

# Veritas Storage Foundation™ and High Availability Release Notes

Solaris

6.0 Rolling Patch 1

# Veritas Storage Foundation and High Availability Release Notes

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

Product version: 6.0 Rolling Patch 1

Document version: 6.0RP1.2

## Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec, the Symantec 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.

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 Web site at the following URL:

[www.symantec.com/business/support/index.jsp](http://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](http://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/](http://www.symantec.com/business/support/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/business/support/](http://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](mailto:customercare_apac@symantec.com)

Europe, Middle-East, and Africa [semea@symantec.com](mailto: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 Web site.

<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](mailto: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 and the uninstallrp scripts .....	12
	The installrp script options .....	12
	The uninstallrp script options .....	15
	Overview of the installation and upgrade process .....	18
	System requirements .....	18
	Supported Solaris operating systems .....	19
	Hardware compatibility list .....	19
	Veritas Storage Foundation for Database features supported in database environments .....	19
	Veritas Storage Foundation memory requirements .....	20
	Number of nodes supported .....	20
	Fixed issues .....	20
	Veritas Volume Manager: Issues fixed in 6.0 RP1 .....	20
	Veritas File System: Issues fixed in 6.0 RP1 .....	22
	Veritas Cluster Server: Issues fixed in 6.0 RP1 .....	24
	Veritas Enterprise Administrator: Issues fixed in 6.0 RP1 .....	24
	Known issues .....	25
	Installation known issues .....	25
	Veritas Dynamic Multi-pathing known issues .....	35
	Veritas Storage Foundation known issues .....	39
	Veritas Cluster Server known issues .....	62
	Veritas Storage Foundation and High Availability known issues .....	81
	Veritas Storage Foundation Cluster File System High Availability known issues .....	81
	Veritas Storage Foundation for Oracle RAC known issues .....	86
	Veritas Storage Foundation for Sybase ASE CE known issues .....	91
	Symantec VirtualStore known issues .....	93
	Software limitations .....	97

	List of patches .....	97
	Documentation errata .....	99
	Veritas Storage Foundation Cluster File System High Availability manual page .....	100
Chapter 2	Installing the products for the first time .....	101
	Installing the Veritas software using the script-based installer .....	101
	Installing Veritas software using the Web-based installer .....	102
	Starting the Veritas Web-based installer .....	103
	Obtaining a security exception on Mozilla Firefox .....	103
	Installing 6.0 RP1 with the Veritas Web-based installer .....	103
Chapter 3	Installing the products using JumpStart .....	107
	Installing the products using JumpStart .....	107
	Generating the finish scripts .....	107
	Overview of JumpStart installation tasks .....	111
	Preparing installation resources .....	112
Chapter 4	Upgrading to 6.0 RP1 .....	115
	Prerequisites for upgrading to 6.0 RP1 .....	115
	Downloading required software to upgrade to 6.0 RP1 .....	115
	Supported upgrade paths .....	116
	Upgrading to 6.0 RP1 .....	116
	Performing a full upgrade to 6.0 RP1 on a cluster .....	116
	Upgrading to 6.0 RP1 on a standalone system .....	125
	Upgrading Veritas products using Live Upgrade .....	128
	Performing a rolling upgrade using the installer .....	134
	Verifying software versions .....	142
Chapter 5	Rolling back Veritas Storage Foundation and High Availability Solutions .....	145
	About rolling back Veritas Storage Foundation and High Availability Solutions 6.0 RP1 .....	145
	Rolling back using the uninstallrp script .....	146
	Rolling back manually .....	148
	Rolling back Storage Foundation or Storage Foundation and High Availability manually .....	148
	Rolling back Storage Foundation Cluster File System High Availability manually .....	151
	Rolling back Veritas Cluster Server manually .....	154
	Rolling back Symantec VirtualStore manually .....	158

Rolling back Dynamic Multi-Pathing manually ..... 161



# About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [Introduction](#)
- [About the installrp and the uninstallrp scripts](#)
- [Overview of the installation and upgrade process](#)
- [System requirements](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [List of patches](#)
- [Documentation errata](#)

## Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 6.0 Rolling Patch 1 release.

## About the installrp and the uninstallrp scripts

Veritas Storage Foundation and High Availability Solutions 6.0 RP1 provides an upgrade script.

See [“Supported upgrade paths”](#) on page 116.

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 installation options. A system name is required for all options. If not specified, the command prompts for a system name.
<code>-precheck</code>	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
<code>-postcheck</code>	Checks any issues after installation or upgrading on the system.
<code>-responsefile response_file</code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<code>-logpath log_path</code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
<code>-tmppath tmp_path</code>	Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.

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

Command Line Option	Function
<code>-timeout <i>timeout_value</i></code>	The <code>-timeout</code> option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the <code>-timeout</code> option overrides the default value of 1200 seconds. Setting the <code>-timeout</code> option to 0 will prevent the script from timing out. The <code>-timeout</code> option does not work with the <code>-serial</code> option
<code>-keyfile <i>ssh_key_file</i></code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i <i>ssh_key_file</i></code> to every SSH invocation.
<code>-hostfile <i>full_path_to_file</i></code>	Specifies the location of a file that contains a list of hostnames on which to install.
<code>-patchpath <i>patch_path</i></code>	Designates the path of a directory that contains all patches to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
<code>-nim</code>	Produces a NIM configuration file for installing with NIM.
<code>-jumpstart <i>dir_path</i></code>	Produces a sample finish file for Solaris JumpStart installation. The <i>dir_path</i> indicates the path to the directory in which to create the finish file.
<code>-rootpath <i>root_path</i></code>	Specifies an alternative root directory on which to install packages.
<code>flash_archive&lt;<i>flash_archive_path</i>&gt;</code>	The <code>-flash_archive</code> option is used to generate Flash archive scripts which can be used by Solaris Jumpstart Server for automated Flash archive installation of all packages and patches for every product, an available location to store the post deployment scripts should be specified as a complete path. The <code>-flash_archive</code> option is supported on Solaris only.
<code>-serial</code>	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.

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

Command Line Option	Function
-rsh	Specifies this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
-redirect	Displays progress details without showing the progress bar.
-pkgset	Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.
-pkgtable	Displays product's packages in correct installation order by group.
-pkginfo	Displays a list of packages and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS packages.
-listpatches	The -listpatches option displays product patches in correct installation order.
-makeresponsefile	Use the -makeresponsefile option only to generate response files. No actual software installation occurs when you use this option.
-comcleanup	The -comcleanup option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.

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

Command Line Option	Function
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
ignorepatchreqs	The -ignorepatchreqs option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The -rollingupgrade_phase1 option is used to perform rolling upgrade Phase-I. In the phase, the product kernel packages get upgraded to the latest version
-rollingupgrade_phase2	The -rollingupgrade_phase2 option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."

## The uninstallrp script options

Veritas Storage Foundation and High Availability Solutions 6.0 RP1 provides a new uninstallation script.

See [About rolling back Veritas Storage Foundation and High Availability Solutions 6.0 RP1](#) for release versions and products that support rolling back.

Symantec recommends that you use the new uninstallation script. The `uninstallrp` script uninstalls all the patches associated with packages installed, and starts the processes.

**Table 1-2** 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 installation options. A system name is required for all options. If not specified, the command prompts for a system name.
<i>-responsefile response_file</i>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<i>-logpath log_path</i>	Specifies a directory other than <i>/opt/VRTS/install/logs</i> as the location where installer log files, summary files, and response files are saved.
<i>-tmppath tmp_path</i>	Specifies a directory other than <i>/var/tmp</i> as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
<i>-timeout timeout_value</i>	The <i>-timeout</i> option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the <i>-timeout</i> option overrides the default value of 1200 seconds. Setting the <i>-timeout</i> option to 0 will prevent the script from timing out. The <i>-timeout</i> option does not work with the <i>-serial</i> option
<i>-keyfile ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes <i>-i ssh_key_file</i> to every SSH invocation.
<i>-hostfile full_path_to_file</i>	Specifies the location of a file that contains a list of hostnames on which to install.
<i>-rootpath root_path</i>	Specifies an alternative root directory on which to install packages.

**Table 1-2** The command line options for the product upgrade script (*continued*)

Command Line Option	Function
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-rsh	Specifies this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
-redirect	Displays progress details without showing the progress bar.
-listpatches	The <code>-listpatches</code> option displays product patches in correct installation order.
-makeresponsefile	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.
-comcleanup	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.
ignorepatchreqs	The <code>-ignorepatchreqs</code> option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.

# 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

- 1 Skip this step if you are upgrading to 5.1 SP1 RP2. If you are installing 5.1 SP1 RP2 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 5.1SP1 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
```

- 2 Download SFHA 5.1 SP1 RP2 from <http://sort.symantec.com/patches> and extract it to a directory called /tmp/sfha51sp1rp2.
- 3 Check <http://sort.symantec.com/patches> to see if there are patches available for the 5.1SP1RP2 installer. Download applicable P-patches and extract them to the /tmp directory.
- 4 Change the directory to /tmp/sfha51sp1rp2:

```
#cd /tmp/sfha51sp1rp2
```

- 5 Install 5.1SP1 RP2:

```
./installer -require complete_path_to_SP1RP2_installer_patch
```

## System requirements

This section describes the system requirements for this release.

## Supported Solaris operating systems

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

[Table 1-3](#) shows the supported operating systems for this release.

**Table 1-3** Supported operating systems

Operating systems	Levels	Chipsets
Solaris 10	Update 8, 9, and 10.	SPARC
Solaris 10	Update 8, 9, and 10.	x64

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

## Hardware compatibility list

The compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware go to the following URL:

<http://www.symantec.com/docs/TECH170013>

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

## Veritas Storage Foundation for Database features supported in database environments

Veritas Storage Foundation for Database (SFDB) features are supported for the following database environments:

**Table 1-4** SFDB features database support for 6.0 RP1

SFDB feature	DB2	Oracle	Sybase
Oracle Disk Manager, Cached Oracle Disk Manager	No	Yes	No
Quick I/O, Cached Quick I/O	Yes	Yes	Yes
Concurrent I/O	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes

**Table 1-4** SFDB features database support for 6.0 RP1 (*continued*)

SFDB feature	DB2	Oracle	Sybase
Flashsnap	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes
Database Storage Checkpoints	No	Yes	No
Database Flashsnap	No	Yes	No
SmartTier for Oracle	No	Yes	No

Review current documentation for your database to confirm the compatibility of your hardware and software.

For the most current information on Storage Foundation products and single instance Oracle versions supported, see:

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

## Veritas Storage Foundation memory requirements

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

## Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

## Fixed issues

This section covers the incidents that are fixed in this release.

### Veritas Volume Manager: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.0 RP1.

**Table 1-5** Veritas Volume Manager 6.0 RP1 fixed issues

Fixed issues	Description
2680604	vxconfigbackupd does not work correctly with NUM_BK.

**Table 1-5** Veritas Volume Manager 6.0 RP1 fixed issues (*continued*)

<b>Fixed issues</b>	<b>Description</b>
2674465	Data Corruption while adding/removing LUNs.
2666163	A small portion of possible memory leak introduced due to addition of enhanced messages.
2657797	Starting 32TB RAID5 volume fails with unexpected kernel error in configuration update.
2649958	vxmpadm dumps core due to null pointer reference.
2647795	Intermittent data corruption after a vxassist move.
2635476	Failure in recovering a DMP failed path.
2632120	vxdiskadm utility does not update default DM name when multiple disks are specified for encapsulation.
2627056	vxmake -g <DGNAME> -d <desc-file> fails with very large configuration due to memory leaks.
2626741	Using vxassist -o ordered and mediatype:hdd options together do not work as expected.
2621465	When detached disk after connectivity restoration is tried to reattach gives 'Tagid conflict' error.
2620556	I/O hung after SRL overflow.
2620555	I/O hang due to SRL overflow and CVM reconfig.
2613425	Encapsulation issue - vxdiskadm should not have a default disk format of cdsdisk, it should be sliced.
2608849	Logowner local I/O starved with heavy I/O load from Logclient.
2607519	Secondary master panics in case of reconfig during autosync.
2607293	Primary master panic'ed when user deleted frozen RVG.
2600863	vxtune doesn't accept tunables correctly in human readable format.
2590183	write fails on volume on slave node after join which earlier had disks in "Ifailed" state.
2589569	vxdisksetup on EFI disk is taking ~2-4 mins.

**Table 1-5** Veritas Volume Manager 6.0 RP1 fixed issues (*continued*)

Fixed issues	Description
2576602	vxdg listtag should give error message and display correct usage when executed with wrong syntax.
2575581	vxtune -r option is printing wrong tunable value.
2574752	Support utility vxfmrmap (deprecating vxfmrshowmap) to display DCO map contents and verification against possible state corruptions.
2565569	read/seek i/o errors during init/define of nopriv slice.
2562416	vxconfigbackup throws script errors due to improper handling of arguments.
2556467	disabling all paths and rebooting host causes /etc/vx/.vxdmprawdev record loss.
2530698	after "vxdg destroy" hung (for shared DG), all vxcommands hang on master.
2526498	Memory leaks seen in some I/O code path.
2516584	startup scripts use 'quit' instead of 'exit', causing empty directories in /tmp.
2348180	Failure during validating mirror name interface for linked mirror volume.

## Veritas File System: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas File System (VxFS) in 6.0 RP1.

**Table 1-6** Veritas File System 6.0 RP1 fixed issues

Fixed issues	Description
2679361	Network Customization screen doesn't show any NICs in I18N-level0 environment.
2678096	The fiostat command dumps core when the count value is 0.
2663750	Abrupt messages are seen in engine log after complete storage failure in cvm resiliency scenario.
2655786	'Shared' extents are not transferred as 'shared' by the replication process.
2655754	Deadlock because of wrong spin lock interrupt level at which delayed allocation list lock is taken.

**Table 1-6** Veritas File System 6.0 RP1 fixed issues (*continued*)

<b>Fixed issues</b>	<b>Description</b>
2653845	When the fsckptadm(1M) command with the '-r' and '-R' option is executed, two mutually exclusive options gets executed simultaneously.
2645441	Native filesystem migrated to vxfs disk layout 8 where layout version 9 is the default.
2645435	The following error message is displayed during the execution of the fsmap(1M) command: 'UX:vxfs fsmap: ERROR: V-3-27313'.
2645112	write operation on a regular file mapping to shared compressed extent results in corruption.
2645109	In certain rare cases after a successful execution of vxfilesnap command, if the source file gets deleted in a very short span of time after the filesnap operation, then the destination file can get corrupted and this could also lead to setting of VX_FULLFCK flag in the super block.
2645108	In certain cases write on a regular file which has shared extent as the last allocated extent can fail with EIO error.
2630954	The fsck(1M) command exits during an internal CFS stress reconfiguration testing.
2630754	64-bit vxfsutil.so on solaris x86 doesn't load.
2624459	Listing of a partitioned directory using the DMAPi does not list all the entries.
2613884	Metadata corruption may be seen after recovery.
2609002	The De-duplication session does not complete.
2600168	The -p option of cp_vxfs command does not work correctly in solaris.
2599590	Expanding or shrinking a DLV5 file system using the fsadm(1M) command causes a system panic.
2583197	Upgrade of a file system from version 8 to 9 fails in the presence of partition directories and clones.
2563251	fsmigadm "commit/status" error messages should be clear.
2552095	The system may panic while re-organizing the file system using the fsadm(1M) command.
2536130	The fscdsconv(1M) command which is used to convert corrupted or non-VxFS file systems generates core.

**Table 1-6** Veritas File System 6.0 RP1 fixed issues (*continued*)

Fixed issues	Description
2389318	Enabling delayed allocation on a small file system sometimes disables the file system.

## Veritas Cluster Server: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas Cluster Server (VCS) in 6.0 RP1.

**Table 1-7** Veritas Cluster Server 6.0 RP1 fixed issues

Fixed issues	Description
2684822	If a pure local attribute like PreOnline is specified before SystemList in main.cf then it gets rejected when HAD is started.
2680435	IPMultiNICB agent does not carry Options attribute in local zone because of RunInContainer set to 0.
2662766	The clean entry point for Mount agent fails to un-mount a file system of type nfs.
2653668	The high availability daemon (HAD) process unexpectedly terminates.
2644483	"VCS ERROR V-16-25-50036 The child service group came online (recovered) before the parent was offlined." message is logging as ERROR message
2635211	AMF calls VxFS API with spinlock held.
2632576	Unable to stop to load amf driver even if amf in SMF is disabled
2616497	Faults in multiple tiers are not handled.

## Veritas Enterprise Administrator: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas Enterprise Administrator (VEA) in 6.0 RP1.

**Table 1-8** Veritas Enterprise Administrator 6.0 RP1 fixed issues

Fixed issues	Description
2672039	vxsvc process crashes and dumps core.
2677191	File placement policy operations are not available in VEA GUI.

## Known issues

This section covers the known issues in this release.

### Installation known issues

This section describes the known issues in this release of installation.

#### Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups (2591399)

The service groups freeze due to upgrading using the product installer if you stopped the installer after the installer already stopped some of the processes and then resumed the upgrade.

**Workaround:** You must unfreeze the service groups manually after the upgrade completes.

##### To unfreeze the service groups manually

- 1 List all the frozen service groups

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

- 2 Unfreeze all the frozen service groups:

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

#### Installed 5.0 MP3 without configuration, then upgrade to 6.0, installer cannot continue (2016346)

If you install 5.0MP3 without configuration, you cannot upgrade to 6.0. This upgrade path is not supported.

**Workaround:** Uninstall 5.0 MP3, and then install 6.0.

## Live Upgrade fails when you try to upgrade to Solaris 10 9/10 or later

When you try to upgrade to Solaris 10 9/10 or later, Live Upgrade fails. The Live Upgrade command, `luupgrade`, requires the `-k auto-registration-file` option, which Symantec's `vxlustart` script does not support.

### To resolve this issue

- 1 Copy the `luupgrade` command that failed during the execution of the `vxlustart` command. For example:

