initiates a bunker replay, Symantec recommends that you set the value of the `OnlineRetryLimit` attribute to a non-zero value for `RVGPrimary` resource when the primary site has a bunker configured.

LVM SG transition fails in all paths disabled status

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

Workaround: You need to restore at least one path.

The CmdServer process may not start in IPv6 environments in secure clusters

In an IPv6 environment on secure clusters, the `CmdServer` process may not start. In addition, security may not function correctly. If it does not start on a particular node, modify that node's `/etc/hosts` file so that the `localhost` resolves to `::1`.

Workaround: In the `/etc/hosts` file, add the following:

```
::1          localhost
```

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

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

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

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

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

Issues related to the VCS engine

Systems with multiple CPUs and copious memory shut-down time may exceed the ShutdownTimeout attribute

The time taken by the system to go down may exceed the default value of the ShutdownTimeout attribute for systems that have a large numbers of CPUs and memory. [1472734]

Workaround: Increase the value of the ShutdownTimeout attribute based on your configuration.

VCS Engine logs messages when it eventually connects to a remote cluster

Description: In a global cluster, if a local cluster fails to connect with a remote cluster in the first attempt but succeeds eventually, then you may see the following warning messages in the engine logs. [2110108]

```
VCS WARNING V-16-1-10510 IpmHandle: pen Bind Failed.  
unable to bind to source address 10.209.125.125. errno = 67
```

Workaround: There is currently no workaround for this issue. This issue has no impact on any functionality.

Agent framework can reject `hares -action` command

When a probed resource is disabled and later enabled then, the agent framework can reject `hares -action` command till the agent successfully monitors the resource.

New nodes get added to SystemList and AutoStartList attributes of ClusterService even if AutoAddSystemToCSG is disabled

The AutoAddSystemToCSG attribute determines whether the newly joined or added systems in a cluster become part of the SystemList of the ClusterService service group if the service group is configured. The value 1 (default) indicates that the new systems are added to SystemList of ClusterService.

AutoAddSystemToCSG has an impact only when you execute the `hasys -add` command or when a new node joins the cluster. [2159139]

However, when you use the installer to add a new node to the cluster, the installer modifies the SystemList and AutoStartList attributes irrespective of whether AutoAddSystemToCSG is enabled or disabled. The installer adds the new system to the SystemList and AutoStartList. To add nodes, the installer uses the following commands that are not affected by the value of AutoAddSystemToCSG:

```
# hagr -modify ClusterService SystemList -add newnode n
# hagr -modify ClusterService AutoStartList -add newnode
```

Workaround

The installer will be modified in future to prevent automatic addition of nodes to SystemList and AutoStartList.

As a workaround, use the following commands to remove the nodes from the SystemList and AutoStartList:

```
# hagr -modify ClusterService SystemList -delete newnode
# hagr -modify ClusterService AutoStartList -delete newnode
```

The hacf -cmdtocf command generates a broken main.cf file (1728738)

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

Workaround:

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

Issues related to the bundled agents

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

LVM Logical Volume gets auto activated during the I/O path failure. This causes the VCS agent to report "Concurrency Violation" errors, and make the resource groups offline/online temporarily. This is due to the behavior of Native LVM. [2140342]

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

The DNS resource enters an UNKNOWN state after upgrade from VCS 5.0 MP3 to 5.1

The DNS resource enters an UNKNOWN state after an upgrade from VCS 5.0 MP3 to 5.1.

Workaround: Manually update the value of the DNS agent's ResRecord attribute using ha commands and verify if DNS resource can be online.

```
# haconf -makerw
# hares -modify resname ResRecord -add alias hostname
# haconf -dump -makero
```

Bring the corresponding resources or service groups online if needed.

Issues related to the VCS database agents

VCS agent for Oracle: Health check monitoring does not work with Oracle 10.2.0.4

The health check monitoring in Oracle agent does not work with Oracle 10.2.0.4 due to incompatibility of the health check APIs provided by Oracle. [2101570]

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

VCS agent for Oracle: Make sure that the ohasd has an entry in the init scripts

Make sure that the ohasd process has an entry in the init scripts so that when the process is killed or the machine is rebooted, this automatically restarts the process.[1985093]

Workaround: Respawn of ohasd process. Add the ohasd process in the `/etc/inittab` file to ensure that this process is automatically restarted when killed or the machine is rebooted.

VCS agent for Oracle: Intentional Offline does not work

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

Health check monitoring for Oracle agent on SLES11 platform

Oracle agent does not support health check monitoring on SLES11 platform. [1938167]

The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups

The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups.

Workaround:

Have a copy of the pfile/spfile in the default `$GRID_HOME/dbs` directory to make sure that this would be picked up during the ASM Instance startup.

VCS agent for ASM: Health check monitoring is not supported for ASMInst agent

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

Issues related to the agent framework

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

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.

Agent may fail to heartbeat under heavy load (2073018)

Description: An agent may fail to heartbeat with the VCS engine under heavy load.

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

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

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

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

In rare cases, the agent framework does not detect if all service threads hang inside a C entry point. In this case it may not cancel them successfully. [1511211]

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

The ArgListValues attribute values for dependent resources may not populate correctly when a target resource is deleted and re-added

For resource attributes, deleting a resource prevents a dependent attribute's value from refreshing in the dependent resource's value.

For example, you have resource (*rD*), which depends on a resource's attribute value (*rT:Attr_rt*). When you delete the target resource (*rT*), and re-add it (*rT*), the dependent resource (*rD*) does not get the correct value for the attribute (*Attr_rt*). [1539927]

Workaround: Set the value of the reference attribute (*target_res_name*) to an empty string.

```
# hares -modify rD target_res_name ""
```

Where *rD* is the name of the dependent resource, and *target_res_name* is the name of the reference attribute that contains the name of the target resource.

Set the value of the reference attribute (*target_res_name*) to the name of the target resource (*rT*).

```
# hares -modify rD target_res_name rT
```

Agent performance and heartbeat issues

Depending on the system capacity and the number of resources configured under VCS, the agent may not get enough CPU cycles to function properly. This can prevent the agent from producing a heartbeat synchronously with the engine. If you notice poor agent performance and an agent's inability to heartbeat to the engine, check for the following symptoms.

Navigate to `/var/VRTSvcs/diag/agents/` and look for files that resemble:

```
FFDC_AGFWMMain_729_agent_type.log  FFDC_AGFWTimer_729_agent_type.log  core
FFDC_AGFWSvc_729_agent_type.log   agent_typeAgent_stack_729.txt
```

Where *agent_type* is the type of agent, for example Application or FileOnOff. If you find these files, perform the next step.

Navigate to `/var/VRTSvcs/log/` and check the `engine_*.log` file for messages that resemble:

```
2009/10/06 15:31:58 VCS WARNING V-16-1-10023 Agent agent_type
not sending alive messages since Tue Oct 06 15:29:27 2009
2009/10/06 15:31:58 VCS NOTICE V-16-1-53026 Agent agent_type
ipm connection still valid
2009/10/06 15:31:58 VCS NOTICE V-16-1-53030 Termination request sent to
agent_type agent process with pid 729
```

Workaround: If you see that both of the above criteria are true, increase the value of the `AgentReplyTimeout` attribute value. (Up to 300 seconds or as necessary.)
[1853285]

Issues related to VCS in Japanese locales

This section covers the issues that apply to VCS 5.1 in a Japanese locale.

The gcoconfig script displays error messages in English

The gcoconfig script incorrectly displays English error messages. [1416136]

Issues related to global clusters

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

When the WAC process runs in secure mode on one site, and the other site does not use secure mode, the engine log file on the secure site gets logs every five seconds. [1539646]

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

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

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. (2107386)

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.

Issues related to LLT

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

LLT port stats sometimes shows recvcnt larger than recvbytes

With each received packet, LLT increments the following variables:

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

Both these variables are integers. With constant traffic, recvbytes hits and rolls over MAX_INT quickly. This can cause the value of recvbytes to be less than the value of recvcnt. [1788315]

This does not impact the LLT functionality.

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

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This

occurs in some corner cases even though a port is not even registered on the peer node. [1809827]

Issues related to GAB

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

Trace messages from the gablogd daemon on the console for RHEL5 Update 5 or later

On RHEL5 Update 5 or later, the `gablogd` daemon prints informational and trace messages similar to the following [2139883]:

```
INFO: task gablogd:22812 blocked for more than 120 seconds.
"echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
gablogd      D ffff81004100daa0      0 22812      1      23106 22809 (NOTLB)
  ffff810faf539e38 00000000000000082 0000000000000084c 0000000000000001
  ffff810faf539de8 00000000000000007 ffff810fc2a130c0 ffff810138ee8100
  000019f130082599 00000000000018572 ffff810fc2a132a8 00000001f76c3d63
Call Trace:
 [<ffffffff88ee3690>] :gab:gab_linux_sv_wait+0x53/0x68
 [<ffffffff8008e68d>] default_wake_function+0x0/0xe
 [<ffffffff88ecd4c8>] :gab:gab_daemonlog+0xae1/0xc52
 [<ffffffff88ee326c>] :gab:gab_linux_ioctl1+0x10e/0x1a3
 [<ffffffff88ee331d>] :gab:gab_linux_compat_ioctl1+0x1c/0x20
 [<ffffffff800fbe53>] compat_sys_ioctl1+0xc5/0x2b2
 [<ffffffff8006249d>] sysenter_do_call+0x1e/0x76
```

Workaround: As the operating system message indicates, set the following:

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

Issues related to I/O fencing

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

All nodes in a sub-cluster panic if the node that races for I/O fencing panics

At the time of a network partition the lowest node in each sub-cluster races for the coordination points on behalf of that sub-cluster. If the lowest node is unable to contact a majority of the coordination points or the lowest node itself unexpectedly panics during the race, then all the nodes in that sub-cluster will panic. [1965954]

Preferred fencing does not work as expected for large clusters in certain cases

If you have configured system-based or group-based preferred fencing policy, preferred fencing does not work if all the following cases are true:

- The fencing setup uses customized mode with one or more CP servers.
- The application cluster has more than eight nodes.
- The node weight for a single node (say galaxy with node id 0) is more than the sum total of node weights for the rest of the nodes.
- A network fault occurs and the cluster partitions into two with the single node (galaxy) on one part and the rest of the nodes on the other part.

Under such circumstances, for group-based preferred fencing, the single node panics even though more high priority services are online on that node. For system-based preferred fencing, the single node panics even though more weight is assigned to the node. [2161816]

See the *Veritas product Administrator's Guide* for more information on preferred fencing.

Server-based I/O fencing fails to start after configuration on nodes with different locale settings

On each (application cluster) node, the vxfen module retrieves and stores the list of the UUIDs of coordination points. When different nodes have different locale settings, the list of UUIDs on one (application) node does not match with that of the other (application) nodes. Hence, I/O fencing does not start after configuration. [2112742]

Workaround: Start I/O fencing after fixing the locale settings to use the same values on all the (application) cluster nodes.

Reconfiguring Storage Foundation HA with I/O fencing fails if you use the same CP servers

When you reconfigure an application cluster that uses server-based I/O fencing (customized fencing mode), the installer does not remove the application cluster information from the CP servers before the reconfiguration. As a result, if you reconfigure the application cluster and choose to configure I/O fencing in customized mode using the same CP servers, then reconfiguration of server-based fencing for the application cluster fails. [2076240]

Workaround: Manually remove the application cluster information from the CP servers after you reconfigure Storage Foundation HA but before you reconfigure server-based I/O fencing for the application cluster.

See the *Veritas Cluster Server Administrator's Guide* for instructions to remove the application cluster information from the CP servers.

CP server cannot bind to multiple IPs (2085941)

Coordination point server (CP server) binds only to a single virtual IP and listens on the same. Application clusters cannot access the CP server if it fails to establish connection to this virtual IP. Therefore, if the connection fails because of the subnet in which the virtual IP of the CP server exists, you cannot access the CP server even if there is another subnet through which the client can connect to the CP server over a different IP.

Resolution: No known resolution for this issue.

Issues related to Symantec Product Authentication Service with VCS

This section covers the known issues related to Symantec Product Authentication Service (AT) in this release.

The vcsat and cpsat commands may appear to be hung

The following commands may appear to be hung when you invoke them from the command shell:

- `/opt/VRTScps/bin/cpsat`
- `/opt/VRTSvc/bin/vcsat`

This issue occurs when the command requires some user interaction. [1841185]

Workaround:

- To fix the issue for vcsat, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTSvc
# /opt/VRTSvc/bin/vssatvcs command_line_argument
# unset EAT_HOME_DIR
```

- To fix the issue for cpsat, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTScps
# /opt/VRTScps/bin/vssatcps command_line_argument
# unset EAT_HOME_DIR
```

Issues related to Veritas Cluster Server agents for Veritas Volume Replicator

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 SP1 RP3 release.

RVGPrimary online script does not function correctly (1949293)

The RVGPrimary online script does not function correctly.

Issues related to IMF

Failure messages of resource un-registration with IMF appear in agent or engine logs after performing online or offline operations on the resource (2909184)

When a resource is registered with IMF for monitoring, any online or offline operation triggers un-registration of the resource from IMF. During such operations, agent may record an error message in the agent or engine logs stating that the un-registration failed. This issue is also observed for multiple resources.

Workaround:

There is no workaround. These failure messages are false positives and no resolution is required. Agent registers resources with IMF again after some time.

Issues related to AMF

Issues with the amfstat output (2926158)

The `amfstat` output displays an extra column in the Registered Reapers list and the `amfstat -n` output displays the header twice.

Workaround:

This issue does not have any effect on the functionality of AMF. It has been fixed in VCS 6.0 and onwards.

Veritas Storage Foundation Cluster File System known issues

This section describes the Veritas Storage Foundation Cluster File System known issues in 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

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 checkpoints. This issue can be seen using the `fsckptadm` command:

```
# fsckptadm getquotalimit /mnt1
Filesystem    hardlimit    softlimit    usage    action_flag
/mnt1         10000        10000        18446744073709551614
```

This could cause writes to checkpoints to fail. It could also trigger the removal of removable checkpoints.

Workaround

If this occurs, disabling and re-enabling file set quotas causes VxFS to recalculate the number of blocks used by checkpoints:

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

installer –makeresponsefile detects the wrong product (2044525)

If you generate a response file to upgrade SFCFS or SFCFSHA using the `./installer -makeresponsefile` command, and then choose **G** (Upgrade a Product) option, the installer detects it as SFCFS RAC.

You can safely ignore that the installer detects it as SFCFS RAC.

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 Oracle RAC known issues

This section describes the Veritas Storage Foundation for Oracle RAC known issues in 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

Oracle RAC issues

This section lists the known issues in Oracle RAC.

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

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 Grid Infrastructure installation may fail with the Veritas product installer

When you run the `installsfrac -configure` command to install Oracle Grid Infrastructure for Oracle RAC 11g Release 2, the installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

Workaround: Export the `OUI_ARGS` environment variable, before you run the Veritas product installation program:

```
export OUI_ARGS=--ignoreInternalDriverError
```

For more information, see the Oracle Metalink document: 970166.1

The Oracle Interface Configuration (OIFCFG) command fails if private IP address is plumbed on a virtual interface

In Oracle RAC 11.2.0.1 installations, plumbing the private IP addresses on virtual interfaces cause the Oracle Interface Configuration (`oifcfg`) command to fail with the following error:

```
# oifcfg delif -global eth3/192.168.1.0:cluster_interconnect
PRIF-26: Error in update the profiles in the cluster
```

This issue may be observed during failover of IP addresses by the MultiPrivNIC agent. This is because the MultiPrivNIC agent plumbs the IP address on a virtual interface during failover. [2131345]

Oracle RAC 11g Release 2 installation on SLES 11 may fail if the RPM `libcap1-1.10-6.10.x86_64` is not present

Oracle RAC 11g Release 2 installations on SLES 11 fail if the RPM `libcap1-1.10-6.10.x86_64` is not present.

Workaround: Before installing Oracle RAC 11g Release 2, verify that the RPM `libcap1-1.10-6.10.x86_64.rpm` is present on the system:

```
# rpm -qa | grep libcap
libcap1-1.10-6.10
libcap2-2.11-2.15
```

If the RPM is not installed, install it using the instructions in the Oracle Metalink document: 952051.1

Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software

The Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software. If the failure indicates that the OCR and vote device locations are not shared, ignore the message.

Changing the Veritas agent for Oracle error handling

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

For a description of the actions:

See the *Symantec High Availability Agent for Oracle Installation and Configuration Guide*.

Currently, the file specifies the `NOFAILOVER` action for the following Oracle errors: `ORA-00061`, `ORA-02726`, `ORA-6108`, `ORA-06114`

The `NOFAILOVER` action means that the agent sets the state of the resource to `OFFLINE` and freezes the service group. If you want to change this behavior, you can stop the agent, edit `oraerror.dat`, and change the `NOFAILOVER` action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

Veritas product issues

This section lists the known issues in Veritas product for this release.

Incorrect ownership assigned to the parent directory of ORACLE_BASE causes Oracle Clusterware/Grid Infrastructure installations to fail

When you use the Veritas product installation program to install Oracle Clusterware/Grid Infrastructure, the ownership of the parent directory of ORACLE_BASE/GRID_BASE that is created by the installation program is incorrectly set to root. This causes the Oracle Universal Installer to display errors when it creates the `oraInventory` directory as the `oraInventory` directory must be created on behalf of the `oracle` user (Oracle RAC 10g Release 2/Oracle RAC 11g Release 1) or `grid` user (Oracle RAC 11g Release 2).

Workaround:

1. Log into each node in the cluster as the root user.
2. Perform the following operations:
 - If you have not yet installed Oracle Clusterware/Grid Infrastructure, create the directory and set the correct ownership as follows before you invoke the installation program:

```
# mkdir -p oracle_base
# chown user_name:oraInventory_group_name
  oracle_base/..
```

where:

`oracle_base` is the name of the Oracle base directory.