```
# luupgrade -u -n dest.18864 \  
-s /net/lyptus-new/image/solaris10/update9_GA63521  
blocksminiroot filesystem is <lofs>  
Mounting miniroot at  
  </net/lyptus-new/image/solaris10/update9_GA/Solaris_10/Tools/Boot>  
ERROR: The auto registration file <> does not exist or incomplete.  
The auto registration file is mandatory for this upgrade.  
Use -k <filename> argument along with luupgrade command.  
cat: cannot open /tmp/.liveupgrade.11624.24307/.lmz.list  
ERROR: vxlustart: Failed: luupgrade -u -n dest.18864  
  -s/net/lyptus-new/image/solaris10/update9_GA
```

In this example, you would copy the `luupgrade -u -n dest.18864 -s/net/lyptus-new/image/solaris10/update9_GA` command.

- 2 Paste the command, and append the command with the `-k auto-registration-file` option. For example:

```
# luupgrade -u -n dest.18864 \  
-s /net/lyptus-new/image/solaris10/update9_GA -k /regfile  
  
/regfile is absolute path for the auto-registration file.
```

- 3 Mount the destination boot environment to `/altroot.5.10`. Do the following:

- Display the source and destination boot environment. Enter:

```
# lustatus
```

- Mount the boot environment. Enter:

```
# lumount destination_boot_environment /altroot.5.1
```

- 4 After `luupgrade` completes and after mounting alternate boot environment, upgrade the Storage Foundation packages on the alternate root path using the following command:

```
# installsf -rootpath /altroot.5.10 -upgrade
```

If you are upgrading from Solaris 9 to 10, do the following in the order presented:

- Remove the currently installed Storage Foundation packages. Enter:

```
# uninstallsf -rootpath /altroot.5.10
```

- Upgrade Storage Foundation to 6.0 RP1. Enter:

```
# installsf -rootpath /altroot.5.1
```

- 5 Activate the destination boot environment. Do the following in the order presented:

- Display the source and destination boot environment. Enter:

```
# lustatus
```

- Unmount the source and destination boot environment alternate root path. Enter:

```
# luumount destination_boot_environment
```

- Activate the destination boot environment. Enter:

```
# luactivate
```

- 6 If the system was encapsulated, manually encapsulate the destination boot environment after it is booted.

## During Live Upgrade, installer displays incorrect message about VRTSaa package removal

If you use Live Upgrade to upgrade SFHA 5.0MP1 to SFHA 6.0, the installer may display a message that the VRTSaa package failed to uninstall.

### Workaround:

Verify whether the VRTSaa package was removed correctly from the alternate boot disk.

```
# pkginfo -R alternate_root_path -l VRTSaa
```

For example, run the following command

```
# pkginfo -R /altroot.5.10 -l VRTSaa
```

If the VRTSaa package was removed, you can ignore this error.

If the VRTSaa package was not removed, remove the package manually:

```
# pkgrm -R alternate_root_path -l VRTSaa
```

For example, run the following command

```
# pkgrm -R /altroot.5.10 -l VRTSaa
```

## After Live Upgrade to Solaris 10 Update 10, boot from alternate boot environment may fail (2370250)

If your setup involves volumes in a shared disk group that are mounted as CFS in a cluster, then during Live Upgrade using the `vxlustart` command from any supported Solaris version to Solaris 10 Update 10, boot from an alternate boot environment may fail.

**Workaround:** Run the `vxlufinish` command. Before rebooting the system, manually delete the entries of all the volumes of shared disks that are mounted as CFS in the `/altroot.5.10/etc/vfstab` directory.

## Live Upgrade to Solaris 10 Update 10 fails in the presence of zones (2521348)

SFCFSHA Live Upgrade from Solaris 10 Update 7 5.1SP1 to Solaris 10 Update 10 using the `vxlustart` commands fails in the presence of zones with the following error message:

```
ERROR: Installation of the packages from this media of the media failed;
pfinstall returned these diagnostics:
Processing default locales
    - Specifying default locale (en_US.ISO8859-1)
Processing profile
ERROR: This slice can't be upgraded because of missing usr packages for the
following zones:
ERROR:    zone1
ERROR:    zone1
ERROR: This slice cannot be upgraded because of missing usr packages for
```

one or more zones.

The Solaris upgrade of the boot environment <dest.27152> failed.

This is a known issue with the Solaris `luupgrade` command.

**Workaround:** Check with Oracle for possible workarounds for this issue.

## On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail (2424410)

On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail with the following error:

```
Generating file list.
```

```
Copying data from PBE <source.24429> to ABE <dest.24429>.
```

```
99% of filenames transferredERROR: Data duplication process terminated unexpectedly.
```

```
ERROR: The output is </tmp/lucreate.13165.29314/lucopy.errors.29314>.
```

```
29794 Killed
```

```
Fixing zonepaths in ABE.
```

```
Unmounting ABE <dest.24429>.
```

```
100% of filenames transferredReverting state of zones in PBE <source.24429>.
```

```
ERROR: Unable to copy file systems from boot environment <source.24429> to BE <dest.24429>.
```

```
ERROR: Unable to populate file systems on boot environment <dest.24429>.
```

```
Removing incomplete BE <dest.24429>.
```

```
ERROR: Cannot make file systems for boot environment <dest.24429>.
```

This is a known issue with the Solaris `lucreate` command.

**Workaround:** Check with Oracle for possible workarounds for this issue.

## Flash archive installed through JumpStart causes new system to go into maintenance mode on reboot (2379123)

If a Flash archive is created on a golden host with encapsulated root disks, when this Flash archive is installed onto another host through JumpStart, the new system may go to maintenance mode when you initially reboot it.

This problem is caused by the predefined root disk mirror in the Flash archive. When the archive is applied to a clone system, which may have different hard drives, the newly cloned system may get stuck at root disk mirroring during reboot.

**Workaround:** Create the Flash archive on a golden host with no encapsulated root disks. Run `vxunroot` to clean up the mirrored root disks before you create the Flash archive.

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

If you install or configure SFHA and then close the Web installer, if you have other browser windows open, the Web installer does not ask for authentication in the subsequent sessions. Since there is no option to log out of the Web installer, the session remains open as long as the browser is open on the system.

**Workaround:** Make sure that all browser windows are closed to end the browser session and subsequently log in again.

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

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

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

### **sfmh discovery issue when you upgrade your Veritas product to 6.0 (2622987)**

If a host is not reporting to any management server but sfmh discovery is running before you upgrade to 6.0, sfmh-discovery may fail to start after the upgrade.

**Workaround:**

If the host is not reporting to VOM, stop sfmh-discovery manually before upgrading to 6.0 by executing the following command on the managed host:

```
/opt/VRTSsfmh/adm/vxvmdiscovery-ctrl.sh stop
```

### **When you uninstall CommandCentral Storage Managed Host from a system where Veritas Storage Foundation 6.0 is installed, SF 6.0 reconfiguration or uninstallation fails (2631486)**

On a system where Veritas Storage Foundation (SF) 6.0 is installed, if you uninstall CommandCentral Storage (CCS) Managed Host (MH) using the installer script from the CCS media, the installer script removes the contents of `/opt/VRTSperl`. As a result, SF 6.0 reconfiguration or uninstallation using

```
/opt/VRTS/install/install_sf_product_name or  
/opt/VRTS/install/uninstall_sf_product_name fails, because the installer  
script removed the contents of /opt/VRTSperl.
```

**Workaround:** To uninstall CCS MH from a system where SF 6.0 is installed, before you perform the uninstallation, perform the procedure in the following CCS TechNote:

<http://www.symantec.com/business/support/index?page=content&id=HOWTO36496>

## Incorrect server names sometimes display if there is a clock synchronization issue (2627076)

When you install a cluster with the Web-based installer, you choose to synchronize your systems with an NTP server due to a clock synchronization issue, you may see the NTP server name in messages instead of your server names.

Workaround:

Ignore the messages. The product is still installed on the correct servers.

## Erroneous uninstallation error message when you install CommandCentral and Storage Foundation (2628165)

If you install a Veritas CommandCentral Management Server product on a Solaris machine, and then you try to install Storage Foundation software on this machine, you may see the following erroneous message that VRTSsfmh will be uninstalled:

```
CPI WARNING V-9-40-3866 The VRTSsfmh package on hostname will be uninstalled
```

Note that the system *hostname* is reporting

to the following management servers:

```
ccs://hostname
```

Workaround:

Ignore this erroneous message.

## Manual upgrade of VRTSvlic loses keyless product levels [2115662]

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. # pkgrm VRTSvlic

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

```
# pkgadd -d VRTSvlic.pkg
```

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

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

## Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

After upgrading from 5.1 to higher versions of VCS, some keyless licenses may be left in the system. As a result, you may see periodic reminders being logged if the VOM server is not configured.

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

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.

2. Set the product level to *NONE* with the command:

```
# vxkeyless set NONE
```

3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic:`

- Verify if the key has `VXKEYLESS` feature Enabled using the following command:

```
# vxlicrep -k <license_key> | grep VXKEYLESS
```

- Delete the key if and only if `VXKEYLESS` feature is Enabled.

---

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

---

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

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

## Upgrade or uninstallation of Veritas Storage Foundation HA may encounter module unload failures

When you upgrade or uninstall Veritas Storage Foundation HA, some modules may fail to unload with error messages similar to the following messages:

```
fdd failed to stop on node_name  
vxfs failed to stop on node_name
```

The issue may be observed on any one or all the nodes in the sub-cluster.

**Workaround:** After the upgrade or uninstallation completes, follow the instructions provided by the installer to resolve the issue.

## Secure WAC communication needs to be disabled explicitly [2392568]

If you have WACs communicating securely where VCS is configured in secure mode and if you disable the VCS security, the WAC where VCS security is disabled continues attempting to communicate securely without success. Therefore, you need to explicitly disable WAC security when you disable VCS security.

**Workaround:** No workaround. Secure WAC communication needs to be disabled explicitly.

## Web installer has no option to remove node from a cluster

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

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

## Flash Archive installation not supported if the target system's root disk is encapsulated

Symantec does not support SFHA installation using Flash Archive if the target system's root disk is encapsulated.

Make sure that the target system's root disk is unencapsulated before starting the installation.

## After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon you change the locale of nodes that use it. The vxconfig daemon starts at boot. If you have changed locale, you need to restart the daemon.

**Workaround:**

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

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

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

If the CVM protocol version does not upgrade successfully, upgrade the CVM protocol on the CVM master node.

#### **To upgrade the CVM protocol on the CVM master node**

- 1 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

- 2 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

### **Unable to stop some SFHA processes (2329580)**

If you install and start SFHA, but later configure SFHA using `installvcs`, some drivers may not stop successfully when the installer attempts to stop and restart the SFHA drivers and processes. The reason the drivers do not stop is because some dependent SFHA processes may be in the running state.

Workaround: To re-configure the product, use the corresponding `installproduct` command to re-configure the product. Otherwise some processes may fail to stop or start.

For example, use `installsfrac` to re-configure SFHA rather than using `installvcs`.

### **The Configure Sybase ASE CE Instance in VCS option creates duplicate service groups for Sybase binary mount points (2560188)**

The CPI installer does not check to see if Sybase binary mount points are already configured on systems, nor does it give an error message. It creates a duplicate service group for Sybase binary mount points.

This issue will be resolved in a later release.

## Veritas Dynamic Multi-pathing known issues

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

### Creating a zpool fails with a incorrect disk size error (2277875)

When the tunable parameter `dmp_native_support` is turned on, creating a zpool on DMP devices may fail with the following error:

```
one or more devices is less than the minimum size (64 M)
```

This error may occur even if the device size is greater than the required minimum size.

#### Work-around:

To resolve this issue, use one of the following commands:

- # `vxdisk scandisks`
- # `format -e dmp_device`

### DMP aggregates EFI labelled LUNS to a 0\_0 disk (2558408)

While performing `vxdiskunsetup` of some luns, if you format and label the disks as EFL, all the EFI labelled luns are aggregated to a 0\_0 disk.

#### Workaround:

When changing the label of a disk from SMI to EFI, or vice-versa, Symantec recommends that the label be changed on all accessible paths to a disk. That is, use the `format -e` command to stamp the new label on all accessible paths. For Active/Passive (A/P) class of arrays, this should be done only on the active paths. For other arrays, all paths should be labeled.

Symantec also recommends the installation of the patch provided by Oracle for EFI label issues (IDR144101-01 or IDR144249-01 or release kernel patch 142909-17). If this patch is installed, you can run the `format -e` command only on one path. After that, perform a read operation (such as `dd if=/dev/rdisk/<path> of=/dev/null count=1`) on the other accessible paths to propagate the label.

### Splitting a mirror from a zpool causes a core dump (2273367)

The following operation to split a mirror from a zpool fails:

```
# zpool split my_pool new_pool mirror
```

This issue is an Oracle issue with zpool. This issue occurs whether DMP is controlling the devices or not. That is, whether the `dmp_native_support` tunable is on or off.

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

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

### **Suppressing the primary path of an encapsulated SAN boot disk from Veritas Volume Manager causes the system reboot to fail (1933631)**

If you suppress the primary path of an array from VxVM control and then reboot the system, the system boot fails.

If you have an encapsulated SAN boot device with multiple primary paths, the issue occurs when you suppress the first primary path. When you configure a SAN boot device, the primary path is set as a boot device. In general, the first path of the SAN boot device corresponds to the first configured path during SAN boot. Even if another primary path is configured as a boot device, suppressing the first device from VxVM causes the boot to fail.

#### **Workaround:**

When the boot device is suppressed from VxVM, change the OS boot device sequencing accordingly.

For Solaris SPARC system, use the `eeeprom boot-device` command to set the boot device sequencing.

For Solaris x86-64 systems, use the `eeeprom bootpath` command to set the boot device sequencing.

### **Changes in enclosure attributes are not persistent after an upgrade to VxVM 6.0 RP1 (2082414)**

The Veritas Volume Manager (VxVM) 6.0 RP1 includes several array names that differ from the array names in releases prior to release 5.1SP1. Therefore, if you upgrade from a previous release to VxVM 6.0 RP1, changes in the enclosure

attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

[Table 1-9](#) shows the Hitachi arrays that have new array names.

**Table 1-9** Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC
TagmaStoreUSPV	Hitachi_USP-V
TagmaStoreUSPVM	Hitachi_USP-VM
<New Addition>	Hitachi_R700
Hitachi AMS2300 Series arrays	New array names are based on the Model Number 8x. For example, AMS_100, AMS_2100, AMS_2300, AMS_2500, etc.

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

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

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

## ZFS pool creation on a DMP device fails when the LUN size is between 1 TB and 2TB (2010919)

Creating a ZFS pool on a DMP device using the whole disk of size > 1TB and < 2TB that contains a SMI SUN label fails. The issue is that `zpool create` on a whole disk

changes the device label from SMI to EFI. This causes confusion between the OS device paths of the same DMP device due to a bug in the Sun SCSI layer. This is due to SUN BugID: 6912703.

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

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

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

## **Issues with DMP upgrade if enclosure-based naming scheme is changed (2632422)**

Veritas Dynamic Multi-Pathing (DMP) support for native devices requires that the naming scheme be set to enclosure-based naming (EBN). DMP 6.0 does not allow changing the naming scheme from EBN when support for native devices is enabled.

Due to a bug in DMP 5.1 Service Pack 1 (5.1SP1), the naming scheme could be set to operating-system based naming (OSN). However, this is not a supported configuration.

Before upgrading DMP 5.1SP1 to DMP 6.0, make sure that the naming scheme is set to enclosure-based naming.