`user_name` is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).

`oraInventory_group_name` is the name of the `oraInventory` group.

Complete the other preparatory tasks before you invoke the installation program. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

- If you faced this issue during the installation of Oracle Clusterware/Grid Infrastructure, open another terminal session, and modify the ownership of the directory on all nodes in the cluster as follows:

```
# chown user_name:oraInventory_group_name
  oracle_base/..
```

where:

`oracle_base` is the name of the Oracle base directory.

`user_name` is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).

`oraInventory_group_name` is the name of the `oraInventory` group.

Return to the former session and proceed with the installation.

Deporting issues with shared disk groups

If you manually deport a shared disk group, the CVMVolDg agent does not automatically reimport it as a shared disk group. You must manually reimport it as a shared disk group.

SCSI reservation errors during bootup

If you reboot a node of an SF Oracle RAC cluster, SCSI reservation errors may be observed during bootup. [255515]

For example:

```
Nov 23 13:18:28 galaxy kernel: scsi3 (0,0,6) : RESERVATION CONFLICT
```

This message is printed for each disk that is a member of any shared disk group which is protected by SCSI-3 I/O fencing. The message may be safely ignored.

Network interfaces change their names after restart

On SUSE systems, network interfaces change their names after restarting even with `HOTPLUG_PCI_QUEUE_NIC_EVENTS=yes` and `MANDATORY_DEVICES="..."` set.

Workaround: Set the interface names in the following files:

```
SLES 10          /etc/udev/rules.d/30-net_persistent_names.rules
```

```
SLES 11          /etc/udev/rules.d/70-persistent-net.rules
```

Kernel warning messages when Veritas modules load

For SLES 10, a warning message resembling the following may be displayed in the console or the system log when Veritas modules are loaded into the kernel.

```
Warning: module not supported by Novell, setting U taint flag.  
module license 'Proprietary. Send bug reports to  
support@symantec.com' taints kernel.
```

These warning messages are displayed because the Veritas modules are proprietary. They can safely be ignored.

DBED features are not integrated with GCO

DBED features are not integrated with Global Cluster Option (GCO). After GCO migration, be aware that DBED features will not be functional. [1241070]

Issue with format of the last 8-bit number in private IP addresses

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. [1164506]

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.

When master node loses access to complete storage, detached sites remain in RECOVER state even after reattaching and recovering the sites

In a campus cluster environment, if the master node loses access to complete storage, all but one of the sites is detached and the DCO volumes may get detached if the `dgfailpolicy` is set to `dgdisable`. If the detached sites are reattached and recovered, the site still remains in RECOVER state. [1828142]

Workaround: Change the status of the site as described in the following procedure to resolve the issue.

To change the status of the site

- 1 Log onto the CVM master node.
- 2 Reattach the detached sites:

```
# vxdbg -g dg_name reattachsite site_name
```

The site remains in RECOVER state.

- 3 Restore DCO volumes by unpreparing and preparing the volumes.

Unprepare the volumes:

```
# vxsnap -g dg_name -f unprepare vol_name
```

Prepare the volumes:

```
# vxsnap -g dg_name prepare vol_name drl=on
```


4 Reattach the detached sites:

```
# vxdg -g dg_name reattachsite site_name
```

5 Verify that the state of the detached sites is now ACTIVE:

```
# vxprint
```

Failure to set the MTU (Maximum Transmission Unit) size in LLT over UDP environments causes issues with the PrivNIC/MultiPrivNIC agents (2557144)

If the MTU size field is not set explicitly when you configure the PrivNIC/MultiPrivNIC agents in an LLT over UDP environment, the agents may fail to plumb the private IP addresses during their operations or may configure incorrect MTU size on the LLT interfaces.

The agents use the `lltstat -l` command to retrieve MTU size information for LLT interfaces. In an LLT over UDP environment, the command retrieves 8192 as the MTU size. When the PrivNIC/MultiPrivNIC agents use this size information to plumb the IP addresses, the operation may fail causing the agents to fault. However, even if the plumbing operation succeeds, the incorrect MTU configuration may still cause issues in the cluster later.

Workaround:

To update the PrivNIC/MultiPrivNIC resource configuration in an LLT over UDP environment

- 1 Retrieve the MTU size of the network interfaces configured under PrivNIC/MultiPrivNIC agents:

```
For Linux: # ifconfig eth1
```

- 2 Set the MTU attribute for the PrivNIC/MultiPrivNIC resource:

```
# haconf -makerw
```

Run the following command for all the network interfaces configured under PrivNIC/MultiPrivNIC agents:

```
# hares -modify resource_name MTU -add interface_name mtu_size
```

Where:

resource_name is the name of the PrivNIC/MultiPrivNIC resource

interface_name is the name of the network interface for which the MTU size is set

mtu_size is the MTU size retrieved in step 1.

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

The MultiPrivNIC agent configures IP addresses on virtual interfaces though the UseVirtualIP parameter is set 0 (2566142)

The MultiPrivNIC agent configures the IP addresses on virtual interfaces though the UseVirtualIP parameter is set to 0 while configuring the MultiPrivNIC resource in the `main.cf` file. However, no functional issues are observed in the cluster.

11.2.0.1 crsd.bin Fails in clsCclClscWait (2933706)

The CRS processes `crsd.bin`, `evmd.bin`, `ohasd.bin` are dead.

Workaround:

Apply the 11814167 patch for this bug. For more information see Metalink ID 1326008.1 on the ORACLE support site.

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

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

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

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

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

3. Remove the `vxfsckd-pid` file:

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

4. Bring the `vxfsckd` resource online:

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

vxdumpasm cannot create the "auto:ASM" TYPE for the ASM disk (2944387)

The `vxdumpasm` command cannot create the "auto:ASM" TYPE for the ASM disk.

Workaround:

There is no workaround for this issue.

Software limitations

This section covers the software limitations of this release.

Veritas Volume Manager software limitations

This is the Veritas Volume Manager software limitation in the 5.1 SP1 RP3 release:

To install or upgrade to 5.1 SP1 RP3 on RHEL6, the version of `VRTSaslapm` should be later than 5.1.100.500.

Veritas Cluster Server software limitations

This is the Veritas Cluster Server software limitations in the 5.1 SP1 RP3 release.

Sometimes parent group will not restart with OnlineRetryLimit set (2279845)

With OnlineRetryLimit set, a child service group that has recovered from a fault will not be able to restart its faulted parent service group, if the parent service group's fault is detected before the child service group's fault.

This scenario may happen typically when:

- You have single system in systemlist of child and parent groups.
- The group has mix of faulted persistent and non-persistent resources.

Workaround:

If the group with OnlineRetryLimit does not restart or failover, manually clear the fault and run the `online` command.

The Ldom agent do not gracefully detect guest domain migration

If the guest domain managed by the LDom agent is migrated to another host, it will not be gracefully detected by the agent and the resource will be marked FAULTED.

Workaround:

Upgrade to VCS version 6.0 or higher for the Ldom agent capability to gracefully responds to migration.

List of RPMs

This section lists the RPMs for 5.1 SP1 RP3.

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

```
./installrp -listpatches
```

Table 1-29 RPMs for Red Hat Enterprise Linux 5

Name	Version	Arch	Size in bytes
VRTSamf	5.1.133.000	x86_64	769098

Table 1-29 RPMs for Red Hat Enterprise Linux 5 (*continued*)

Name	Version	Arch	Size in bytes
VRTScavf	5.1.133.000	i386	180574
VRTScps	5.1.132.000	i686	15383939
VRTSdbac	5.1.133.000	x86_64	933299
VRTSdbed	5.1.133.000	i686	6167644
VRTSfssdk	5.1.133.000	x86_64	350647
VRTSgab	5.1.133.000	x86_64	939317
VRTSglm	5.1.132.100	x86_64	106348
VRTSllt	5.1.133.000	x86_64	908240
VRTSvvmconv	5.1.133.000	i686	70899
VRTSob	3.4.312	i686	33422315
VRTSodm	5.1.133.000	x86_64	236778
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcsc	5.1.133.000	i686	46761530
VRTSvcscag	5.1.133.000	i686	686279
VRTSvcscdr	5.1.133.000	x86_64	198994
VRTSvcsea	5.1.133.000	i686	3826483
VRTSvcxfen	5.1.133.000	x86_64	452160
VRTSvcxfs	5.1.133.000	x86_64	9368732
VRTSvcxvm	5.1.133.000	x86_64	27197137

Table 1-30 RPMs for Red Hat Enterprise Linux 6

Name	Version	Arch	Size in bytes
VRTSamf	5.1.133.000	x86_64	538256
VRTSaslapm	5.1.100.501	x86_64	153408
VRTScavf	5.1.133.000	i386	164368
VRTScps	5.1.132.000	x86_64	12056456

Table 1-30 RPMs for Red Hat Enterprise Linux 6 (*continued*)