```
# vxddladm get namingscheme
```

If the naming scheme displays as OSN, set the naming scheme to EBN.

```
# vxddladm set namingscheme=ebn
```

## **After changing the preferred path from the array side, the secondary path becomes active (2490012)**

For EVA arrays, DMP requires that the prefer bit is static. If the prefer bit is not static, issues like the following may occur. After changing the prefer path of LUN from the array side, and performing a disk discovery (`vxdisk scandisks`) from the host, the secondary path becomes active for the LUN.

**Work-around:**

**To work around this issue**

- 1 Set the pref bit for the LUN.
- 2 Perform disk discovery again:

```
# vxdisk scandisks
```

## Veritas Storage Foundation known issues

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

### Veritas Storage Foundation known issues

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

#### **Not all the objects are visible in the VOM GUI (1821803)**

After upgrading SF stack from 5.0MP3RP2 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the SFM GUI.

**Workaround:****To resolve this known issue**

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

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

#### **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 SmartTier placement policy if you do not tag the volume with the placement classes prior to constructing a volume set for the volume.

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

## NULL pointer dereference panic with Solaris 10 Update 10 on x86 and Hitachi Data Systems storage (2616044)

Due to a limitation with Solaris 10 Update 10 on x86, when the server is connected to Hitachi Data storage, the system panics due to a NULL pointer dereference during the boot cycle with the following stack trace:

```
fffffe8000988570 unix:die+da ()
fffffe8000988650 unix:trap+5e6 ()
fffffe8000988660 unix:cmntrap+140 ()
fffffe8000988870 scsi_vhci:hds_sym_path_get_opinfo+62 ()
fffffe8000988920 scsi_vhci:vhci_update_pathinfo+5b ()
fffffe80009889a0 scsi_vhci:vhci_pathinfo_online+2df ()
fffffe8000988a10 scsi_vhci:vhci_pathinfo_state_change+202 ()
fffffe8000988a70 genunix:i_mdi_pi_state_change+148 ()
fffffe8000988ab0 genunix:mdi_pi_online+32 ()
fffffe8000988b20 fcp:ssfcp_online_child+ff ()
fffffe8000988b90 fcp:ssfcp_trigger_lun+2b0 ()
fffffe8000988bc0 fcp:ssfcp_hp_task+88 ()
fffffe8000988c40 genunix:taskq_thread+295 ()
fffffe8000988c50 unix:thread_start+8 ()
```

For more information, see Oracle bug ID 7079724.

**Workaround:** Disable Solaris I/O multipathing on the server to avoid the system panic.

### To disable Solaris I/O multipathing on the server

- 1 Disable Solaris I/O multipathing:

```
# stmsboot -d
```

- 2 Reboot the server:

```
# reboot
```

## Veritas Volume Manager known issues

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

### Tunable values doesn't change as per the values applied through vxtune (2698035)

The `vxtune` command displays tunable values in terms of default unit. When a new value is provided to a tunable, `vxtune` accepts it in terms of default unit. For example, `voldrl_min_regionsz` is stored in terms of blocks (OS block size). So if

a value of 1024 is provided, then `vxtune` interprets it as 1024 blocks; if a value of 2M is provided, it interprets it as 2M blocks (2097152 blocks), not 4096 blocks (2MB).

**Workaround:**

It is recommended to provide exact tunable value without any suffix. This will be addressed in future releases.

**Expanding a LUN to a size greater than 1 TB fails to show correct expanded size (2123677)**

This issue occurs when you perform a Dynamic LUN Expansion for a LUN that is smaller than 1 TB and increase the size to greater than 1 Tb. After the expansion, Veritas Volume Manager (VxVM) fails ongoing I/O, and the public region size is reset to original size. After you run the `vxdisk scandisks` command, VxVM does not show the correct expanded size of the LUN. The issue is due to underlying Solaris issues. Refer to Sun Bug Id 6929449 and Sun Bug Id 6912703.

**Workaround:** There is no workaround.

**1 TB luns goes in error state with Solaris x86 (2706776)**

If you label a disk device as EFI using `format` on a subset of the paths or on the DMP device, Solaris will not be able to propagate the label to all the other paths of the LUN. This will lead the device to appear in the error state under `'vxdisk list'`.

**Workaround:** There is no workaround for this issue.

**The failure of one disk causes all the sub-disks to be relocated (2641510)**

In a campus cluster environment, failure of one disk causes all the sub-disks, on the storage devices tagged with the specific site, to be relocated.

Ideally, the disks that encounter an I/O failure only are marked for hot-relocation (RLOC). However, the `vxrelocd` daemon does not try to recover the volumes that become available after the devices at the specific sites are partially reattached. It checks if the sub-disks/plexes are still detached and then relocates all the sub-disks.

There is no workaround for this issue.

**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.

**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
```

**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
```

### Co-existence check might fail for CDS disks

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

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

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

**Workaround:** There is no workaround for this issue.

### 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
```

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

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

**Workaround:** Rerun the `vxsnap addmir` command.

### Probing vxio with DTrace fails on Sparc machines. (2180635)

This issue exists because of inability of DTrace to load a module whose text size is greater than 2MB on Sparc machines. While trying to load `vxio` with DTrace you may see following warning messages on console:

```
dtrace: WARNING: couldn't allocate SDT table for module vxio  
fbt: WARNING: couldn't allocate FBT table for module vxio
```

There is no workaround for this issue.

### **The vxassist maxsize option fails to report the maximum size of the volume that can be created with given constraints when the disk group has the siteconsistent flag set (2563195)**

The vxassist maxsize option fails to report the maximum size of volume that can be created with given constraints when the disk group has the siteconsistent flag set. The following error is reported:

```
# vxassist -g dgroup maxsize  
VxVM vxassist ERROR V-5-1-752 No volume can be created within the given  
constraints
```

#### **Workaround:**

Specify the size explicitly to the vxassist make command.

### **Disks on the LDOM guest are claimed under other\_disks category (2354005)**

The disks on the LDOM guest are claimed under "other\_disks" enclosure, because these disks are not capable of being multi-pathed by DMP. This is expected because these devices represent VxVM volumes in the host. By design, devices under other\_disks enclosure have their name based on underlying OS path regardless of the DDL naming scheme.

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

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

#### **Workaround:**

To configure path-level attributes

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

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

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In SFHA 6.0, if this command is run while the volumes are online and being actively used, the shown %dirty may lag from actual percentage dirty for instant snap data change object (DCO) volumes. That is, the command output may show less %dirty than actual.

### Encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path names are not the same (2607706)

The encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path name are not the same.

For example:

Dmpnode:sdh

Paths: sda sdb

#### Work-around:

Before running the encapsulation command (`vxencap`), run the following command:

```
# vxddladm assign names
```

### The vxdisksetup command fails to initialize disks in cdsdisk format for LDOM disks greater than 1 TB (2557072)

The `vxdisksetup` command fails to initialize disks in `cdsdisk` format for LDOM disks greater than 1 TB. This issue is due to an LDOM operating system command which fails when the number of partitions in the GUID partition table (GPT) label is greater than 9. The `cdsdisk` format requires at least 128 partitions to be compatible with Linux systems.

**Workaround:** There is no workaround for this issue.

### pkgchk displays errors for VRTSaslapm package after upgrading (2616459)

If you upgrade from the Veritas Storage Foundation 5.1 SP1 release to the 6.0 release using the installer, the `pkgchk -n VRTSaslapm` command displays the following errors:

```
# pkgchk -n VRTSaslapm
ERROR: /etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.10
      pathname does not exist
ERROR: /etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.9
      pathname does not exist
ERROR: /etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.10
```

```
pathname does not exist
ERROR: /etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.9
pathname does not exist
ERROR: /kernel/drv/vxapm/dmpnetapp.SunOS_5.10
pathname does not exist
ERROR: /kernel/drv/vxapm/dmpnetapp.SunOS_5.9
pathname does not exist
ERROR: /kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
pathname does not exist
ERROR: /kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
pathname does not exist
ERROR: /usr/lib/vxvm/root/kernel/drv/vxapm/dmpnetapp.SunOS_5.10
pathname does not exist
ERROR: /usr/lib/vxvm/root/kernel/drv/vxapm/dmpnetapp.SunOS_5.9
pathname does not exist
ERROR: /usr/lib/vxvm/root/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
pathname does not exist
ERROR: /usr/lib/vxvm/root/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
pathname does not exist
```

These errors display because the installer installs the new `VRTSaslapm` packages without first removing the old packages. This causes the `/var/sadm/install/contents` directory to contain some old packages, which cause the errors.

**Workaround:**

Perform a fresh installation of the Veritas Storage Foundation 6.0 release.

## Veritas File System known issues

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

### A cluster mounted file system may cause system panic showing `vx_tflush_map` in the stack trace (2558892)

VxFS causes the server to panic. The subroutine initiating the panic is `vx_tflush_map`.

There is no workaround for this issue.

### Full file system check takes over a week (2628207)

On a large file system with many checkpoints, a full file system check using the `fsck_vxfs(1m)` command may appear to be hung.

There is no workaround for this issue.

### **umount(1m) on a CFS file system causes system panic (2107152)**

In rare corner cases, the system panics while unmounting a cluster mounted file system.

There is no workaround for this issue.

### **fscckptadm(1m) fails with ENXIO (1956458)**

The `fscckptadm(1m)` fails with the ENXIO error and the file system is marked for full file system check.

There is no workaround for this issue.

### **vxfsd consumes a lot of CPU resources after deleting some directories (2129455)**

The `vxfsd` daemon consumes 100% CPU after some directories are deleted.

There is no workaround for this issue.

### **vxfsckd resource fails for start when a cluster node is rebooted or when it is killed manually (2720034)**

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

#### **Workaround:**

Use the following commands for this situation:

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

Kill all `vxfsckd` processes:

```
fsclustadm cfsdeinit  
hastart
```

### **fsdedupadm status shows FAILED on already deduped filesystem (2715413)**

If a file system is already deduped, then after dryrun `fsdedupadm` status shows FAILED all time.

**Workaround:** Only do the dryrun on a fresh file system. Or If you do dryrun on an already deduped file system, then start dedup again:

```
fsdedupadm start mountpoint
```

### Taking a FileSnap over NFS multiple times with the same target name can result in the 'File exists' error (2353352)

The "File exists" error occurs as a result of the caching behavior of the NFS client. Because the link operation is successful, the NFS client assumes that a file with the specified target name, such as `file2::snap:vxfs:`, was created.. As a result, the NFS client caches a file with this name.

**Workaround:** Remove the target file after a snapshot is created. This forces the NFS client to remove the name from the cache. For example:

```
# ln file1 file2::snap:vxfs:
# rm file2::snap:vxfs:
```

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

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

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

### A mutex contention in `vx_worklist_lk()` can use up to 100% of a single CPU (2086902)

A mutex contention in the `vx_worklist_lk()` call can use up to 100% of a single CPU.

**Workaround:** There is no workaround for this issue.

### Deduplication can fail with error 110 (2591473)

In some cases, data deduplication fails with a message similar to the following example:

```
Saving      Status      Node           Type           Filesystem
-----
00%         FAILED      node01         MANUAL         /data/fs1
                2011/10/26 01:38:58 End full scan with error
```

In addition, the deduplication log contains an error similar to the following example:

```
2011/10/26 01:35:09 DEDUP_ERROR AddBlock failed. Error = 110
```

These errors indicate that the deduplication process is running low on space and needs more free space to complete.

**Workaround:** Make more space available on the file system.

### **vxresize fails while shrinking a file system with the "blocks are currently in use" error (2437138)**

The `vxresize` shrink operation may fail when active I/Os are in progress on the file system and the file system is being shrunk to a size closer to its current usage. You see a message similar to the following example:

```
UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink /dev/vx/rdisk/dg1/voll -  
blocks are currently in use.  
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume  
voll, in diskgroup dg1
```

**Workaround:** Rerun the shrink operation after stopping the I/Os.

### **Warning message sometimes appear in the console during system startup (2354829)**

During system startup, following messages sometimes appear in system console:

```
WARNING: couldn't allocate SDT table for module vxfs  
WARNING: couldn't allocate FBT table for module vxfs  
Loading smf(5) service descriptions: 2/2
```

These warnings indicate that the SDT and FBT DTrace probes might not be available for the VxFS module. The VxFS module still loads and works correctly. Dtrace SDT/FBT has limits on the size of module that it can support. Since the VxFS module exceeds the size that Dtrace can support, SDT and FBT Dtrace probes might not work for VxFS.

**Workaround:** There is no workaround for this issue.

## **Replication known issues**

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

### **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 voll fe80::221:5eff:fe49:ad10:dg1:voll  
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.

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

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

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

**A snapshot volume created on the Secondary, containing a VxFS file system may not mount in read-write mode and performing a read-write mount of the VxFS file systems on the new Primary after a global clustering site failover may fail (1558257)****Issue 1:**

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

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

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

**Issue 2:**

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

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

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

**Workaround:** The following workarounds resolve these issues.

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

For example:

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

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

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

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

When VVR is used for replicating shared disk groups in 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:

```
vradmin not reachable on cluster peer
```

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

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

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

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

```
# /lib/svc/method/vras-vradmind.sh stop
# /lib/svc/method/vras-vradmind.sh start
```

### While `vradmin` commands are running, `vradmind` may temporarily lose heart beats (2162625, 2275444)

This issue may occasionally occur when you use `vradmin` commands to administer VVR. While the `vradmin` commands run, `vradmind` may temporarily lose heartbeats, and the commands terminate with the following error message:

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

### Workaround:

#### To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

### 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.

### **Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (2478684)**

Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (DCM), even if you have enough disk space.

**Workaround:** Add a LUN to the diskgroup before creating the primary diskgroup.

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

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

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

**Workaround:** There are two workarounds for this issue.

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

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

When VVR is used for asynchronous replication in shared disk group environment, one of the nodes of the cluster at the primary site is chosen as the logowner. When the logowner node is on a node which is a slave node for the underlying CVM cluster, in the presence of heavy I/O from a node that is not the logowner, it is possible to get into a replication hang. This is due to an internal defect which will be fixed in later releases.

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

### 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.

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

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

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

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

There is no workaround for this issue.

## Secondary-host related operations may fail in a pure IPv6 environment (2625708)

In a pure IPv6 configuration, using the `--secondary_host` option with `vxsfadm` or the equivalent `-t` option for `dbed_vmchecksnap` may fail with the following error message if name resolution is configured to use `/etc/hosts` rather than DNS:

```
SFDB vxsfadm ERROR V-81-0161 The given secondary host <hostname>
is not reachable.
```

Action: Specify a different secondary host, or remove the `-t` option if the secondary host is not running at this time.

### Workaround

Use one of the following workarounds:

- Configure name resolution to use DNS instead of a hosts file, if a DNS server is available.
- Do not use the `--secondary_host` option.  
This option is not mandatory for off-host operations. Off-host operations are fully supported without specifying this option.

## `dbed_update` or `vxsfadm` may fail in a pure IPv6 environment (2619958)

In a pure IPv6 environment, the `dbed_update` or the `vxsfadm` command may fail with the following error:

```
SFDB vxsfadm ERROR V-81-0685
Oracle instance SIDB does not seem to be running
```

### Workaround

Ensure that an IPv4 address is configured for localhost. This IP does not need to be reachable from other hosts.

#### To configure an IPv4 localhost address

- ◆ Run the following command as root:

```
# ifconfig lo0 plumb
```

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

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

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

### Workaround

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

### Incorrect error message if wrong host name is provided (2585643)

If you provide an incorrect host name with the `-r` option of `vxsfadm`, the command fails with an error message similar to one of the following:

```
FSM Error: Can't use string ("") as a HASH ref while "strict refs"  
in use at /opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 776.
```

```
SFDB vxsfadm ERROR V-81-0609 Repository location is invalid.
```

The error messages are unclear.

### Workaround

Provide the name of a host that has the repository database, with the `-r` option of `vxsfadm`.

### FlashSnap validate reports snapshot unsplitable (2534422)

The FlashSnap validation operation fails with the following error if the mirrors for data volumes and archive log volumes share the same set of disks:

```
SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.
```

### Workaround

Ensure that snapshot plexes for data volumes and snapshot plexes for archive log volumes reside on separate set of disks.

### Attempt to use SmartTier commands fails (2332973)

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

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

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

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

#### **`dbed_vmclonedb` ignores new clone SID value after cloning once (2580318)**

After you have done FlashSnap cloning using a snapplan, any further attempts to create a clone from the same snapplan using the `dbed_vmclonedb` continue to use the original clone SID, rather than the new SID specified using the `new_sid` parameter.

This issue is also observed when you resynchronize the snapplan, take a snapshot again without specifying the new clone SID, and then try to clone with the new SID.

#### **Workaround**

You can use one of the following workarounds:

- After the snapshot is resynchronized, delete the snapplan using the `dbed_vmchecksnap -o remove` command. You can then use a new clone SID by creating a new snapplan, which may have the same name, and using the snapplan for taking more snapshots.
- Use the `vxsfadm` command to take the snapshot again and specify the clone SID with the snapshot operation so that the clone operation can be done with the new clone SID.

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

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

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

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

- BALANCE
- CHECKPOINT
- METADATA

#### **Workaround**

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

#### **User authentication fails (2579929)**

The `sfae_auth_op -o auth_user` command, used for authorizing users, fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username>
```

Reattempting the operation fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0372 AT broker failed to start:
```

The authentication setup might have been run with a strict umask value, which results in the required files and directories being inaccessible to the non-root users.

### Workaround

If you have not done authentication setup, set umask to a less strict value before running the `sfae_auth_op -o setup` or `sfae_auth_op -o import_broker_config` commands.

To set umask to a less strict value

- ◆ Use the command:

```
# umask 022
```

If you have already done authentication setup, perform the following steps.

To resolve the problem if you have already done authentication setup

- 1 Shut down the authentication broker, if it is running.

```
# /opt/VRTSdbed/at-broker/bin/sfaeatd.sh stop
```

- 2 Change the permissions for files and directories that are required to be readable by non-root users.

```
# chmod o+r /etc/vx/vxdbed/admin.properties
```

```
# chmod o+rx /var/vx/vxdba/auth/users
```

```
# find /opt/VRTSdbed/at-broker -type d -exec chmod o+rx {} \;
```

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

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

### Workaround

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

- For FlashSnap, resync the snapshot and try the clone operation again.

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

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

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

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

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

#### **Workaround**

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

### **Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 RP1 (2184482)**

The `sfua_rept_migrate` command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.0 RP1.

When upgrading from SFHA version 5.0 to SFHA 6.0 RP1 the `S*vxdbs3` startup script is renamed to `NO_S*vxdbs3`. The `S*vxdbs3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbs3` startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbs3 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*vxdbs3` to `S*vxdbs3`.

### **Clone command fails if PFILE entries have their values spread across 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.

## **Veritas Cluster Server known issues**

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

### **You may be unable to configure the MultiNICB agent to perform probe-based detection of resource state (2509227)**

You can ordinarily configure the MultiNICB agent to perform link-based detection and probe-based detection by setting the LinkTestRatio attribute to '1' and the IgnoreLinkStatus attribute to '0'. However, when you set these values, the agent may fail to send out ICMP requests to determine the resource state. As a result, the agent may report an erroneous resource state and probe-based detection may fail.

#### **Workaround:**

Set the LinkTestRatio attribute to '2'. MultiNICB agent now sends an ICMP request after every alternate monitor cycle, and is able to determine the correct resource state.

### **An event registered with AMF still remains registered after you apply or remove the VRTSamf patch (2700494)**

If you add or remove the VRTSamf patch (version 6.0RP1) either manually or through CPI, an event which is registered with AMF outside of VCS control (not through VCS agent), remains registered even after the addition or removal of the patch.

#### **Workaround:**

- If you are monitoring a resource outside of VCS control, you need to unregister the resource from AMF before you apply or remove the VRTSamf patch.
- Apply or remove the VRTSamf patch only by using CPI.

### **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.

## Operational issues for VCS

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

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

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

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the *Veritas Cluster Server Installation Guide* for the list of ports and services used by VCS.
- Configure the firewall policy such that the TCP ports required by VCS are not blocked. Refer to your respective firewall or OS vendor documents for the required configuration.

## Issues related to the VCS engine

### Extremely high CPU utilization may cause HAD to fail to heartbeat to GAB

When CPU utilization is very close to 100%, HAD may fail to heartbeat to GAB. [1818687]

### Missing host names in engine\_A.log file

The GUI does not read the engine\_A.log file. It reads the engine\_A.ldf file, gets the message id from it, and then queries for the message from the bmc file of the appropriate locale (Japanese or English). The bmc file does not have system names present and so they are read as missing. [1736295]

### 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.

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

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

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

### **Service group is not auto started on the node having incorrect value of EngineRestarted [2397532]**

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

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

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

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

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

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

### **NFS resource goes offline unexpectedly and reports errors when restarted [2490404]**

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

Workaround: Forcefully stop the agent process.

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

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

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

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

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

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

**If secure and non-secure WAC are connected the engine\_A.log receives logs every 5 seconds [1539646]**

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

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

**Oracle group fails to come online if Fire Drill group is online on secondary cluster [2556835]**

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

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

**POSTONLINE and POSTOFFLINE triggers are not enabled by default [2567387]**

Before VCS 6.0, POSTONLINE and POSTOFFLINE triggers were enabled by default, so the triggers got executed whenever a service group came online. In VCS 6.0, you must explicitly enable the POSTONLINE and POSTOFFLINE triggers whenever you upgrade to VCS 6.0.

**Alternatively, if you want the triggers to execute after the upgrade:**

- 1 so that HAD does not start after the upgrade.
- 2 Upgrade the existing VCS to VCS 6.0.
- 3
- 4 Start VCS on each node using `hastart`.
- 5 Set `TriggersEnabled` in `main.cf` for required groups as follows:

```
TriggersEnabled @<systemname>={POSTONLINE, POSTOFFLINE}
```

### Example of trigger behavior:

```
group scriptfileonoff (
    SystemList = { vcssx235 = 0, vcssx236 = 1 }
    AutoStartList = { vcssx235, vcssx236 }
    TriggersEnabled @vcssx235 = { POSTONLINE }
)
MyFileOnOff MFileOnOff (
    PathName = "/tmp/mf1"
)
MyFileOnOff MFileOnOff1 (
    PathName = "/tmp/mf2"
```

### Two CmdServer instances seen running on a node [2399292]

You may see two instances of CmdServer running on a node. One of these using IPv4 and the other IPv6.

This does not impact functionality in any way.

Workaround: No workaround.

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

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

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

## Issues related to the bundled agents

### Entry points that run inside a zone are not cancelled cleanly [1179695]

Cancelling entry points results in the cancellation of only the `zlogin` process. The script entry points that run inside a zone are forked off using the `zlogin` command. However, the `zlogin` command forks off an `sh` command, which runs in the context of the Solaris zone. This shell process and its family do not inherit the group id of the `zlogin` process, and instead get a new group id. Thus, it is difficult for the agent framework to trace the children or grand-children of the shell process, which translates to the cancellation of only the `zlogin` process.

Workaround: Oracle must provide an API or a mechanism to kill all the children of the `zlogin` process that was started to run the entry point script in the local-zone.

## Solaris mount agent fails to mount Linux NFS exported directory

The Solaris mount agent mounts the mount directories. At this point, if it tries to mount a Linux NFS exported directory, the mount fails showing the following error:

```
nfs mount: mount: <MountPoint>: Not owner
```

This is due to system NFS default version mismatch between Solaris and Linux.

The workaround for this is to configure `MountOpt` attribute in mount resource and set `vers=3` for it.

### Example

```
root@north $ mount -F nfs south:/test /logo/
nfs mount: mount: /logo: Not owner
root@north $
Mount nfsmount (
    MountPoint = "/logo"
    BlockDevice = "south:/test"
    FSType = nfs
    MountOpt = "vers=3"
)
```

## The zpool command runs into a loop if all storage paths from a node are disabled

The Solaris Zpool agent runs `zpool` commands to import and export zpools. If all paths to the storage are disabled, the `zpool` command does not respond. Instead, the `zpool` export command goes into a loop and attempts to export the `zpool`. This continues till the storage paths are restored and `zpool` is cleared. As a result, the offline and clean procedures of Zpool Agent fail and the service group cannot fail over to the other node.

Workaround: You must restore the storage paths and run the `zpool clear` command for all the pending commands to succeed. This will cause the service group to fail over to another node.

## Application agent cannot handle a case with user as root, envfile set and shell as csh [2584285]

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

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

### **IMF registration fails for Mount resource if the configured MountPoint path contains spaces [2442598]**

If the configured MountPoint of a Mount resource contains spaces in its path, then the Mount agent can online the resource correctly, but the IMF registration for ONLINE monitoring fails. This is due to the fact that the AMF driver does not support spaces in the path. Leading and trailing spaces are handled by the Agent and IMF monitoring can be done for such resources.

Workaround: Symantec recommends to turn off the IMF monitoring for a resource having spaces in its path. For information on disabling the IMF monitoring for a resource, refer to Veritas Cluster Server Administrator's Guide.

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

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

Workaround:

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

### **Concurrency violation in the service group [2555306]**

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

This happens when:

- In a cluster environment/configuration, if cluster wide UseFence attribute is set to SCSI3 and service group contains Volume resource and DiskGroup resource with the PanicSystemOnDGLoss attribute set to 0 (zero).
- If storage connectivity is lost or all paths under VxDMP are disabled, VCS fails over the service group. If storage connectivity is restored on the node on which the service group was faulted and DG is not deported manually, then volume may get started if disk group is not deported during the service group failover.

So volume resource shows state as online on both the nodes and thus cause concurrency violation. This may lead to data corruption.

Workaround: Ensure that the disk group is deported soon after storage connectivity is restored.

You are recommended to always configure Volume resource whenever Disk group resources is configured and set the attribute PanicSystemOnDGLoss to 1 or 2 as per requirement.

### **Coordpoint agent remains in faulted state [2555191]**

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

Workaround: Clear the fault and reconfigure fencing.

### **No IPv6 support for NFS [2022174]**

IPv6 is not supported for NFS.

Workaround: No workaround.

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

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

## **Issues related to the VCS database agents**

### **Health check monitoring does not work with VCS agent for Oracle [2101570, 1985055]**

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

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

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

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

### **Make sure that the ohasd has an entry in the init scripts [1985093]**

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.

**Workaround:** Respawn off the `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.

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

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

**Workaround:**

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

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

The ASMInst agent does not support health check monitoring.

**Workaround:** Set the MonitorOption attribute to 0.

### **NOFAILOVER action specified for certain Oracle errors**

The Veritas High Availability agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file `oraerror.dat`, which consists of a list of Oracle errors and the actions to be taken.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for a description of the actions.

Currently, the reference file specifies the NOFAILOVER action when the following Oracle errors are encountered:

ORA-00061, ORA-02726, ORA-6108, ORA-06114

The NOFAILOVER action means that the agent sets the resource's state to OFFLINE and freezes the service group. You may stop the agent, edit the `oraerror.dat` file, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

### **ASM instance does not unmount VxVM volumes after ASMDG resource is offline**

In configurations where ASMInstance resource is part of a separate parallel service group, the ASM instance does not unmount the volumes even after the ASMDG resource is taken offline. Therefore, the Volume resource cannot be taken offline. This issue occurs when you use VxVM volumes as ASM disk groups. [918022]

Workaround: Configure the ASMInstance resource as part of the failover service group where ASMDG resource is configured.

## 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.

### The agent framework does not detect if service threads hang inside an entry point [1511211]

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

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

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

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

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

Workaround: No workaround.

## Issues related to VCS in Japanese locales

This section covers the issues that apply to SFHA in a Japanese locale.

### The `hares -action` command displays output in English [1786742]

The `hares -action` command incorrectly displays output in English.

## Issues related to global clusters

### **The engine log file receives too many log messages on the secure site in global cluster environments [1539646]**

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

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

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

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

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

## 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 (1788315)**

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

### **LLT may incorrectly declare port-level connection for nodes in large cluster configurations (1809827)**

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.

## Issues related to GAB

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

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

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

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

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

### Cluster panics during reconfiguration (2590413)

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

## Issues related to I/O fencing

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

### Delay in rebooting Solaris 10 nodes due to vxfen service timeout issues (1897449)

When you reboot the nodes using the `shutdown -i6 -g0 -y` command, the following error messages may appear:

```
svc:/system/vxfen:default:Method or service exit
timed out. Killing contract 142
svc:/system/vxfen:default:Method "/lib/svc/method/vxfen stop"
failed due to signal Kill.
```

This error occurs because the `vxfen` client is still active when VCS attempts to stop I/O fencing. As a result, the `vxfen stop` service times out and delays the system reboot.

**Workaround:** Perform the following steps to avoid this `vxfen stop` service timeout error.

### To avoid the vxfen stop service timeout error

- 1 Stop VCS. On any node in the cluster, run the following command:

```
# hastop -all
```

- 2 Reboot the systems:

```
# shutdown -i6 -g0 -y
```

### CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

See the *Veritas Storage Foundation and High Availability Administrator's Guide* for more details.

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

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

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

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

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

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

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

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily, or if the node is a single-node VCS configuration where LLT is not configured, then the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

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

### **When I/O fencing is not up, the svcs command shows VxFEN as online (2492874)**

Solaris 10 SMF marks the service status based on the exit code of the start method for that service. The VxFEN start method executes the `vxfen-startup` script in the background and exits with code 0. Hence, if the `vxfen-startup` script subsequently exits with failure then this change is not propagated to SMF. This behavior causes the `svcs` command to show incorrect status for VxFEN.

**Workaround:** Use the `vxfenadm` command to verify that I/O fencing is running.

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

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

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

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

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

```
operation failed.  
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@node1,  
domaintype vx; not allowing action
```

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

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

### **The `vxfsnwap` utility does not detect failure of coordination points validation due to an RSH limitation (2531561)**

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

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

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

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

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

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

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

**Workaround:** Start VxFEN again after some time.

### The `cpsadm` command fails after upgrading CP server to 6.0 in secure mode (2478502)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTS`at` package is not removed from the system, the `cpsadm` command loads the old security libraries present on the system. As the installer runs the `cpsadm` command on the CP server to add or upgrade the SFHA cluster (application cluster), the installer also fails.