Name	Version	Arch	Size in bytes
VRTSdbac	5.1.133.000	x86_64	795496
VRTSdbed	5.1.133.000	i686	6167644
VRTSfssdk	5.1.133.000	x86_64	291152
VRTSgab	5.1.133.000	x86_64	777276
VRTSglm	5.1.132.100	x86_64	111748
VRTSllt	5.1.133.000	x86_64	740984
VRTSvmconv	5.1.133.000	i686	64484
VRTSob	3.4.312	i686	33422315
VRTSodm	5.1.133.000	x86_64	260880
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcsc	5.1.133.000	i686	30699960
VRTSvcscag	5.1.133.000	i686	531260
VRTSvcscdr	5.1.133.000	x86_64	225484
VRTSvcsea	5.1.133.000	i686	3399016
VRTSvcxfen	5.1.133.000	x86_64	905428
VRTSvcxfs	5.1.133.000	x86_64	6913416
VRTSvcxvm	5.1.133.000	x86_64	17106960

Table 1-31 RPMs for SUSE Linux Enterprise Server 10

Name	Version	Arch	Size in bytes
VRTSsamf	5.1.133.000	x86_64	4153682
VRTScavf	5.1.133.000	i386	153902
VRTScps	5.1.132.000	i686	14430499
VRTSdbac	5.1.133.000	x86_64	1764429
VRTSdbed	5.1.133.000	i586	5519549
VRTSfssdk	5.1.133.000	x86_64	307658

Table 1-31 RPMs for SUSE Linux Enterprise Server 10 (*continued*)

Name	Version	Arch	Size in bytes
VRTSgab	5.1.133.000	x86_64	2062384
VRTSglm	5.1.132.100	x86_64	177667
VRTSllt	5.1.133.000	x86_64	1553923
VRTSvvmconv	5.1.133.000	i586	61947
VRTSob	3.4.312	i686	33422315
VRTSodm	5.1.133.000	x86_64	358834
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcs	5.1.133.000	i586	42265072
VRTSvcsag	5.1.133.000	i586	548750
VRTSvcsdr	5.1.133.000	x86_64	404894
VRTSvcssea	5.1.133.000	i586	3606834
VRTSvxfen	5.1.133.000	x86_64	2219877
VRTSvxfs	5.1.133.000	x86_64	11315766
VRTSvxvm	5.1.133.000	x86_64	25770910

Table 1-32 RPMs for SUSE Linux Enterprise Server 11

Name	Version	Arch	Size in bytes
VRTSamf	5.1.133.000	x86_64	1159244
VRTScavf	5.1.133.000	i386	157651
VRTScps	5.1.132.000	i686	11569516
VRTSdbac	5.1.133.000	x86_64	1162246
VRTSdbed	5.1.133.000	i586	5519549
VRTSfssdk	5.1.133.000	x86_64	268304
VRTSgab	5.1.133.000	x86_64	1298761
VRTSglm	5.1.132.100	x86_64	166658
VRTSllt	5.1.133.000	x86_64	1156053

Table 1-32 RPMs for SUSE Linux Enterprise Server 11 (*continued*)

Name	Version	Arch	Size in bytes
VRTSsvmconv	5.1.133.000	i586	63030
VRTSob	3.4.312	i686	33422315
VRTSodm	5.1.133.000	x86_64	275615
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcsc	5.1.133.000	i686	30446411
VRTSvcscag	5.1.133.000	i686	497174
VRTSvcscdr	5.1.133.000	x86_64	329748
VRTSvcscsea	5.1.133.000	i686	2637345
VRTSvcscfen	5.1.133.000	x86_64	1404452
VRTSvcscfs	5.1.133.000	x86_64	6182484
VRTSvcscvsm	5.1.133.000	x86_64	17227970

Downloading the 5.1 SP1 RP3 archive

The patches that are included in the 5.1 SP1 RP3 release are available for download from the Symantec website. After downloading the 5.1 SP1 RP3 rolling patch, use `gunzip` and `tar` commands to uncompress and extract it.

For the 5.1 SP1 RP3 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

<http://www.symantec.com/docs/TECH75506>

Installing the products for the first time

This chapter includes the following topics:

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

Installing the Veritas software using the script-based installer

This section describes how to install a 5.1 SP1 RP3 Veritas Storage Foundation and High Availability Solutions product for the first time on a host. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 5.1 SP1 *Installation Guide* and *Release Notes* for your product for more information.

See “[Upgrading to 5.1 SP1 RP3](#)” on page 134.

To install the Veritas software for the first time

- 1 Download Storage Foundation and High Availability Solutions 5.1 SP1 from <http://fileConnect.symantec.com>.
- 2 Extract the tar ball into a directory called `/tmp/sfha51sp1`.
- 3 Check <http://sort.symantec.com/patches> to see if there are any patches available for the 5.1 SP1 Installer. Download applicable P-patches and extract them to the `/tmp` directory.

- 4 Change the directory to `/tmp/sfha51sp1`:

```
# cd /tmp/sfha51sp1
```
- 5 Run the installer to install SFHA 5.1 SP1. See the Installation Guide for instructions on installing the 5.1 SP1 version of this product.

```
# ./installer -require complete_path_to_SP1_installer_patch
```
- 6 Download SFHA 5.1 SP1 RP3 from <http://sort.symantec.com/patches>.
- 7 Extract it to a directory called `/tmp/sfha51sp1rp3`.
- 8 Check <http://sort.symantec.com/patches> to see if there are patches available for the 5.1 SP1 RP3 installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 9 Change the directory to `/tmp/sfha51sp1rp3`:

```
# cd /tmp/ sfha51sp1rp3
```
- 10 Invoke the `installrp` script to install 5.1 SP1 RP3:

```
# installrp -require complete_path_to_SP1RP3_installer_patch
```

See “[About the installrp script](#)” on page 12.
- 11 If you did not configure the product after the 5.1 SP1 installation, the installer prompts you to configure the product during RP installation. If you do not want to configure the product now, answer **n** when prompted. To configure the product in the future, run the product installation script from the 5.1 SP1 installation media or from `/opt/VRTS/install` directory with the `-configure` option

Installing Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 5.1 SP1 RP3 using the Web-based installer. For detailed instructions on how to install 5.1 SP1 using the Web-based installer, follow the procedures in the 5.1 SP1 Installation Guide and Release Notes for your products.

See “[Upgrading to 5.1 SP1 RP3](#)” on page 134.

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

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

```
# ./webinstaller start
```

The webinstaller script displays a URL.

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

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

Obtaining a security exception on Mozilla Firefox

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

To obtain a security exception

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

Installing 5.1 SP1 RP3 with the Veritas Web-based installer

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

To install Veritas product

- 1 The 5.1 SP1 version of the Veritas product must be installed before upgrading to 5.1 SP1 RP3.

See “[Prerequisites for upgrading to 5.1 SP1 RP3](#)” on page 133.

- 2 On the **Select a task and product** page, select **Install 5.1 SP1 RP3** from the **Task** drop-down list, and click **Next**.
- 3 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
- 4 You have the option to let the installer configure SSH or RSH communications between the systems. If you choose to allow this configuration, select the shell and provide the root passwords for each system.
- 5 After the validation completes successfully, click **Next** to install 5.1 SP1 RP3 patches on the selected system.
- 6 The installer prompts you to configure the cluster.

If you select n, you can exit the installer. You must configure the product before you can use Veritas product.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

- 7 Select the checkbox to specify whether you want to send your installation information to Symantec.

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

Click **Finish**.

Upgrading to 5.1 SP1 RP3

This chapter includes the following topics:

- [Prerequisites for upgrading to 5.1 SP1 RP3](#)
- [Downloading required software to upgrade to 5.1 SP1 RP3](#)
- [Supported upgrade paths](#)
- [Upgrading to 5.1 SP1 RP3](#)
- [Verifying software versions](#)

Prerequisites for upgrading to 5.1 SP1 RP3

The following list describes prerequisites for upgrading to the 5.1 SP1 RP3 release:

- For any product in the Veritas Storage Foundation stack, you must have the 5.1 SP1 (or later) installed before you can upgrade that product to the 5.1 SP1 RP3 release.
- Each system must have sufficient free space to accommodate patches.
- For RHEL6, un-install the VRTSaslapm package if it is already installed.
- The full list of prerequisites can be obtained by running `./installrp -precheck`
- Make sure to download the latest patches for the installer.
See [“Downloading required software to upgrade to 5.1 SP1 RP3”](#) on page 133.

Downloading required software to upgrade to 5.1 SP1 RP3

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

To download required software to upgrade to 5.1 SP1 RP3

- 1 Download SFHA 5.1 SP1 RP3 from <http://sort.symantec.com/patches>.
- 2 Extract it to a directory such as `/tmp/sfha51sp1rp3`.
- 3 Check <http://sort.symantec.com/patches> to see if there are patches available for the 5.1 SP1 RP3 installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 4 When you run the `installrp` script, use the `-require` option and specify the location where you downloaded the 5.1 SP1 RP3 installer patches.