Workaround : Perform the following steps on all the nodes of the CP server:

- Rename `cpsadm` to `cpsadmbin`.

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

- Create a file `/opt/VRTScps/bin/cpsadm` with the following content:

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

- Provide the following permissions to the new file:

```
# chmod 755 /opt/VRTScps/bin/cpsadm
```

### Server-based fencing comes up incorrectly if default port is not mentioned (2403453)

When you configure fencing in customized mode and do not provide default port, fencing comes up. However, the `vxfenconfig -l` command output does not list the port numbers.

Workaround: Retain the "`port=<port_value>`" setting in the `/etc/vxfenmode` file, when using customized fencing with at least one CP server. The default port value is 14250.

### Secure CP server does not connect from localhost using 127.0.0.1 as the IP address (2554981)

The `cpsadm` command does not connect to the secure CP server on the localhost using 127.0.0.1 as the IP address.

Workaround: Connect the secure CP server using any of the virtual IPs that is configured with the CP server and is plumbed on the local node.

### Unable to customize the 30-second duration (2551621)

When the `vxcpsserv` process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: No workaround.

### NIC resource gets created with incorrect name while configuring CPSSG with the `configure_cps.pl` script (2585229)

The name of the NIC resource created by the `configure_cps.pl` script does not come out correct when, for example,  $m^{\text{th}}$  VIP is mapped to  $n^{\text{th}}$  NIC and every  $m$  is not equal to  $n$ . In this case, although CPSSG continues to function without any problem, when you unconfigure CPSSG using `configure_cps.pl`, it fails.

Workaround: To unconfigure CPSSG, you must remove the CPSSG configuration from the VCS configuration.

### CP server configuration fails while setting up secure credentials for CP server hosted on an SFHA cluster (2621029)

When you configure CP server using the `configure_cps.pl` utility, the configuration fails while setting up secure credentials for CP server that is hosted on an SFHA cluster. You may see the following error:

```
Creating softlink to credential directory /etc/VRTSvcs/db/CPSEVER
on node nodename.
Unable to connect to node nodename using /usr/bin/ssh.
Please configure ssh communication and retry. Exiting.
```

Workaround: You can use any of the following options:

- Before running the `configure_cps.pl` utility, change the default shell for root user to either `KSH` or `bash`.
- Perform the following steps after running the `configure_cps.pl` utility on each node of the cluster:

- Manually remove the old credential directory or softlink. For example:

```
# rm -rf /var/VRTSvcs/vcsauth/data/CPSEVER
```

- Create a new soft-link to the shared location of the credential directory:

```
# ln -s path_of_CP_server_credential_directory \
/var/VRTSvcs/vcsauth/data/CPSEVER
```

- Start the CPSSG service group:

```
# hagr -online CPSSG -any
```

## Issues related to Intelligent Monitoring Framework (IMF)

### Registration error while creating a Firedrill setup [2564350]

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

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

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

Workaround: No workaround.

### Pearl errors seen while using `haimfconfig` command

Pearl errors seen while using `haimfconfig` command:

```
Pearl errors seen while using haimfconfig command
```

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

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

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

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

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

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

## Issues related to the Cluster Manager (Java Console)

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

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

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

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

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

### **VCS Cluster Manager (Java Console) does not encrypt Sybase and SybaseBk agent passwords [2379510]**

If `isvcsagentcrypt` flag is set to `True` in `Sybase.xml` and `SybaseBk.xml` files, the attribute values get encrypted. However, the password attributes of Sybase and SybaseBk agents do not have the `isvcsagentcrypt` flag set to `True` in `Sybase.xml` and `SybaseBk.xml` files.

**Workaround:** Sybase and SybaseBk agents are modified to encrypt the password by default. As a result, you need not encrypt passwords if you use the VCS Cluster Manager (Java Console) to configure attributes.

## **Issues related to Virtual Business Services (VBS)**

### **Child service group fault missed in case of firm dependency (2673289)**

In case of firm dependency, if the child service group has faulted, the parent service group is brought offline. Subsequently when the child service group recovers, the parent service group is brought online.

However, if the child service group faults again before the parent has recovered, this new child fault event is missed. This could happen if the parent recovery takes time. As a result, the parent service group may remain online.

**Workaround:** There is no workaround.

### **VBS version mismatch across tiers may affect its functionality (2695412)**

If you have different versions of VBS across various tiers, VBS may not function as expected. Therefore, you must ensure that when you upgrade a particular tier of VBS, you must also upgrade the other tiers to the same version, such that VBS version across all tiers is the same.

**Workaround:** Maintain the same VBS versions on all tiers.

### **Fault propagation for Virtual Business Services with shared service groups and different controllers [2407832]**

Fault propagation may not work for certain configurations having shared service groups and distinct controllers.

Workaround: No workaround.

### **Virtual Business Services fail to start if a participating service group has multiple children with the LOCAL FIRM dependency type [2490098]**

Virtual Business Services fail to start if a participating service group has multiple children with the LOCAL FIRM dependency type. This occurs because Veritas Cluster Server (VCS) does not support propagating dependencies.

**Workaround:** Pull the dependent VCS groups into the Virtual Business Services without any dependencies. The Virtual Business Services will recognize the VCS dependencies and treat them as soft Virtual Business Services dependencies.

## Veritas Storage Foundation and High Availability known issues

For known issues of Veritas Storage Foundation and High Availability, refer to the section called “Veritas Storage Foundation known issues” and the section called “Veritas Cluster Server known issues”.

## Veritas Storage Foundation Cluster File System High Availability known issues

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

### **Veritas Storage Foundation Cluster File System High Availability known issues**

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

#### **CFS commands might hang when run by non-root (2403263)**

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

#### **Workaround**

##### **To resolve this issue**

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

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

### Miscalculated file set usage (2123429)

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

```
# fscckptadm getquotalimit /mnt1
Filesystem  hardlimit  softlimit  usage  action_flag
/mnt1       10000     10000     18446744073709551614
```

This could cause writes to Storage Checkpoints to fail. It could also trigger the removal of removable Storage Checkpoints.

### Workaround

If this occurs, disabling and re-enabling file set quotas causes VxFS to recalculate the number of blocks used by Storage Checkpoints:

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

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

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

---

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

---

### Workaround

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

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

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

### Workaround

Create a resource dependency between the various CFSmount resources.

### 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.

### NFS issues with VxFS Storage Checkpoint (1974020)

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

```
Stale NFS file handle
```

This is a result of major numbers of VxFS Storage Checkpoints not necessarily being the same on all SFHA cluster nodes.

**Workaround:** There is no workaround at this time.

### Panic due to null pointer de-reference in vx\_bmap\_lookup() (2582232)

A null pointer dereference in the `vx_bmap_lookup()` call can cause a panic.

**Workaround:** Resize the file system with the `fsadm` command from the primary node of the cluster.

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

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

**Workaround:** There is no workaround for this issue.

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

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

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

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

**Workaround:** There is no workaround for this issue.

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

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

#### **Work-around:**

##### **To work around this issue**

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

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

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

#### **To run disk discovery**

- ◆ Run the following command:

```
# vxdisk scandisks
```

### **The vxdsconvert utility is supported only on the master node (2616422)**

The `vxdsconvert` utility should be run only from the master node, not from the slave nodes of the cluster.

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

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

- The shared disk groups on the disconnected storage are marked as `dgdisabled` on the slave node only.

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

**Work-arounds:**

Use one of the following work-arounds:

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

**The "vxdg listclone" command output may not list all the disks with "clone\_disk" or "udid\_mismatch" flag set (235456)**

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

**Workaround:**

Administrator has to run "vxdisk scandisks" or "vxdisk -o alldgs list" followed by "vxdg listclone" to get all the disks containing "clone\_disk" or "udid\_mismatch" flag on respective host.

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

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

**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 -F 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 -F vxfs /dev/vx/dsk/dg/data_volume
```

**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.

## Veritas Storage Foundation for Oracle RAC known issues

This section describes the known issues in this release of Veritas Storage Foundation for Oracle RAC.

### Oracle RAC issues

This section lists the known issues in Oracle RAC.

## Oracle Grid Infrastructure installation may fail with internal driver error

The Oracle Grid Infrastructure installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

### Workaround:

Perform one of the following steps depending on the type of installer you use for the installation:

#### ■ Script-based installer

Export the `OUI_ARGS` environment variable, before you run the SFHA installation program:

```
export OUI_ARGS=-ignoreInternalDriverError
```

For more information, see the Oracle Metalink document: 970166.1

#### ■ Web-based installer

When you run the Web-based installer, in the **Enter the arguments to be passed to the Oracle installer** text box, enter the value `-ignoreInternalDriverError`.

For more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

## 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

## Enabling ODM in Oracle RAC 11 Release 2 installations causes errors (1913013)

Enabling ODM in Oracle RAC 11 Release 2 installations causes the following error:

```
'ODM ERROR V-41-4-1-253-12 Not enough space'
Oracle instance may also crash with same error.
```

The error is observed if the DISM (Dynamic Intimate Shared memory) feature is enabled. In Solaris, the Oracle database uses DISM if it is available on the system, and if the value of the `sga_max_size` initialization parameter is larger than the size required for all SGA components combined.

**Workaround:** Make sure that the file `ORACLE_HOME/bin/oradism` is owned by the root user with "execute" and "setuid" permissions. If the problem persists after correcting the permissions, uncomment the `sga_max_size` and `memory_target` `init.ora` parameters.

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

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

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

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

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

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

### Oracle Cluster Verification utility displays a warning message

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

For example:

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

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

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

```
=====
```

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

### **Oracle Database Configuration Assistant displays an error**

The Database Configuration Assistant utility displays the following error:

```
SGA size cannot be greater than maximum shared memory  
segment size (0).
```

**Workaround:** Ignore this message and manually configure the database memory parameters for Oracle. In the "Memory" tab of the Oracle Database Creation Assistant (DBCA), select a Custom and Manual shared memory management configuration and enter the appropriate values.

## **SFHA issues**

This section lists the known issues in SFHA for this release.

### **Verification of Oracle binaries incorrectly reports as failed during Oracle Grid Infrastructure installation**

The verification of Oracle binaries may incorrectly report as failed during the Oracle Grid Infrastructure installation using the SFHA installer. The message is erroneously reported due to a break in passwordless SSH communication. The SSH communication fails because execution of the `root.sh` script changes the owner of the operating system root directory to the grid user directory.

### **SFHA installer does not support use of fully qualified domain names (2585899)**

The SFHA installer does not support the use of fully qualified domain names (FQDN). Specifying the fully qualified domain name of a system results in the following error:

```
The node galaxy doesn't seem to be part of the cluster,  
or CVM is not running on the node galaxy.
```

**Workaround:** Use only the host name of the system when you specify the system name.

### **PrivNIC and MultiPrivNIC agents not supported with Oracle RAC 11.2.0.2**

The PrivNIC and MultiPrivNIC agents are not supported with Oracle RAC 11.2.0.2.

For more information, see the following Technote:

<http://www.symantec.com/business/support/index?page=content&id=TECH145261>

### **Issue with format of the last 8-bit number in private IP addresses (1164506)**

The PrivNIC/MultiPrivNIC resources fault if the private IP addresses have a leading 0 in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1 or X.0X.X.1 or 0X.X.X.1, where X is an octet of the IP address.

When you configure private IP addresses for Oracle Clusterware, ensure that the IP addresses have a format as displayed in the following two-node example:

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

Confirm the correct format by viewing the PrivNIC or MultiPrivNIC resource in the `/etc/VRTSvcs/conf/config/main.cf` file.

### **Warning message displayed on taking cssd resource offline if LANG attribute is set to "eucJP" (2123122)**

When you take the `cssd` resource offline using the `hares -offline cssd` command and the `LANG` attribute is set to "eucJP", the following message may be observed in the `hamsg engine_A` command output:

```
VCS INFO V-16-2-13716 Could not find message V-16-2-13716
```

You may ignore the message.

### **Error displayed on removal of VRTSjadba language package (2569224)**

Removal of the VRTSjadba language package displays the following error on the screen:

```
Executing postremove script.  
Generating BMC map file...  
bmcmap ERROR V-33-1000-10001 Unable to create BMC map
```

You may ignore the error.

### **Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool (2392741)**

If the cardinality of a policy-managed Oracle RAC database is set to a number lesser than the number of nodes in the server pool, and if the Oracle agent tries to bring the database online on all the nodes in the server pool, the operation fails on some of the nodes in the server pool. The resource on respective nodes move to the faulted state.

### **Removal of SAN cable from any node in a global cluster setup takes application service groups offline on all nodes (2580393)**

In a replicated global cluster setup, the removal of SAN cable from any node in the cluster causes the CFS mount points to fault. As a result, dependent application groups are taken offline and replication to the secondary site is adversely affected.

## Veritas Storage Foundation for Sybase ASE CE known issues

This section describes the known issues in this release of Veritas Storage Foundation for Sybase ASE CE.

### **Stale .vxfsdargs file lets hashadow restart vxfsd in Sybase mode (2554886)**

When I/O fencing is configured in customized mode, vxfsd, the user mode daemon of I/O fencing, creates the `/opt/VRTSvcs/lock/.vxfsdargs` file. VCS uses this file to restart the vxfsd daemon when it gets killed. However, VCS does not use this file when I/O fencing is configured in Sybase mode. This file is not removed from the system when I/O fencing is unconfigured.

If user configures I/O fencing in Sybase mode and an old `/opt/VRTSvcs/lock/.vxfsdargs` file is present in the system from an earlier configuration of I/O fencing in customized mode, then VCS attempts to restart the vxfsd daemon every time it is killed. This interferes with the functioning of I/O fencing in the Sybase mode.

**Workaround:** Before you configure I/O fencing in Sybase mode, delete the `/opt/VRTSvcs/lock/.vxfsdargs` file if it is present in the system.

### **Sybase Agent Monitor times out (1592996)**

**Problem:** The Sybase Agent Monitor has issues with timing out, in cases where qrmutil reports delay.

The Sybase Agent monitor times out, if qrmutil fails to report the status to the agent within the defined MonitorTimeout for the agent.

**Solution:** If any of the following configuration parameters for Sybase Database is increased, it will require a change in its MonitorTimeout value:

- quorum heartbeat interval (in seconds)
- Number of retries

If the above two parameters are changed, Symantec recommends that the MonitorTimeout be set to a greater value than the following:  $((\text{number of retries} + 1) * (\text{quorum heartbeat interval})) + 5$ .

## Installer warning (151550)

**Problem:** During configuration of Sybase instance under VCS control, if the quorum device is on CFS and is not mounted, the following warning message appears on the installer screen:

```
Error: CPI WARNING V-9-40-4563 The quorum file /qrmnt/qfile
cannot be accessed. This may be due to a file system not being mounted.
```

The above warning may be safely ignored.

## Unexpected node reboot while probing a Sybase resource in transition (1593605)

**Problem:** A node may reboot unexpectedly if the Sybase resource is probed while the resource is still in transition from an online to offline state.

Normally the monitor entry point for Sybase agent completes with 5-10 seconds. The monitor script for the Sybase agent uses the qrmutil binary provided by Sybase. During a monitor, if this utility takes longer time to respond, the monitor entry point will also execute for longer duration before returning status.

**Resolution:** During the transition time interval between online and offline, do not issue a probe for the Sybase resource, otherwise the node may reboot.

## Unexpected node reboot when invalid attribute is given (2567507)

**Problem:** A node may reboot unexpectedly if the Home, Version, or Server attributes are modified to invalid values while the Sybase resources are online in VCS.

**Resolution:** Avoid setting invalid values for the Home, Version, or Server attributes while the Sybase resources are online in VCS, to avoid panic of the node.

## 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.

## Resources not brought online after had process is killed and restarted

If the `had` process is killed either manually or by some other means, the `hashadow` process restarts it. During `had` restart time, the resources remain in the same state

unless their states are changed manually, but after restarting the `had`, process VCS may not bring some resources online.

#### To work around this issue

- ◆ Manually restart the resources that were not started during `had` restart.

### **AutoFailOver = 0 attribute absent in the sample files at /etc/VRTSagents/ha/conf/Sybase (2615341)**

**Problem:** `AutoFailOver = 0` attribute is not present in the sample files at `/etc/VRTSagents/ha/conf/Sybase`.

**Resolution:** If you copy the `main.cf` file from the `/etc/VRTSagents/ha/conf/Sybase` location, add the `AutoFailOver = 0` attribute to the `binmnt` and `sybasece` service groups.

## Symantec VirtualStore known issues

This section describes the known issues in this release of Symantec VirtualStore.

### **The `svsiscsiadm create lun` command fails if you create a LUN greater than the available space on the file system (2567517)**

The `svsiscsiadm create lun` command fails if you create a LUN of a size greater than the total amount of space available on the file system. The underlying `iscsitadm` command fails with the following error message:

```
iscsitadm: Error Requested size is too large for system
```

The report of this error is logged in the `/var/VRTSvcs/log/engine_A.log` file.

If you then try to create a LUN on the same target, the LUN creation call fails again with the following error message:

```
iscsitadm: Error Failed to create a symbolic link to the backing store
```

The report of this error is logged in the `/var/VRTSvcs/log/engine_A.log` file.

This makes the target unusable.

#### **Workaround**

### To resolve this issue

- 1 Note the TargetID and LunID on which the `svsiscsiadm create lun` command failed. To find the failed LunID, note the last LunID for the target on which `svsiscsiadm create lun` command failed with the use of the `svsiscsiadm list` command. To calculate the failed LunID, add 1 to last LunID seen by the `svsiscsiadm list` command.

- 2 Go to the configuration directory for the TargetID:

```
# cd /etc/iscsi/TargetID .
```

- 3 Delete the symlink pointing to path of LUN backing file which failed to get added. The below LunID is the failed LunID, which is a result of calculation in point 1:

```
# rm -f /etc/iscsi/TargetID/lun.($lunid + 1)
```

After removal of the symlink you should be able to add LUNs on the unusable target.

## VirtualStore machine clones created while the VirtualStore cluster reboots will probably not start (2164664)

In some cases when you clone while rebooting the SVS nodes, you may receive several of the following error messages:

```
clone vms could not start X server
```

### Workaround

Delete all the clones that got created while the node crashed and redo the cloning operation.

## Cloning may not work (2348628)

If you cannot clone and you are using the VMware vAPP and OVF templates, then you must disable the vApp.

### Workaround

To disable the vAPP

- 1 In VI Client, right-click on the virtual machine > **Edit Settings** > **Options** > **vApp Options**.
- 2 Click **Disable**.

## **Need intelligent NDMP/NBU backups for virtual machines (2378396)**

When using NDMP or the NBU client to backup a virtual machine, the space consumed by the backup is equivalent to the size of the disks in the virtual machine, even though not all of the disk space in the virtual machine is used.

If a VMDK (Virtual Machine Disk) file is 10GB in size, but only consumes 1GB of disk space, an backup done by NDMP or the NBU client generates 10GB of backup data, even though the original VMDK file contains 9GB of unassigned disk space.

### **Workaround**

Use VMware-specific backup applications (such as NetBackup for VMware) to create space-efficient backups.

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

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

For example, if you do the following:

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

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

### **Workaround**

To resolve this issue:

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

## **Virtual machines created by the Symantec Quick Clone Virtual Machine Wizard might not boot correctly if during the process the FileStore cluster node, the ESX Server, or the vCenter Server reboots (2164664, 2374229)**

In some cases when you clone using the wizard, and one of the following servers crashes or reboots while the clone process is in progress, the clones might not get created correctly:

- FileStore nodes

- ESX host on which the clones are being created
- vCenter Server

Even if the clones appear in the vCenter inventory as created, the clones GuestOS might not be able to boot.

#### **Workaround**

Delete all of the clones that were created when the servers crashed or were rebooted, and redo the wizard operation.

### **Error message does not always display when you select an incorrect cluster to clone (2372713)**

In cases where multiple FileStore clusters are registered with the same Virtual Center, the Symantec Quick Clone Virtual Machine Wizard might not provide a warning that you selected an incorrect cluster to clone a golden image. This could happen if all of the FileStore clusters are exporting the same file system path, such as `/mnt`. Instead of an advanced warning that you selected the wrong cluster, you instead see an error on the final page of the wizard when the wizard attempts to clone the disks (vmdks) of the golden image. The error that displays is similar to the following example:

```
/mnt/goldvm/goldvm.vmdk no such file or directory...
```

#### **Workaround**

There is no workaround for this issue.

### **Cloning issue in a Japanese environment (2623471)**

You might be unable to clone with Guest OS Customization or VMware View integration. While using the FileSnap wizard, options may be missing or in an error state.

#### **Workaround**

The workaround involves temporarily running the vCenter Server in the English locale.

**To resolve this issue**

- 1 From the vCenter Server, stop the following services using the Task Manager or services.msc:

```
VMware VCMSD
VMware VirtualCenter Server
VMware VirtualCenter Management Webservices
VMware vCenter Update Manager Services
```

- 2 Rename the following language directories ja to ja-x:

```
C:\Program Files\VMware\Infrastructure\VirtualCenter Server\ja
C:\Program Files\VMware\Infrastructure\VirtualCenter Server\locale\ja
C:\Program Files\VMware\Infrastructure\VirtualCenter Server\imgres\ja
```

- 3 Restart the services from step 1.
- 4 Take FileSnap clones with customization and/or View integration using the FileSnap wizard.
- 5 To return the vCenter Server to the Japanese locale, reverse steps 1-3.

## Software limitations

There are no software limitations in this release.

## List of patches

This section lists the patches and packages for 6.0 RP1.

**Table 1-10** Patches and packages for Solaris SPARC

Patch ID	Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
146917-01	VRTSvcs	SFCFSHA,SFHA, SFSYBASECE, SVS,VCS	133 MB		X
146918-01	VRTSvcsag	SFCFSHA,SFHA, SFSYBASECE, SVS,VCS	4.1 MB		X

**Table 1-10** Patches and packages for Solaris SPARC (*continued*)

Patch ID	Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
147852-01	VRTSvxfs	FS,SF,SFCFSHA, SFHA,SFSYBASECE, SVS	37 MB		X
147853-02	VRTSvxvm	DMP,SF,SFCFSHA, SFHA,SFSYBASECE, SVS,VM	194 MB	X	X
147883-01	VRTSamf	SFCFSHA,SFHA, SFSYBASECE, SVS,VCS	2.6 MB		X
147887-01	VRTSvbs	SFCFSHA,SFHA, SFSYBASECE, SVS,VCS	53 MB	X	X
147893-01	VRTSfcpi60	DMP,FS,SF, SFCFSHA,SFHA, SFSYBASECE, SVS,VCS,VM	4.5 MB		X
148452-01	VRTSfsadv	FS,SF,SFCFSHA, SFHA,SFSYBASECE, SVS	15 MB		X
148453-01	VRTSsvs	SVS	151 MB		X
148457-01	VRTScavf	SFCFSHA,SFSYBASECE,SVS	873 KB		X
148462-01	VRTSob	FS,SF,SFCFSHA, SFHA,SFSYBASECE, SVS,VM	47 MB	X	X

**Table 1-11** Patches and packages for Solaris x64

Patch ID	Package Name	Products Affected	Patch Size	Solaris 10
147849-01	VRTSvcs	SFCFSHA,SFHA, SVS,VCS	228 MB	X
147850-01	VRTSvcsag	SFCFSHA,SFHA, SVS,VCS	6.2 MB	X
147854-02	VRTSvxvm	DMP,SF,SFCFSHA, SFHA,SVS,VM	201 MB	X
147855-01	VRTSvxfs	FS,SF,SFCFSHA, SFHA,SVS	31 MB	X
147886-01	VRTSamf	SFCFSHA,SFHA, SVS,VCS	2.3 MB	X
147888-01	VRTSvbs	SFCFSHA,SFHA, SVS,VCS	85 MB	X
148450-01	VRTSsfcp60	DMP,FS,SF,SFCFSHA, SFHA,SVS,VCS,VM	4.3 MB	X
148454-01	VRTSsvs	SVS	142 MB	X
148458-01	VRTScavf	SFCFSHA, SVS	880 KB	X
148463-01	VRTSob	FS,SF,SFCFSHA, SFHA,SVS,VM	44 MB	X

---

**Note:** You can also view the list using the `installrp` command: `./installrp -listpatches`

---

## Documentation errata

The following sections cover additions or corrections for the product documentation. These additions or corrections may be included in later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).

## Veritas Storage Foundation Cluster File System High Availability manual page

The following errata applies to the Veritas Storage Foundation Cluster File System High Availability 6.0 manual page:

### **cfsshare manual page**

The following argument is missing from the cfsshare manual page:

*network\_hosts* Hosts on the network that receive pings to determine the state of the NIC. The specified hosts must be pingable.

# Installing the products for the first time

This chapter includes the following topics:

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

## Installing the Veritas software using the script-based installer

This section describes how to install a 6.0 RP1 Veritas Storage Foundation and High Availability Solutions product for the first time on a host. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 6.0 *Installation Guide* and *Release Notes* for your product for more information.

### To install the Veritas software for the first time

- 1 Download Storage Foundation and High Availability Solutions 6.0 from <http://fileConnect.symantec.com>.
- 2 Extract the tar ball into a directory called `/tmp/sfha6.0`.
- 3 Check <http://sort.symantec.com/patches> to see if there are any patches available for the 6.0 Installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 4 Change to the `/tmp/sfha6.0` directory:

```
# cd /tmp/sfha6.0
```

- 5 Run the installer to install SFHA 6.0.

See the *Installation Guide* for instructions on installing the 6.0 version of this product.

```
#./installer -require complete_path_to_6.0_installer_patch
```

---

**Note:** If the P-patch is not available for 6.0 installer, use the `installer` script without `-require` option.

---

- 6 Download SFHA 6.0 RP1 from <http://sort.symantec.com/patches>.
- 7 Extract it to a directory called `/tmp/sfha6.0RP1`.
- 8 Check <http://sort.symantec.com/patches> to see if there are patches available for the 6.0 RP1 installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 9 Change to the `/tmp/sfha6.0RP1` directory:

```
# cd /tmp/sfha6.0RP1
```

- 10 Invoke the `installrp` script to install 6.0 RP1:

```
# ./installrp -require complete_path_to_6.0RP1_installer_patch
```

---

**Note:** Note: If the P-patch is not available for 6.0 RP1 installer, use the `installrp` script without `-require` option.

---

- 11 If you did not configure the product after the 6.0 installation, the installer prompts you to configure the product during RP installation. If you do not want to configure the product now, answer `n` when prompted. To configure the product in the future, run the product installation script from the 6.0 installation media or from `/opt/VRTS/install` directory with the `-configure` option.

## Installing Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 6.0 RP1 using the Web-based installer. For detailed instructions on how to install 6.0

using the Web-based installer, follow the procedures in the 6.0 Installation Guide and Release Notes for your products.

See “[Upgrading to 6.0 RP1](#)” on page 116.

## Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

### To start the Web-based installer

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

```
# ./webinstaller start
```

The webinstaller script displays a URL.

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

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

## Obtaining a security exception on Mozilla Firefox

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

### To obtain a security exception

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

## Installing 6.0 RP1 with the Veritas Web-based installer

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

### To install SFHA

- 1 The 6.0 version of the Veritas product must be installed before upgrading to 6.0 RP1.

See “Prerequisites for upgrading to 6.0 RP1” on page 115.

- 2 On the **Select a task and a product** page, select **Install 6.0 RP1** 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 **Next**.
- 4 After the validation completes successfully, click **Next** to install 6.0 RP1 patches on the selected system.
- 5 After the installation completes, you must choose your licensing method. On the license page, select one of the following tabs:

- Keyless licensing

---

**Note:** The keyless license option enables you to install without entering a key. However, in order to ensure compliance you must manage the systems with a management server.

For more information, go to the following website:

<http://go.symantec.com/sfhakeyless>

---

Complete the following information:

Choose whether you want to install Standard or Enterprise mode.

Choose whether you want to enable Veritas Volume Replicator.

For Storage Foundation High Availability, choose whether you want to enable Global Cluster option.

Choose whether you want to enable Global Cluster option.

Click Register.

- Enter license key

If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

- 6 For Storage Foundation, click Next to complete the configuration and start the product processes.

For Storage Foundation High Availability, the installer prompts you to configure the cluster.

Note that you are prompted to configure only if the product is not yet configured.

If you select n, you can exit the installer. You must configure the product before you can use SFHA.

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 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 SFHA.

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.

- 8 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**.



# Installing the products using JumpStart

This chapter includes the following topics:

- [Installing the products using JumpStart](#)

## Installing the products using JumpStart

These JumpStart instructions assume a working knowledge of JumpStart. See the JumpStart documentation that came with your operating system for details on using JumpStart. Only fresh installations of SFHA are supported using JumpStart. Upgrading is not supported. The following procedure assumes a stand-alone configuration.

For the language pack, you can use JumpStart to install packages. You add the language packages in the script, and put those files in the JumpStart server directory.

## Generating the finish scripts

Perform these steps to generate the finish script to install SFHA.

### To generate the finish script

- 1 Mount the 6.0 RP1 media and run the `installrp` program to generate the scripts.

```
# installrp -jumpstart directory_to_generate_scripts
```

where the *directory\_to\_generate\_scripts* is the location where you want to put the scripts.

For example:

```
# ./installrp -jumpstart /js_scripts
```

- 2 When you are prompted to encapsulate the root disk automatically, choose **yes** to do so. If you do not want to encapsulate it automatically, choose **no** and go to step 6.
- 3 Specify a disk group name for the root disk.

```
Specify the disk group name of the root disk to be encapsulated:  
rootdg
```

- 4 Specify private region length.

```
Specify the private region length of the root disk to be  
encapsulated: (65536)
```

- 5 Specify the disk's media name of the root disk to encapsulate.

```
Specify the disk media name of the root disk to be encapsulated:  
(rootdg_01)
```

**6** JumpStart finish scripts of Veritas products, and encapsulation scripts are generated in the directory you specified in step 1. Output resembles:

```
The finish scripts for DMP is generated at
/js_scripts/jumpstart_dmp.fin
The finish scripts for FS is generated at
/js_scripts/jumpstart_fs.fin
The finish scripts for SF is generated at
/js_scripts/jumpstart_sf.fin
The finish scripts for SFCFSHA is generated at
/js_scripts/jumpstart_sfcfsha.fin
The finish scripts for SFHA is generated at
/js_scripts/jumpstart_sfha.fin
The finish scripts for SF Oracle RAC is generated at
/js_scripts/jumpstart_sfrac.fin
The finish scripts for SVS is generated at
/js_scripts/jumpstart_svs.fin
The finish scripts for VCS is generated at
/js_scripts/jumpstart_vcs.fin
The finish scripts for VM is generated at
/js_scripts/jumpstart_vm.fin
The encapsulation boot disk script for VM is generated at
/js_scripts/encap_bootdisk_vm.fin
```

List the `js_scripts` directory.