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 5.1 SP1 to 5.1 SP1 RP3
- 5.1 SP1 PR1 to 5.1 SP1 RP3
- 5.1SP1PR2 to 5.1 SP1 RP3 on RHEL6
- 5.1 SP1 RP1 to 5.1 SP1 RP3
- 5.1 SP1 RP2 to 5.1 SP1 RP3
- 5.1 SP1 P-patch to 5.1 SP1 RP3
- 5.1 SP1 RP1 P-patch to 5.1 SP1 RP3
- 5.1 SP1 RP2 P-patch to 5.1 SP1 RP3

Upgrading to 5.1 SP1 RP3

This section describes how to upgrade from 5.1 SP1 (or later) to 5.1 SP1 RP3 on a cluster or a standalone system.

- [Performing a full upgrade to 5.1 SP1 RP3 on a cluster](#)
Use the procedures to perform a full upgrade to 5.1 SP1 RP3 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System (SFCFS), Veritas Storage Foundation for Oracle RAC (SFRAC), Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC), or Symantec VirtualStore (SVS) installed and configured.
- [Upgrading to 5.1 SP1 RP3 on a standalone system](#)
Use the procedure to upgrade to 5.1 SP1 RP3 on a system that has SF installed.
- [Performing a rolling upgrade using the installer](#)

Use the procedure to upgrade your Veritas product with a rolling upgrade.

- [Performing a phased upgrade to SFCFSA and SFRAC](#)

Use the procedure to perform a phase upgrade using the installer.

See [“Installing the Veritas software using the script-based installer”](#) on page 129.

Performing a full upgrade to 5.1 SP1 RP3 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 (SFCFS) and Cluster Volume Manager (CVM), the SFCFS and CVM services remain available.

Depending on your cluster’s configuration, select one of the following procedures to upgrade to 5.1 SP1 RP3:

- [Performing a full upgrade to 5.1 SP1 RP3 on a Veritas Cluster Server](#)

- [Performing a full upgrade to 5.1 SP1 RP3 on an SFHA cluster](#)

- [Performing a full upgrade to 5.1 SP1 RP3 on an SFCFS cluster](#)

- [Performing a full upgrade to 5.1 SP1 RP3 on an SFCFS RAC cluster](#)

- [Performing a full upgrade to 5.1 SP1 RP3 on an SF Oracle RAC cluster](#)

See [“Downloading required software to upgrade to 5.1 SP1 RP3 ”](#) on page 133.

Performing a full upgrade to 5.1 SP1 RP3 on a Veritas Cluster Server

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

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

To upgrade VCS

- 1 Make sure you have downloaded the latest software required for the upgrade.
See [“Downloading required software to upgrade to 5.1 SP1 RP3 ”](#) on page 133.
- 2 Log in as superuser.

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

See [“System requirements”](#) on page 16.

- 3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP3 rolling patch binaries, change to the directory that contains the `installrp` script. Start the pre-upgrade check:

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

See [“About the installrp script”](#) on page 12.

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

```
# ./installrp node1 node2 ... nodeN
```

See [“About the installrp script”](#) on page 12.

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

Performing a full upgrade to 5.1 SP1 RP3 on an SFHA cluster

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

To perform a full upgrade to 5.1 SP1 RP3 on an SFHA cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
See [“Downloading required software to upgrade to 5.1 SP1 RP3”](#) on page 133.
- 2 Log in as superuser.
- 3 From the directory that contains the extracted and untarred 5.1 SP1 RP3 rolling patch binaries, change to the directory that contains the `installrp` script. Start the upgrade.

```
# ./installrp node1 node2
```

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

Performing a full upgrade to 5.1 SP1 RP3 on an SFCFS cluster

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

To perform a full upgrade to 5.1 SP1 RP3 on an SFCFS cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
See [“Downloading required software to upgrade to 5.1 SP1 RP3”](#) on page 133.
- 2 Log in as superuser.

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

```
# mount | grep vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# mount | grep vxfs
```

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

```
# umount /filesystem
```

Note: If file system is CFS mounted then use `cfsumont` command.

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvrg` stop command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

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

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

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

- 6 Stop activity to all VxVM volumes.

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

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

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

Verify that no volumes remain open:

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

- 8 If required, apply the OS kernel patches.
- 9 From the directory that contains the extracted and untarred 5.1 SP1 RP3 rolling patch binaries, change to the directory that contains the `installrp` script. Start the upgrade.

```
# ./installrp node1 node2
```

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

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

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

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

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

- 14 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

- 15 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

Performing a full upgrade to 5.1 SP1 RP3 on an SFCFS RAC cluster

To prepare for a full upgrade to 5.1 SP1 RP3 on an SFCFS RAC cluster

- 1 Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 5.1 SP1 RP3”](#) on page 133.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
- 4 Stop the applications that are not managed by VCS. Use native application commands to stop the application.
- 5 Stop Oracle Clusterware.

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

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

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

- 7 Unmount all file systems:

```
# umount /filesystem
```

- 8 Check if the VEA service is running:

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

- 9 If the VEA service is running, stop it:

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

- 10 If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.

- 11 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.

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

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

- 13 To verify that no volumes remain open, use the following command:

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

- 14 Disable the startup scripts before upgrading the operating system.

```
# chkconfig vcs off
# chkconfig vxodm off
# chkconfig vxfen off
# chkconfig vxgms off
# chkconfig vxglm off
# chkconfig gab off
# chkconfig llt off
```

- 15 Upgrade the operating system. For instructions, see the operating system documentation.

To upgrade Storage Foundation Cluster File System for Oracle RAC

- 1 From the directory that contains the extracted and untarred 5.1 SP1 RP3 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

```
# ./installrp node1 node2 ... nodeN
```

- 2 After the initial system checks are complete, press Return to start the requirement checks.
- 3 When the Upgrade is complete, note the locations of the summary, log, and response files indicated by the installer.
- 4 Shut down and reboot the systems.
- 5 Upgrade Oracle RAC, if required.
- 6 Relink the Oracle's ODM library with Veritas ODM library.

- For Oracle RAC 10g:

- Change to the \$ORACLE_HOME/lib directory:

```
# cd $ORACLE_HOME/lib
```

- Back up libodm10.so file.

```
# mv libodm10.so libodm10.so.oracle-`date +%m_%d_%y-%H_%M_%S`
```

- Link libodm10.so file with the Veritas ODM library:

```
# ln -s /opt/VRTSodm/lib64/libodm.so libodm10.so
```

■ For Oracle 11g:

- Change to the \$ORACLE_HOME/lib directory:

```
# cd $ORACLE_HOME/lib
```

- Back up libodm11.so file.

```
# mv libodm11.so libodm11.so.oracle-`date '+%m_%d_%y-%H_%M_%S'`
```

- Link libodm11.so file with the Veritas ODM library:

```
# ln -s /opt/VRTSodm/lib64/libodm.so libodm11.so
```

To bring the upgraded cluster online and restore components

- 1 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.
- 2 If necessary, reinstate any missing mount points in the /etc/fstab file on each node.
- 3 If any VCS configuration files need to be restored, stop the cluster, restore the files to the /etc/VRTSvcS/conf/config directory, and restart the cluster.
- 4 Restart all the volumes by entering the following command for each disk group:

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

- 5 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

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

Performing a full upgrade to 5.1 SP1 RP3 on an SF Oracle RAC cluster

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

To upgrade to 5.1 SP1 RP3 on an SF Oracle RAC cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
See [“Downloading required software to upgrade to 5.1 SP1 RP3”](#) on page 133.

- 2 Log in as superuser.

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

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

```
# haconf -makerw
```

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

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

- 6 Make the configuration read-only:

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

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

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

- 8 Stop VCS.

```
# hastop -all
```

- 9 If required, apply the OS kernel patches.

See [“System requirements”](#) on page 16.

See *Oracle’s* documentation for the procedures.

- 10 From the directory that contains the extracted and untarred 5.1 SP1 RP3 rolling patch binaries, change to the directory that contains the `installrp` script. If ssh key authentication is configured then enter:

```
# ./installrp node1 node2
```

If ssh is not configured then enter:

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

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

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

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

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

Refer to *Veritas Storage Foundation for Oracle RAC 5.1 SP1 or later Installation and Configuration Guide* for more information.

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

```
# haconf -makerw
```

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

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

- 18 Make the configuration read-only:

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

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

```
# hagrps -online service_group -sys nodename
```

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

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

- 21 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

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

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

- 23 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

Upgrading to 5.1 SP1 RP3 on a standalone system

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

To upgrade to 5.1 SP1 RP3 on a standalone system

- 1 Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 5.1 SP1 RP3”](#) on page 133.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 4 If required, apply the OS kernel patches.
- 5 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

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

- 6 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
```

```
# umount /filesystem
```

- 7 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `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.

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

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

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

Verify that no volumes remain open:

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

- 10 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

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

```
# ./installrp nodename
```

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

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

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

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

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

15 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem  
# mount /checkpoint_name
```

16 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

Upgrading to 5.1 SP1 RP3 on a system that has encapsulated boot disk

You can use this procedure to upgrade to 5.1 SP1 RP3 on a system that has encapsulated boot disk.

Note: Upgrading with encapsulated boot disk from 5.1 SP1 to 5.1 SP1 RP3 requires multiple reboots.