```
# ls /js_scripts
```

You could select scripts according to the products you want to install and copy them to the `BUILDSRC` NFS shared location. For example, `/export/config` where you mounted the `BUILDSRC`.

For SF:

```
encap_bootdisk_vm.fin jumpstart_sf.fin
```

For SFHA:

```
encap_bootdisk_vm.fin jumpstart_sfha.fin
```

For SFCFSHA:

```
encap_bootdisk_vm.fin jumpstart_sfcfs.fin
```

For SF Oracle RAC:

```
encap_bootdisk_vm.fin jumpstart_sfrac.fin
```

For VCS:

```
encap_bootdisk_vm.fin jumpstart_vcs.fin
```

For DMP:

```
encap_bootdisk_vm.fin jumpstart_dmp.fin
```

For SVS:

```
encap_bootdisk_vm.fin jumpstart_svs.fin
```

For FS:

```
encap_bootdisk_vm.fin jumpstart_fs.fin
```

For VM:

```
encap_bootdisk_vm.fin jumpstart_vm.fin
```

- 7 Copy the install and uninstall scripts of Veritas product that you are installing on to `BUILDSRC` shared location `/export/config` from 6.0 media.

For example:

```
# cp -p /dvd_mount/product_name/installprod /export/config  
# cp -p /dvd_mount/product_name/uninstallprod /export/config
```

Here *prod* is one of `dmp/fs/sf/sfcfsha/sfha/sfrac/svs/vcs/vm`.

- 8 Modify the JumpStart script according to your requirements. You must modify the `BUILDSRC` and `ENCAPSRC` values. Keep the values aligned with the resource location values.

```
BUILDSRC="hostname_or_ip:/path_to_pkgs_patches_scripts"
```

Example: **BUILDSRC=10.209.100.100:/export/config**

```
// If you don't want to encapsulate the root disk automatically  
// comment out the following line.
```

```
ENCAPSRC="hostname_or_ip:/path_to_encap_script"
```

- 9 ■ If you want to install other products packages in the Storage Foundation and High Availability suite, then use the product-specific install command with one of the following options to get a list of packages in the order to be installed:

- `minpkgs`
- `recpkgs`

- `allpkgs`

Use the list of packages that is generated to replace the package list in the finish scripts.

- If you want to install other products patches in the Storage Foundation and High Availability suite, then use the product-specific install command with one of the following options to get a list of patches in the order to be installed:

```
# ./installrp -listpatches
```

See “[The installrp script options](#)” on page 12.

- 10 Once the installation is complete, refer to installation and configuration guide for the respective product from 6.0 to proceed with the configuration.

## Overview of JumpStart installation tasks

Review the summary of tasks before you perform the JumpStart installation.

### Summary of tasks

- 1 Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.

- 2 Read the JumpStart installation instructions.

- 3 Generate the finish scripts.

See “[Generating the finish scripts](#)” on page 107.

- 4 Prepare shared storage installation resources.

See “[Preparing installation resources](#)” on page 112.

- 5 Modify the rules file for JumpStart.

See the JumpStart documentation that came with your operating system for details.

- 6 Run the installer command from the disc or from directory `/opt/VRTS/install` directory to configure the Veritas software.

```
# /opt/VRTS/install/installprod -configure
```

### Summary of tasks

- 1 Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.

- 2 Read the JumpStart installation instructions.

- 3 Generate the finish scripts.  
See “[Generating the finish scripts](#)” on page 107.
- 4 Modify the rules file for JumpStart.  
See the JumpStart documentation that came with your operating system for details.
- 5 Prepare installation resources.  
See “[Preparing installation resources](#)” on page 112.
- 6 On each client node, run the following command to install the Veritas product packages and patches:

For Solaris SPARC:

```
Ok> boot net - install
```

For Solaris x64:

Press **F12** and select the network boot mode.

---

**Note:** The system is restarted after the packages are installed. If you choose to encapsulate the root disk on your systems, the systems start with an encapsulated root disk.

---

- 7 Run the `installprod` command from the disc or from directory `/opt/VRTS/install` directory to configure the Veritas software.

```
# /opt/VRTS/install/installprod -configure
```

where `installprod` is the product's installation command.

## Preparing installation resources

Prepare resources for the JumpStart installation.

## To prepare the resources

- 1 Copy the contents of 6.0 disk and 6.0 RP1 disk both to buildsrc.

```
# cd /cdrom/cdrom0
# cp -r * BUILDSRC
```

---

**Note:** After you copied the patches, you must uncompress them using the gunzip and tar commands.

---

- 2 Generate the response file for the package and patch list that you found when you generated the finish script.

See [“Generating the finish scripts”](#) on page 107.

To view the patches, packages and operating systems for your Veritas product use the `installrp -listpatches` command, type:

```
# ./installrp -listpatches

# cd BUILDSRC/pkg/
# pkgask -r package_name.response -d /BUILDSRC/pkg/packages_name.pkg
```

- 3 Generate the response file for the package list that you found when you generated the finish script.

See [“Generating the finish scripts”](#) on page 107.

In this example, the packages are: VRTSaslapm, VRTSdbac, and VRTSvxvm.

```
# cd BUILDSRC/pkg/
# pkgask -r package_name.response -d /BUILDSRC/pkg/packages_name.pkg
```

- 4 Create the `adminfile` file under `BUILDSRC/pkgs/` directory. The `adminfile` file's contents follow:

```
mail=  
instance=overwrite  
partial=nocheck  
runlevel=quit  
idepend=quit  
rdepend=nocheck  
space=quit  
setuid=nocheck  
conflict=nocheck  
action=nocheck  
basedir=default
```

- 5 Copy the install and uninstall scripts that you created when you generated the finish script to `BUILDSRC` if you want to configure or uninstall from `/opt/VRTS/install`. Otherwise, you need to configure and uninstall from disc. See [“Generating the finish scripts”](#) on page 107.
- 6 If you want to encapsulate the root disk automatically when perform the JumpStart installation, copy the scripts `encap_bootdisk_vm51.fin` created when you generated the finish script to `ENCAPSRC`. See [“Generating the finish scripts”](#) on page 107.

# Upgrading to 6.0 RP1

This chapter includes the following topics:

- [Prerequisites for upgrading to 6.0 RP1](#)
- [Downloading required software to upgrade to 6.0 RP1](#)
- [Supported upgrade paths](#)
- [Upgrading to 6.0 RP1](#)
- [Verifying software versions](#)

## Prerequisites for upgrading to 6.0 RP1

The following list describes prerequisites for upgrading to the 6.0 RP1 release:

- For any product in the Veritas Storage Foundation stack, you must have the 6.0 installed before you can upgrade that product to the 6.0 RP1 release.
- Each system must have sufficient free space to accommodate patches.
- 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 6.0 RP1](#)” on page 115.

## Downloading required software to upgrade to 6.0 RP1

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

### To download required software to upgrade to 6.0 RP1

- 1 Download SFHA 6.0 RP1 from <http://sort.symantec.com/patches>.
- 2 Extract it to a directory, say /tmp/sfha60rp1.

## Supported upgrade paths

This section describes the supported upgrade paths for this release:

- 6.0 to 6.0 RP1

## Upgrading to 6.0 RP1

This section describes how to upgrade from 6.0 to 6.0 RP1 on a cluster or a standalone system.

- [Performing a full upgrade to 6.0 RP1 on a cluster](#)  
Use the procedures to perform a full upgrade to 6.0 RP1 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System High Availability (SFCFSHA), Veritas Storage Foundation for Oracle RAC (SFRAC), or Symantec VirtualStore (SVS) installed and configured.
- [Upgrading to 6.0 RP1 on a standalone system](#)  
Use the procedure to upgrade to 6.0 RP1 on a system that has SF installed.
- [Upgrading Veritas products using Live Upgrade](#)  
Use the procedure to upgrade your Veritas product with a Live Upgrade.
- [Performing a rolling upgrade using the installer](#)  
Use the procedure to upgrade your Veritas product with a rolling upgrade.

See “[Installing the Veritas software using the script-based installer](#)” on page 101.

## Performing a full upgrade to 6.0 RP1 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 (SFCFSHA) and Cluster Volume Manager (CVM), the SFCFSHA and CVM services remain available.

Depending on your cluster’s configuration, select one of the following procedures to upgrade to 6.0 RP1:

- [Performing a full upgrade to 6.0 RP1 on a Veritas Cluster Server](#)
- [Performing a full upgrade to 6.0 RP1 on an SFHA cluster](#)

- [Performing a full upgrade to 6.0 RP1 on an SFCFSA cluster](#)
- [Performing a full upgrade to 6.0 RP1 on an SF Oracle RAC cluster](#)  
See “[Downloading required software to upgrade to 6.0 RP1](#)” on page 115.

## Performing a full upgrade to 6.0 RP1 on a Veritas Cluster Server

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

### To upgrade VCS

- 1 Make sure you have downloaded the latest VCS 6.0 RP1 patches required for the upgrade.
- 2 Log in as superuser.

---

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

---

- 3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 6.0 RP1 rolling patch binaries, change to the directory that contains the `installrp` script. Start the pre-upgrade check:

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

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

```
# ./installrp node1 node2 ... nodeN
```

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

## Performing a full upgrade to 6.0 RP1 on an SFHA cluster

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

### To perform a full upgrade to 6.0 RP1 on an SFHA cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 2 Log in as superuser.

- 3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and uncompressed 6.0 RP1 rolling patch binaries, change to the directory that contains the `installrp` script. Start the pre-upgrade check:

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

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

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

```
# ./installrp node1 node2 ... nodeN
```

## Performing a full upgrade to 6.0 RP1 on an SFCFSHA cluster

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

### To perform a full upgrade to 6.0 RP1 on an SFCFSHA cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` and `/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 any node:

```
# hagrps -freeze groupname -persistent
```

- 6 Make the configuration read-only:

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

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

```
# df -F vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# df -F vxfs
```

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

```
# umount /filesystem
```

- 9 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.

---

- 10 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.

- 11** On each node, stop all VxVM volumes by entering the following command for each disk group:

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

Verify that no volumes remain open:

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

- 12** Stop VCS:

```
# hastop -all
```

- 13** On each node, stop the VCS command server:

```
# ps -ef | grep CmdServer  
# kill -9 pid_of_CmdServer
```

where *pid\_of\_CmdServer* is the process ID of CmdServer.

- 14** On each node, stop ODM, cluster fencing, GAB, and LLT in the following order:

```
# svcadm disable -t vxfen  
# svcadm disable -t vxodm  
# svcadm disable -t gab  
# svcadm disable -t llt
```

- 15** If required, apply the OS kernel patches.

See Oracle's documentation for the procedures.

- 16** On each node, check if the VEA service is running:

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

If the VEA service is running, stop it:

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

- 17** From the directory that contains the extracted and untarred 6.0 RP1 rolling patch binaries, change to the directory that contains the `installrp` script. Start the upgrade.

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

Review the output.

- 18** If necessary, reinstate any missing mount points in the `/etc/vfstab` file on each node.

19 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

20 Enter the following command on any node to unfreeze HA service group operations:

```
# hagrps -unfreeze groupname -persistent
```

21 Make the configuration read-only:

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

22 Bring the CVM service group online on each node:

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

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

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

24 If you stopped any RVGs in step 9, restart each RVG:

```
# vxrvgs -g diskgroup start rvg_name
```

25 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

26 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

## Performing a full upgrade to 6.0 RP1 on an SF Oracle RAC cluster

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

### To upgrade to 6.0 RP1 on an SF Oracle RAC cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 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 all applications that use VxFS or VxVM disk groups, whether local or CFS.

If the applications are under VCS control:

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

If the applications are not under VCS control:

Use native application commands to stop the application.

9 For Oracle RAC 10g and Oracle RAC 11g:

Stop all Oracle RAC resources.

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

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

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

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

- 10 ■ If the Oracle database is managed by VCS, set the `AutoStart` value to 0 to prevent the database service group from starting automatically when VCS starts:

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

- 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
```

- 11 Unmount the VxFS file system, which is not under VCS control.

```
# mount -v |grep vxfs

# fuser -c /mount_point

# umount /mount_point
```

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

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

- 12 Stop VCS.

```
# hastop -all
```

- 13 If required, apply the OS kernel patches.

See *Oracle's* documentation for the procedures.

- 14 From the directory that contains the extracted and untarred 6.0 RP1 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.

- 15 After the entire cluster is upgraded, follow the installer instructions to proceed further.

**16** If necessary, reinstate any missing mount points in the `/etc/vfstab` file on each node.

**17** Relink the SF Oracle RAC libraries with Oracle.

Refer to the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for 6.0 for more information.

**18** From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

**19** Enter the following command on each node to unfreeze HA service group operations:

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

**20** Make the configuration read-only:

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

**21** Enter the following command on each node to take service groups online:

```
# hagrps -online service_group -sys nodename
```

**22** Restart all the volumes by entering the following command for each disk group:

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

**23** Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

**24** 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
```

**25** 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
```

**26** Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

```
# hagrps -online Oracle_group -any
```

- If the Oracle database is not managed by VCS:

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

**27** ■ 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_groupname 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
```

**28** For upgrade scenarios that involve Oracle RAC 9i, start `gsd` as the Oracle user:

```
$ ORACLE_HOME/bin/psdctl start
```

**29** Upgrade Oracle RAC.

---

**Note:** Oracle RAC 11g Release 1 Clusterware is not supported. Make sure that you install Oracle RAC 11g Release 2 Grid Infrastructure in order to use the Oracle RAC 11g Release 1 database. All database versions starting from Oracle 10g Release 2 and later are supported.

---

For instructions, see the chapter *Upgrading Oracle RAC* in *Veritas Storage Foundation™ for Oracle® RAC Installation and Configuration Guide*.

## Upgrading to 6.0 RP1 on a standalone system

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

### To upgrade to 6.0 RP1 on a standalone system

- 1 Make sure you have downloaded the latest software required for the upgrade.

- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 4 If required, apply the OS kernel patches.  
See Oracle's documentation for the procedures.
- 5 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -F 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 `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.

---

- 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/vxsvcctrl status
```

If the VEA service is running, stop it:

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

- 11 Copy the patch archive downloaded from the patch central to temporary location, untar the archive and browse to the directory containing the installrp installer script. Enter the `installrp` script:

```
# ./installrp nodename
```

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

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

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

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

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

- 15 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
```

```
# mount /checkpoint_name
```

- 16 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

## Upgrading Veritas products using Live Upgrade

This section describes how to upgrade 6.0 to 6.0 RP1 using Live Upgrade.

Supported live upgrade paths:

- Upgrading Veritas Products without Solaris OS upgrade:
  - Upgrading Solaris 10 Update x 6.0 to Solaris 10 Update x 6.0 RP1
- Upgrading Veritas Products with Solaris OS upgrade
  - Upgrading Solaris 10 Update x 6.0 to Solaris 10 Update y 6.0 RP1

Prerequisites to upgrade to 6.0 RP1 using Live Upgrade:

- The node should have an alternate boot disk that is identical to the primary boot disk.
- Installation disc for 6.0 and 6.0 RP1 to be installed on the ABE.
- Installation disc for target OS to be installed on ABE.
- The latest list of required patches is available in the Oracle Solaris Live Upgrade Software:  
Patch Requirements (Doc ID 1004881.1) document in My Oracle Support (<https://support.oracle.com/>).
- If OS upgrade is involved, then remove the currently installed SUNWluu, SUNWlur and SUNWlucfg packages and install SUNWluu, SUNWlur, SUNWlucfg packages from target OS. Also replace SUNWluzone if zones are involved.
- The `vxlustart` script takes around 2-3 hours to complete uninterrupted. Symantec recommends to have a network connection that does not time out in the interim.

### Upgrading Veritas products using Live Upgrade from 6.0 to 6.0 RP1 without OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, Sybase ASE CE, or SF for Oracle RAC from 6.0 to 6.0 RP1 using Live Upgrade where OS upgrade is not involved.

### To upgrade your Veritas product using Live Upgrade

- 1 Ensure that 6.0 is installed and configured on PBE.

See your Veritas product 6.0 Installation Guide for more information.

- 2 Run the `vxlustart -V` command to ensure there are no problems before beginning the Live Upgrade process.

If the `vxlustart -V` command reports success, proceed with running the `vxlustart` command.

If the `vxlustart -V` command reports errors, correct the problem, and run the `vxlustart -V` command again.

---

**Note:** This `vxlustart -V` command does not catch failures that are reported by Solaris Live Upgrade commands.

---

- 3 Run the `vxlustart` command to start the Live Upgrade for your Veritas product:

```
# ./vxlustart -v -u target_os_version -U -d disk_name
```

- 4 Run the `installrp` command to upgrade your Veritas product:

```
# ./installrp -rootpath /altroot_path
```

- 5 Run the `vxlufinish` command to complete the Live Upgrade:

- If the primary root disk is not encapsulated, run the following command:

```
# ./vxlufinish -u target_os_version
```

- If the primary root disk is encapsulated by VxVM, run the following command:

```
# ./vxlufinish -u target_os_version -g diskgroup
```

- 6 If the Oracle database is managed by VCS, modify the VCS configuration file on the alternate root disk (`/altroot.5.10/etc/VRTSvcs/conf/config/main.cf`) to set the `AutoStart` value to 0. This prevents the database service group from starting automatically when VCS starts:

For SFRAC:

```
group oradb_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStart = 0
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)
.
.
```

If the database is not managed by VCS, change the management policy for the database to manual on the primary boot disk:

```
$ svrctl modify database -d db-name -y manual
```

- 7 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

---

**Note:** DO NOT use the `reboot`, `halt`, or `uadmin` commands to reboot the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate boot environment.

---

```
# shutdown -g0 -y -i6
```

- 8 In case of SFRAC, refer to the “Performing post-upgrade Tasks” section to relink Oracle RAC libraries with SF Oracle RAC from 6.0 Installation and Configuration guide.
- 9 For SFRAC, relink the SFHA libraries with the Oracle RAC libraries:
- 10 Start the database group on all nodes:

For SFRAC:

```
# hagrps -online oradb_grpname -any
```

- 11 If the Oracle database is managed by VCS, modify the VCS configuration file (`/etc/VRTSvcs/conf/config/main.cf`) to set the `AutoStart` value to 1.

For SFRAC:

```
group oradb_grp (  
    SystemList = { galaxy = 0, nebula = 1 }  
    AutoStart = 1  
    AutoFailOver = 0  
    Parallel = 1  
    AutoStartList = { galaxy, nebula }  
)  
  
.  
.
```

If the database is not managed by VCS, change the management policy for the database to automatic on the primary boot disk:

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

- 12 If you are on an unsupported version of Oracle RAC, upgrade Oracle RAC. For instructions, see the chapter *Upgrading Oracle RAC* in this document.
- 13 Verify that the alternate boot environment is active.