To upgrade to 5.1 SP1 RP3 on a system that has encapsulated boot disk

- 1 Manually unmount file systems and stop open volumes.
- 2 If required, manually break the mirror.
- 3 Upgrade to 5.1 SP1 RP3 using `installrp` command.
- 4 After upgrading, reboot the system to have the new VM drivers take effect.
- 5 If the mirrors of boot disk are split manually in step 2, re-join the mirrors manually when systems reboot after upgrading to 5.1 SP1 RP3.

Performing a rolling upgrade using the installer

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

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

About rolling upgrades

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

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

You can perform a rolling upgrade for the following products:

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

You can perform a rolling upgrade from 5.1 SP1 or later.

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.
- For SFCFS for Oracle RAC, stop Oracle Clusterware before upgrading Kernel packages on any node.
- Make sure you have downloaded the latest software required for the upgrade. See [“Downloading required software to upgrade to 5.1 SP1 RP3 ”](#) on page 133.

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

Performing a rolling upgrade using the installer

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

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

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

To perform the rolling upgrade on kernel packages: phase 1

- 1 Stop all applications that access volumes.
- 2 Unmount all the file systems that are managed by SF.
- 3 On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.

```
# ./installrp -upgrade_kernelpkgs nodeA
```
- 4 Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 5 The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel packages.
- 6 The installer loads new kernel modules and starts all the relevant processes and brings all the service groups online.
- 7 If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step 8.
- 8 Before you proceed to phase 2, complete step 1 to step 6 on the second subcluster.

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

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

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

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

```
# ./installrp -upgrade_nonkernelpkgs nodeA nodeB nodeC nodeD
```

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

```
# hastatus -sum
```

- 5 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

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

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

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

To perform the rolling upgrade on kernel packages: phase 1

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

```
# ./installrp -upgrade_kernelpkgs nodeA
```

- 2 Note that if the boot disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 3 The installer checks system communications, package versions, product versions, and completes prechecks.
It then upgrades applicable product kernel packages.
- 4 The installer loads new kernel modules.

- 5 The installer starts all the relevant processes and brings all the service groups online.

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.

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

- 6 Relink the SF Oracle RAC libraries with Oracle by choosing the option **Relink Oracle Database Binary** from the program menu.

- 7 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
```

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

- 9 Start all applications that VCS does not manage. Use native application commands to start the applications.

- 10 If the boot disk is encapsulated, reboot the first sub-cluster's system.

- 11 Before you proceed to phase 2, complete step 1 to 10 on the second subcluster.

- 12 ■ If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If VCS does not manage the Oracle database, change the management policy for the database to automatic:

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

- 13 Migrate the SFDB repository database.

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

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

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

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

```
# ./installrp -upgrade_nonkernelpkgs nodeA nodeB nodeC...
```

- 2 The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3 The installer loads the new kernel modules.
- 4 The installer starts all relevant processes and brings all the service groups online.
- 5 Verify the cluster's status:

```
# hastatus -sum
```

- 6 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

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

- 7 If you want to upgrade the operating system, perform the following steps:

- Change to the `/opt/VRTS/install` directory on the node where you want to upgrade the operating system:

```
# cd /opt/VRTS/install
```

- Stop SF Oracle RAC:

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

- Upgrade the operating system. For instructions, see the operating system documentation.
- Reboot the nodes:

```
# shutdown -r now
```

Performing a phased upgrade to SFCFSHA and SFRAC

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

Performing a phased upgrade of SFCFSHA

Performing a phased upgrade involves the following tasks:

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

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

Prerequisites for a phased upgrade

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

Planning for a phased upgrade

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

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

Phased upgrade limitations

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

- While you perform the upgrades, do not start any modules. Also, do not add or remove service groups to any of the nodes.

- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.
- You can perform a phased upgrade when the root disk is encapsulated.

Moving the service groups to the second subcluster

To move the service groups to the second subcluster

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

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

- 2 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the applications.
- 3 On the first half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS.

```
# mount | grep vxfs
```

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

```
# fuser -c mount_point
```

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

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

```
# umount /mount_point
```

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

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

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

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

```
# hagrps -state group_name
```

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

```
# haconf -makerw
# hasys -freeze -persistent sys1
# haconf -dump -makero
```

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

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

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

- 9 Verify that only GAB ports a, b, d and h are open. Enter the following:

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

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

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

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

Verify that no volumes remain open. Enter the following:

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

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

Upgrading the SFCFSHA stack on the first subcluster

To upgrade the SFCFSHA stack on the first subcluster

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

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

```
# ./installrp sys1
```

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

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

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

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

Preparing the second subcluster

To prepare the second subcluster

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

```
# mount | grep vxfs
```

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

```
# fuser -c mount_point
```

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

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

```
# umount /mount_point
```

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

```
# haconf -makerw  
# hagrps -unfreeze group_name -persistent  
# haconf -dump -makero
```

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

```
# hagrps -offline group_name -sys sys4
```

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

```
# hagrps -state group_name
```

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

```
# hastop -local
```

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

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

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

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

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

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

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

Activating the first subcluster

To activate the first subcluster

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

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

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

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
```

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

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

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

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

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

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

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

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

Upgrading the operating system on the second subcluster

To upgrade the operating system on the second subcluster

- ◆ Enter the following.

```
# chkconfig vcs off
# chkconfig vxfen off
# chkconfig gab off
# chkconfig llt off
```

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

Upgrading the second subcluster

To upgrade the second subcluster

- ◆ Enter the following:

```
# ./installrp node_name
```


Completing the phased upgrade

To complete the phased upgrade

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

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

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

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -copy -from_sys \  
node01 -to_sys node03 node04
```

- 2 On each node in the second half of the cluster, enable fencing. Enter the following:

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

3 Enter the following before reboot of nodes:

```
# chkconfig vcs on
# chkconfig vxfs on
# chkconfig gab on
# chkconfig llt on
```

4 Restart the upgraded nodes in the second half of the cluster:

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

When second half of the nodes come up, all the GAB ports a, b, d, h, u, v, w and f are ONLINE. Also all the CFS mounts service groups come online automatically.

5 Manually mount the VxFS and CFS file systems that are not managed by VCS in the second half of the cluster.

6 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

Performing phased upgrade of SF Oracle RAC to version 5.1 SP1 RP3 and later releases

Table 3-1 illustrates the phased upgrade process. Each column describes the steps to be performed on the corresponding subcluster and the status of the subcluster when operations are performed on the other subcluster.

Table 3-1 Summary of phased upgrade

First half of the cluster	Second half of the cluster
---------------------------	----------------------------

SF Oracle RAC cluster before the upgrade:




Table 3-1 Summary of phased upgrade (*continued*)





First half of the cluster	Second half of the cluster
<p>STEP 1: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> ■ Switch failover applications. ■ Stop all parallel applications. <p>See “Step 1: Performing pre-upgrade tasks on the first half of the cluster” on page 164.</p> <p>STEP 2: Upgrade SF Oracle RAC.</p> <p>See “Step 2: Upgrading the first half of the cluster” on page 166.</p>	<p>The second half of the cluster is up and running.</p> 
<p>The first half of the cluster is not running.</p> 	<p>STEP 3: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> ■ Stop all parallel and failover applications. ■ Stop SF Oracle RAC. <p>See “Step 3: Performing pre-upgrade tasks on the second half of the cluster” on page 167.</p> <p>The downtime starts now.</p>
<p>STEP 4: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> ■ Start SF Oracle RAC. ■ Start all applications. <p>See “Step 4: Performing post-upgrade tasks on the first half of the cluster” on page 169.</p> <p>The downtime ends here.</p>	<p>The second half of the cluster is not running.</p> 
<p>The first half of the cluster is up and running.</p> 	<p>STEP 5: Upgrade SF Oracle RAC.</p> <p>See “Step 5: Upgrading the second half of the cluster” on page 170.</p> <p>STEP 6: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> ■ Start SF Oracle RAC. ■ Start all applications. <p>See “Step 6: Performing post-upgrade tasks on the second half of the cluster” on page 171.</p>

Table 3-1 Summary of phased upgrade (*continued*)

First half of the cluster	Second half of the cluster
---------------------------	----------------------------

The phased upgrade is complete and both the first and the second half of the cluster are running.



Step 1: Performing pre-upgrade tasks on the first half of the cluster

Perform the following pre-upgrade steps on the first half of the cluster.

To perform the pre-upgrade tasks on the first half of the cluster

- 1 Back up the following configuration files: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```

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

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

If not, a warning message is displayed.

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

- 3 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

4 Stop the applications configured under VCS. Stop the Oracle RAC database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys sys1
# hagrps -offline oracle_group -sys sys2
```

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

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

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

5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# 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
```

6 Unmount the CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs | grep cluster

# fuser -cu /mount_point
```

- Unmount the non-system CFS file system:

```
# umount /mount_point
```

- 7 Stop the parallel service groups and switch over failover service groups on each of the nodes in the first half of the cluster:

```
# hastop -local -evacuate
```

- 8 Unmount the VxFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
```

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

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 9 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.

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

```
# /opt/VRTS/install/installsfrac -stop sys1 sys2
```

Step 2: Upgrading the first half of the cluster

Perform the following steps to upgrade the first half of the cluster.

To upgrade the first half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

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

- 2 Upgrade the operating system, if required.
For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

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

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 Upgrade SF Oracle RAC. Navigate to the product directory on the installation media. When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

```
# cd product folder

# cd /dvd_mount/specific_os/

# ./installrp sys1 sys2
```

Note: After you complete the upgrade of the first half of the cluster, no GAB ports will be shown in the output when you run the `gabconfig -a` command.

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not updated
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.
```

```
SFRAC is licensed on the systems
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

Step 3: Performing pre-upgrade tasks on the second half of the cluster

Perform the following pre-upgrade steps on the second half of the cluster.

To perform the pre-upgrade tasks on the second half of the cluster

- 1 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

Note: The downtime starts now.

- 2 Stop all applications that are configured under VCS. Stop the Oracle RAC database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys sys3
# hagrps -offline oracle_group -sys sys4
```

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

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

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

3 Unmount the CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

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

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

4 Stop VCS on each of the nodes in the second half of the cluster:

```
# hastop -local
```

5 Unmount the VxFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

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


- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 6 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.

- 7 Stop all ports.

```
# /opt/VRTS/install/installsfrac -stop sys3 sys4
```

Step 4: Performing post-upgrade tasks on the first half of the cluster

Perform the following post-upgrade steps on the first half of the cluster.

To perform the post-upgrade tasks on the first half of the cluster

- 1 On any one node on the first half of the cluster, force GAB to form a cluster.

```
# /etc/init.d/llt start
```

```
# /etc/init.d/gab start
```

```
# gabconfig -x
```

- 2 On the first half of the cluster, start SF Oracle RAC:

```
# cd /opt/VRTS/install
```

```
# ./installsfrac -start sys1 sys2
```

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

- 4 Relink the SF Oracle RAC libraries with Oracle.

5 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrpl -online oracle_group -sys node_name
```

If the Oracle database is not managed by VCS:

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \  
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \  
-i instance_name
```

Note: The downtime ends here.

6 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

Step 5: Upgrading the second half of the cluster

Perform the following steps to upgrade the second half of the cluster.

To upgrade the second half of the cluster

1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

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

2 Upgrade the operating system, if required.

For instructions, see the operating system documentation.

3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

4 Rename the `/etc/llttab` file to its original name:

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

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 On the second half of the cluster, upgrade SF Oracle RAC. Navigate to the product directory on the installation media.

When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

```
# cd /dvd_mount/specific_os/  
# ./installrp sys3 sys4
```

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not  
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.  
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not updated  
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.
```

```
SFRAC is licensed on the systems  
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

Step 6: Performing post-upgrade tasks on the second half of the cluster

Perform the following post-upgrade steps on the second half of the cluster.

To perform the post-upgrade tasks on the second half of the cluster

- 1 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 2 On the second half of the cluster, start SF Oracle RAC:

```
# cd /opt/VRTS/install  
# ./installsfrac -start sys3 sys4
```

- 3 Relink the SF Oracle RAC libraries with Oracle.
- 4 Upgrade VxVM disk group version.
- 5 Upgrade disk layout version.

6 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys sys3
# hagrps -online oracle_group -sys sys4
```

If the Oracle database is not managed by VCS:

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \
-i instance_name
```

7 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# 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
```

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

9 Set or change the product license level, if required.

10 Migrate the SFDB repository database.

11 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1 RP3, make sure that you upgraded all application clusters to version 5.1 SP1 RP3. Then, upgrade VCS or SFHA on the CP server systems.

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

Upgrading the operating system

This section describes how to upgrade the operating system on a Storage Foundation node where you plan to upgrade to 5.1 SP1 RP3 for RHEL 5, RHEL6, SLES11, OEL 5, and OEL6.

To upgrade the operating system to a later version

- 1 Stop Storage Foundation.
- 2 Upgrade to the latest operating system.
- 3 Upgrade to 5.1 SP1 RP3.
See [“Upgrading to 5.1 SP1 RP3”](#) on page 134.
- 4 Start Storage Foundation.

Verifying software versions

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

```
# rpm -qa | egrep VRTS
```


Uninstalling version 5.1 SP1 RP3

This chapter includes the following topics:

- [About removing Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP3](#)
- [Uninstalling Veritas Storage Foundation Cluster File System 5.1 SP1 RP3](#)
- [Uninstalling Veritas Storage Foundation for Oracle RAC](#)
- [Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP3](#)

About removing Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP3

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

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

Uninstalling Veritas Storage Foundation Cluster File System 5.1 SP1 RP3

This section provides procedures for uninstalling Veritas Storage Foundation Cluster File System (SFCFS). You must complete the preparatory tasks before you uninstall SFCFS.

Preparing to uninstall Veritas Storage Foundation Cluster File System

The following procedure prepares your system for uninstalling Veritas Storage Foundation Cluster File System (SFCFS).

To prepare to uninstall Veritas Storage Foundation Cluster File System

- 1 Log in as the root user on any node in the cluster.
- 2 Verify that the following directories are set in your `PATH` environment variable:

```
/opt/VRTS/bin
/opt/VRTSvcs/bin
```

- 3 Back up the following configuration files:

```
# mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
```

- 4 Determine if each node's root disk is under VxVM control and proceed as follows.

- Check if each node's root disk is under VxVM control:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

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


Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

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

```
# vxrvrg -g diskgroup stop rvg_name
```

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

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

- 6 Check if any VxFS file systems or Storage Checkpoints are mounted:

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

- 7 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint1
# umount /filesystem1
```

If file system is mounted in a cluster, then use `cfsumount` command.

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

```
# vxvol -g dgl stopall
```

To verify that no volumes are open:

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

- 9 Stop VCS:

```
# hastop -all
```

Uninstalling Veritas Storage Foundation Cluster File System

The following procedure uninstalls Veritas Storage Foundation Cluster File System (SFCFS).

To uninstall Veritas Storage Foundation Cluster File System

- 1 Log in as the root user on any node in the cluster.
- 2 Navigate to the directory that contains the uninstallation program:

```
# cd /opt/VRTS/install
```

- 3 Run the uninstallation program while specifying each node in the cluster:

```
# ./uninstallsfcfs node1 node2
```

- 4 Confirm the uninstallation:

```
Are you sure you want to uninstall SFCFS [y,n,q] (y)
```

The installer stops the Veritas Storage Foundation Cluster File System processes and uninstalls the packages.

After uninstalling the Veritas Storage Foundation Cluster File System, refer to the *Veritas Storage Foundation Cluster File System 5.1 SP1 Installation Guide* to reinstall the 5.1 SP1 software.

Uninstalling Veritas Storage Foundation for Oracle RAC

The following procedure uninstalls Veritas Storage Foundation for Oracle RAC (SFRAC).

Note: This procedure will remove the complete SFRAC stack from all nodes.

To uninstall Veritas Storage Foundation for Oracle RAC

- 1 On each node, take the Oracle resources in the VCS configuration file (`main.cf`) offline.

```
# hagrps -offline oracle_group -sys node_name
```

If the database is not managed by VCS, stop the Oracle database as follows:

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

- 2 If Oracle Clusterware is not under VCS Control, then enter the following command on each node of the cluster to stop Oracle Clusterware.

- For 10gR2 or 11gR1:

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

- For 11gR2:

```
$ GRID_HOME/bin/crsctl stop crs
```

3 Stop the applications that use CVM or CFS that are not under VCS control

- Using native application commands, stop the applications that use CVM or CFS on all nodes.
- Verify that no processes use the CFS mount point:

```
# fuser -c mount_point
```

4 Unmount CFS file systems that are not under VCS control

- Determine the file systems that need to be unmounted by checking the output of mount command.

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

- Unmount each file system that is not controlled by VCS on each node:

```
# umount mount_point
```

5 Stop VCS to take the service groups on all nodes offline

On any one node execute following command to stop VCS:

```
# hastop -all
```

6 Stopping the applications that use VxVM or VxFS that are not under VCS control

- Using native application commands, stop the applications that use VxVM or VxFS.
- Verify that no processes use the VxFS mount point:

```
# fuser -c mount_point
```

7 Unmounting VxFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

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

Unmount each file system that is not controlled by VCS on each node:

```
# umount mount_point
```

8 Remove SF for Oracle RAC.

- On any one node, navigate to the directory that contains the `uninstallsfrac` program:

```
# cd /opt/VRTS/install
```

- Start the `uninstallsfrac` program:

```
# ./uninstallsfrac
```

9 After uninstalling the SFRAC, refer to the *Veritas Storage Foundation for Oracle RAC 5.1 SP1 Installation and Configuration Guide* document to reinstall the SFRAC 5.1 SP1 software.

Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP3

To prepare to uninstall Veritas Storage Foundation Cluster File System for Oracle RAC from a cluster

- 1 Log in as the root user on any node in the cluster.
- 2 Verify that the following directories are set in your `PATH` environment variable in order to execute the necessary commands:

```
/opt/VRTS/bin  
/opt/VRTSvcs/bin
```

3 Back up the following configuration files:

```
# mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`  
# mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`  
# mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`  
# mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
```

Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP3

- 4 On all the nodes, stop the CFS-dependant applications that are not under VCS control using application specific commands.

For example, to stop Oracle Clusterware:

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

- 5 Stop VCS:

```
# hastop -all
```

- 6 Verify that port h is not open:

```
# gabconfig -a
```

- 7 Check if any VxFS file systems or Storage Checkpoints are mounted:

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

- 8 Unmount all file systems:

```
# umount /filesystem
```

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

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

To verify that no volumes are open:

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

Perform the steps in the following procedure to uninstall Veritas Storage Foundation Cluster File System for Oracle RAC from a cluster.

To uninstall Veritas Storage Foundation Cluster File System for Oracle RAC from a cluster

- 1 Log in as the root user on any node in the cluster.
- 2 Navigate to the directory that contains the uninstallation program:

```
# cd /opt/VRTS/install
```

- 3 Start the uninstallation program:

```
# ./uninstallsfcfsrac node1 node2 ... nodeN
```

- 4 Press Enter to uninstall Veritas Storage Foundation Cluster File System for Oracle RAC.

```
Do you want to uninstall SFCFSRAC from these systems [y,n,q] (y)
```

The installer checks the RPMs installed on the system.

- 5 Confirm the uninstallation at the following prompt:

```
Are you sure you want to uninstall SFCFSRAC [y,n,q] (y)
```

The installer stops the Veritas Storage Foundation Cluster File System for Oracle RAC processes and uninstalls the packages.

After uninstalling the Veritas Storage Foundation Cluster File System for Oracle RAC, refer to the *Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 Installation and Configuration Guide* to reinstall the 5.1 SP1 software.

DMP as a multi-pathing solution in an Oracle VM Server for SPARC environment

This appendix includes the following topics:

- [DMP as a multi-pathing solution in an Oracle VM Server for SPARC environment](#)

DMP as a multi-pathing solution in an Oracle VM Server for SPARC environment

Using Dynamic Multi-Pathing (DMP), you can choose to manage multiple paths to a system's storage in the following domains:

- control domain
- I/O domains
- guest domains

For configurations that have multiple I/O domains, Symantec recommends that you have DMP manage the paths inside of the I/O domains.

Note: Not having a multi-pathing solution can cause accidental data corruption when the same device is accessed.

Using DMP in the control domain lets you use all of the Storage Foundation (SF) features from inside of the guest. You must enable DMP in the control and alternate I/O domains.

See [“Enabling DMP in the control and alternate I/O domains”](#) on page 184.

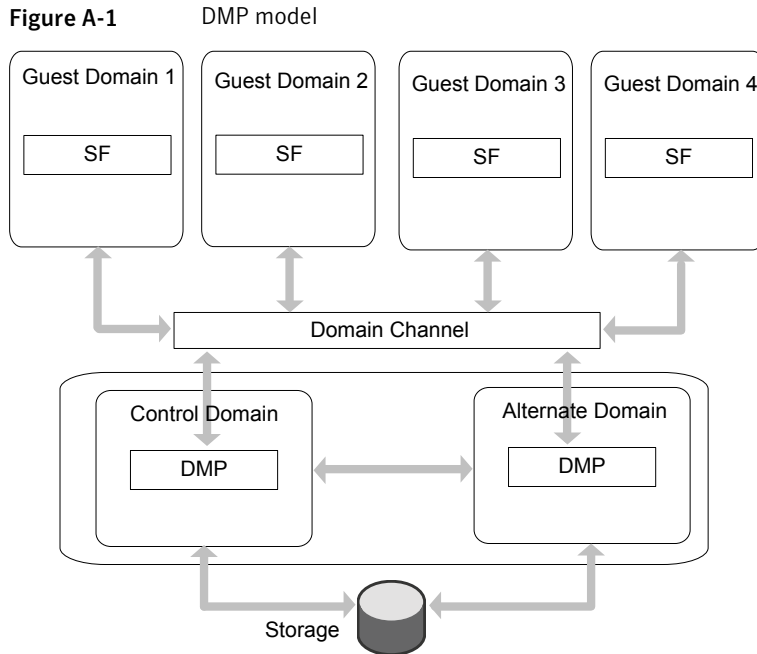
You can enable DMP path failover in the guest domain.

See [“Enabling DMP path failover in the guest domain”](#) on page 185.

Enabling DMP in the control and alternate I/O domains

This section describes how to enable Dynamic Multi-Pathing (DMP) in the control and alternate I/O domains.

[Figure A-1](#) illustrates the architecture of DMP.



To enable DMP in the control and alternate I/O domains

- 1 Install VRTSvxvm and VRTSaslapm packages on both the control and alternate I/O domains.

For information on manual installation of DMP:

See the *Veritas Dynamic Multi-Pathing Installation Guide*.

- 2 Turn on the enclosure-based naming default setting in both the control and alternate domains.

```
# vxddladm set namingscheme=ebn
```

- 3 Create the Virtual Disk Service (VDS) devices on DMP metanodes that can be provisioned to the guest domains.

For example:

```
# ldm add-vdsdev /dev/vx/dmp/xyz vo10015-001-p1@primary-vds0
# ldm add-vdsdev /dev/vx/dmp/xyz vo10015-001-p2@alternate-vds0
```

- 4 While provisioning, export the DMP metanodes from both the control and alternate I/O domain:

```
# ldm add-vdsdev /dev/vx/dmp/xyz vo10015-001-p1@primary-vds0
# ldm add-vdsdev /dev/vx/dmp/xyz vo10015-001-p2@alternate-vds0
# ldm add-vdisk timeout=30 vdisk0015-001-p1 \
vo10015-001-p1@primary-vds0 hssxd0015
# ldm add-vdisk timeout=30 vdisk0015-001-p2 \
vo10015-001-p2@alternate-vds0 hssxd0015
```

This allows DMP in the guest domain to see two access paths, one through the control domain and the other through the alternate domain to the storage.

DMP in the guest domain can take care of the control and alternate I/O domain failures.

Enabling DMP path failover in the guest domain

In Oracle VM Server configurations the Virtual Disk Client (VDC) driver timeout is set to zero by default which signifies infinity. This can cause the failed I/O not to return to the guest domain, if either the control or alternate I/O domain crashes unexpectedly. As a result the guest domain cannot get back the failed I/Os and cannot route them through the alternate domain. If this issue occurs or to avoid this issue, you must set the VDC driver timeout.

There are two ways to set the VDC driver timeout:

Modify globally all of the LUNs that are exported to the current guest domain. This requires a reboot to all the guest domains. See [“To change the VDC driver timeout globally”](#) on page 186.

Manually export every LUN directly to the guest domain and set the timeout parameter to 30 seconds. No reboot is required. See [“To change the VDC driver timeout for each LUN”](#) on page 186.

To change the VDC driver timeout globally

- 1 On each guest domain, edit the `/etc/system` file and add the following line to set the VDC driver timeout to 30 seconds:

```
set vdc:vdc_timeout=30
```

- 2 Reboot the guest domains.

To change the VDC driver timeout for each LUN

- 1 Create the primary domain using four internal disks and get all required SAN LUNs for the guest domains allocated to the primary domain.
- 2 Turn on enclosure-based naming in the primary domain.

```
# vxddladm set namingscheme=ebn
```

- 3 Remove half of the system's I/O from the primary domain:

```
# ldm remove-io pci_X primary_domain_name
```

where `pci_x` is the name of the PCI bus for your system.

where `primary_domain_name` is the name of the primary domain.

For example:

```
# ldm remove-io pci_@400 primary
```

- 4 Create the alternate I/O domain on the other four internal disks and add the I/O that was removed from the primary domain:

```
# ldm add-io pci_X primary_domain_name
```

where `pci_x` is the name of the PCI bus for your system.

where `primary_domain_name` is the name of the primary domain.

For example:

```
# ldm add-io pci_@400 primary
```

5 Turn on enclosure-based naming in the alternate I/O domain.

```
# vxddladm set namingscheme=ebn
```

At this point you have one of the two paths to each SAN LUN in the primary domain and the other path is in the alternate I/O domain. Both will have the same name: `/dev/vx/dmp/enclosure_based_name`.

6 On the primary domain, create the guest domains. In the sample the enclosure-based name of one of the LUNs is `xyz` and the guest domain is

`hsxd0015`:

```
# ldm add-vdsdev /dev/vx/dmp/xyz vol0015-001-p1@primary-vds0
# ldm add-vdsdev /dev/vx/dmp/xyz vol0015-001-p2@alternate-vds0
# ldm add-vdisk timeout=30 vdisk0015-001-p1 \
vol0015-001-p1@primary-vds0 hsxd0015
# ldm add-vdisk timeout=30 vdisk0015-001-p2 \
vol0015-001-p2@alternate-vds0 hsxd0015
```

The same set of four commands for each SAN LUN that gets placed in a guest domain. Use three SAN LUNs for SAN boot in the guest domain and the rest for application data. Each LUN in the guest domain has one path backup through the primary domain and one backup through the alternate domain. That means each LUN only uses one LDC in each domain. Also, since you are using DMP you still only use 1 LDC in each domain, even if the LUN has more than two paths from the array.