```
# lustatus
```

- 14 In a cluster environment, make sure that all the GAB ports are up. Note different ports appear for different products.

```
# gabconfig -a
```

## Upgrading Veritas products using Live Upgrade from 6.0 to 6.0 RP1 with OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, Sybase ASE CE, or SF for Oracle RAC from 6.0 to 6.0 RP1 using Live Upgrade where OS upgrade is involved..

**To upgrade your Veritas product using Live Upgrade**

- 1 Ensure that 6.0 is installed and configured on PBE.  
See your Veritas product 6.0 Installation Guide for more information.
- 2 Run the `vxlustart -v` command to ensure there are no problems before beginning the Live Upgrade process.

If the `vxlustart -v` command reports success, proceed with running the `vxlustart` command.

If the `vxlustart -v` command reports errors, correct the problem, and run the `vxlustart -v` command again.

---

**Note:** This `vxlustart -v` command does not catch failures that are reported by Solaris Live Upgrade commands.

---

- 3 Run the `vxlustart` command to start the Live Upgrade for your Veritas product:

```
# ./vxlustart -v -u target_os_version -s osimage_path -d disk_name
```

- 4 Run the `installrp` command to upgrade your Veritas product:

```
# ./installrp -rootpath /altroot_path
```

- 5 Run the `vxlufinish` command to complete the Live Upgrade:
  - If the primary root disk is not encapsulated, run the following command:

```
# ./vxlufinish -u target_os_version
```

- If the primary root disk is encapsulated by VxVM, run the following command:

```
# ./vxlufinish -u target_os_version -g diskgroup
```

- 6 If the Oracle database is managed by VCS, modify the VCS configuration file on the alternate root disk (`/altroot.5.10/etc/VRTSvcs/conf/config/main.cf`) to set the `AutoStart` value to 0. This prevents the database service group from starting automatically when VCS starts:

For SFRAC:

```
group oradb_grp (  
    SystemList = { galaxy = 0, nebula = 1 }  
    AutoStart = 0  
    AutoFailOver = 0  
    Parallel = 1  
    AutoStartList = { galaxy, nebula }  
)  
  
.  
.
```

If the database is not managed by VCS, change the management policy for the database to manual on the primary boot disk:

```
$ svct1 modify database -d db-name -y manual
```

- 7 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

---

**Note:** DO NOT use the `reboot`, `halt`, or `uadmin` commands to reboot the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate boot environment.

---

```
# shutdown -g0 -y -i6
```

- 8 In case of SFRAC, refer to the “Performing post-upgrade Tasks” section to relink Oracle RAC libraries with SF Oracle RAC from 6.0 Installation and Configuration guide.
- 9 For SFRAC, relink the SFHA libraries with the Oracle RAC libraries:
- 10 Start the database group on all nodes:

For SFRAC:

```
# hagrp -online oradb_grpname -any
```

- 11** If the Oracle database is managed by VCS, modify the VCS configuration file (/etc/VRTSvcs/conf/config/main.cf) to set the AutoStart value to 1.

For SFRAC:

```
group oradb_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStart = 1
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)
.
.
```

If the database is not managed by VCS, change the management policy for the database to automatic on the primary boot disk:

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

- 12** If you are on an unsupported version of Oracle RAC, upgrade Oracle RAC. For instructions, see the chapter *Upgrading Oracle RAC* in this document.
- 13** Verify that the alternate boot environment is active.

```
# lustatus
```

- 14** In a cluster environment, make sure that all the GAB ports are up. Note that different ports appear for different products.

```
# gabconfig -a
```

## 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 on kernel packages for VCS, SFHA and SFCFSHA: phase 1](#)
- [Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA : phase 2](#)

- [Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1](#)
- [Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2](#)

## About rolling upgrades

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

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

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System and High Availability

You can perform a rolling upgrade from 6.0 to 6.0 RP1.

## 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.
- Make sure you have downloaded the latest software required for the upgrade.

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

## Performing a rolling upgrade using the installer

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

## Performing a rolling upgrade on kernel packages for VCS, SFHA and 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 as subcluster1 and nodeB as subcluster2.

### To perform the rolling upgrade on kernel packages: phase 1

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

```
# ./installrp -rollingupgrade_phase1 nodeA
```

- 2 Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 3 The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
- 4 The installer loads new kernel modules and starts all the relevant processes and brings all the service groups online.
- 5 If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step 6.
- 6 After rolling upgrade phase 1 is completed on nodeA, the following message displays:

```
It is recommended to perform rolling upgrade phase 1 on the systems nodeB in the next step.
```

```
Would you like to perform rolling upgrade phase 1 on the systems?  
[y,n,q] (y)
```

If you choose `y`, it continues to run rolling upgrade phase 1 by itself on nodeB.

If you choose `n` or `q`, you need to complete step 1 to step 4 on nodeB.

- 7 After rolling upgrade phase 1 of the cluster, the following message displays:

```
Would you like to perform rolling upgrade phase 2 on the cluster?  
[y,n,q] (y)
```

- If you choose `y`, it continues to run rolling upgrade phase 2 of the cluster by itself. You don't need to run phase 2: See [“Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA : phase 2”](#) on page 137.

After phase 2 upgrade, verify the cluster's status:

```
# hastatus -sum
```

If you want to upgrade CP server systems that use VCS or SFHA to 6.0, make sure that you upgraded all application clusters to version 6.0. 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.

- If you choose `n` or `q`, you need to use the following steps to finish rolling upgrade phase 2 of the cluster: See [“Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA : phase 2”](#) on page 137.

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

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

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

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

```
# ./installrp -rollingupgrade_phase2 nodeA nodeB nodeC nodeD
```

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

```
# hastatus -sum
```

- 5 If you want to upgrade CP server systems that use VCS or SFHA to 6.0, make sure that you upgraded all application clusters to version 6.0. Then, upgrade VCS or SFHA on the CP server systems.

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

### Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1

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

**To perform the rolling upgrade on kernel: phase 1**

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`.

For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \  
    /etc/VRTSvcs/conf/config/main.cf.save  
# cp /etc/VRTSvcs/conf/config/types.cf \  
    /etc/VRTSvcs/conf/config/types.cf.save  
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \  
    /etc/VRTSvcs/conf/config/OracleTypes.cf.save  
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \  
    /var/tmp/PrivNIC.cf.save  
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \  
    /var/tmp/MultiPrivNIC.cf.save
```

- 3 ■ If the Oracle database is managed by VCS, set the `AutoStart` value to 0 to prevent the database service group from starting automatically when VCS starts. Failing to perform this step results in the database attempting to come online after the upgrade; the attempt fails due to the presence of old libraries on the system.

```
# haconf -makerw  
# 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
```

- 4 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

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

If the applications are not under VCS control:

Use native application commands to stop the application.

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

Stop the Oracle RAC resources on each node.

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

```
# hagrps -offline oracle_group -sys nodeA
# hagrps -offline oracle_group -sys nodeB
```

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

For Oracle RAC 11.2.0.2:

```
$ 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
```

- 6 Switch over all failover service groups to the other nodes in the cluster:

```
# hagrps -switch grp_name -to node_name
```

- 7 Take all the VCS service groups offline:

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

- 8 Unmount all the VxFS file system which is not under VCS control.

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

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

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

---

**Note:** Installer will automatically stop all the applications, database instances, filesystems and volumes which are under VCS control on nodes, while using the `rollingupgrade_phase1` option.

---

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

```
# ./installrp -rollingupgrade_phase1 nodeA nodeB
```

- 10 Note that if the boot-disk is encapsulated, you do not need to perform an unencapsulation for upgrades.
- 11 The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
- 12 Relink the SF Oracle RAC libraries with Oracle.

Refer to the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide for 6.0* for more information.

- 13 If the boot disk is encapsulated, the installer strongly recommends a reboot of the nodes. Reboot the nodes as prompted by the installer.

---

**Note:** Before you reboot the nodes, ensure that the boot device is set to the disk containing the upgraded version of the product.

```
# eeprom
```

---

- 14 After rolling upgrade phase 1 is completed on nodeA and nodeB, the following message displays:

It is recommended to perform rolling upgrade phase 1 on the systems nodeC and nodeD in the next step.

Would you like to perform rolling upgrade phase 1 on the systems?  
 [y,n,q] (y)

- If you choose `y`, first complete step 4 to step 8 on the remaining subcluster. Then it continues to run rolling upgrade phase 1 on nodeC and nodeD by itself.
- If you choose `n` or `q`, go to step 15.

**15** After rolling upgrade phase 1 of the cluster, the following message displays:

Would you like to perform rolling upgrade phase 2 on the cluster?  
 [y,n,q] (y)

- If you choose `y`, it continues to run rolling upgrade phase 2 of the cluster by itself. You don't need to run phase 2: See [“Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2”](#) on page 142. After rolling upgrade phase 2, complete step 14 to step 20 (except step 19) and verify the cluster's status:

```
# hastatus -sum
```

If you want to upgrade CP server systems that use VCS or SFHA to 6.0, make sure that you upgraded all application clusters to version 6.0. 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.

- If you choose `n` or `q`, you need to complete step 14 to step 20 and run rolling upgrade phase 2: See [“Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2”](#) on page 142.

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

**17** 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 instance -d db_name
```

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

**19** Before you proceed to phase 2, complete step 4 to 18 on the remaining subcluster.

**20** Perform one of the following steps:

- If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If the VCS does not manage the Oracle database, change the management policy for the database to automatic:

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

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

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

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

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

```
./installrp -rollingupgrade_phase2 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 6.0, make sure that you upgraded all application clusters to version 6.0. Then, upgrade VCS or SFHA on the CP server systems.

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

## Verifying software versions

To verify the version of the software, enter the following command:

```
# pkginfo -l pkgname
```



# Rolling back Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [About rolling back Veritas Storage Foundation and High Availability Solutions 6.0 RP1](#)
- [Rolling back using the `uninstallrp` script](#)
- [Rolling back manually](#)

## About rolling back Veritas Storage Foundation and High Availability Solutions 6.0 RP1

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

Roll back of version 6.0 RP1 to the 6.0 release is supported for the following products:

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

## Rolling back using the `uninstallrp` script

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

---

**Note:** If any of the systems that you plan to roll back have encapsulated boot disks, you must reboot them after rollback.

---

### To roll back

- For SFRAC:

- 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 CRS is not under VCS Control, then enter the following command on each node of the cluster to stop CRS.

- For 10gR2:

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

- For 11gR2:

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

- 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 node execute following command to stop VCS:

```
# hstop -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 Run the `uninstallrp` command, type:

```
# ./uninstallrp node A node Bnode C...
```

- ## 9 If you performed a roll back on a system that has an encapsulated boot disk, you must reboot the system. After reboot, you may need to run `hagrp -list Frozen=1` to get the frozen SG list . Then run `hagrp -unfreeze <group> -persistent` to unfreeze all the frozen SGs manually.

- For other products:

- 1 Run the `uninstallrp` command, type:  

```
# ./uninstallrp system_list
```
- 2 If you performed a roll back on a system that has an encapsulated boot disk, you must reboot the system. After reboot, you may need to run `hagrpl -list Frozen=1` to get the frozen SG list . Then run `hagrpl -unfreeze <group> -persistent` to unfreeze all the frozen SGs manually.

## Rolling back manually

Use one of the following procedures to roll back to 6.0 manually.

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

---

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

---

## Rolling back Storage Foundation or Storage Foundation and High Availability manually

Use the following procedure to roll back to 6.0 manually.

To roll back SF or SFHA

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

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

- 4 Check if the root disk is under VxVM control by running this command:

```
# 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
```

---

**Note:** 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. There must be at least one other disk in the rootdg disk group in addition to the root disk for `vxunroot` to succeed.

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

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

- 6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

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

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

---

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

---

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

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

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

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

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

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

```
# svcadm disable -t vcs
```

- 10 If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```

- 11 Unmount /dev/odm:

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

- 12 Unload the ODM module:

```
# svcadm disable -t odm  
# modinfo | grep odm  
# modunload -i odm_mod_id
```

- 13 Unload the cluster fencing (vxfen) module:

```
# svcadm disable -t vxfen  
# modinfo | grep vxfen  
# modunload -i vxfen_mod_id
```

**14** Stop GAB and LLT in the following order:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

**15** 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
```

**16** Remove the SF 6.0 RP1 patches.

- Get the list of 6.0 RP1 patches, type:

```
# ./installrp -listpatches
```

- Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```

## Rolling back Storage Foundation Cluster File System High Availability manually

Use the following procedure to roll back to 6.0 manually.

**To roll back SFCFSHA manually**

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

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

- 4** Check if the root disk is under VxVM control by running this command:

```
# 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
```

---

**Note:** 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. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

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

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

- 6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the `vxrvg stop` command to stop each RVG individually:

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

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

```
# vxlink -g diskgroup status rlink_name
```

---

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

---

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

**8** Stop all VxVM volumes by entering the following command for each disk group:

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

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

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

**9** Stop VCS along with all its resources. Then, stop the remaining resources manually:

```
# svcadm disable -t vcs
```

**10** If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```

**11** Unmount /dev/odm:

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

**12** Unload the ODM module:

```
# svcadm disable -t odm
# modinfo | grep odm
# modunload -i odm_mod_id
```

**13** Unload the cluster fencing (vxfen) module:

```
# svcadm disable -t vxfen
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

14 Stop GAB and LLT in the following order:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

15 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
```

16 Remove the SFCFS 6.0 RP1 patches.

■ Get the list of 6.0 RP1 patches, type:

```
# ./installrp -listpatches
```

■ Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```

## Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 6.0 RP1 to VCS 6.0 on your cluster manually. To uninstall VCS, see the *Veritas Cluster Server Installation Guide*.

---

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

---

### To roll back VCS manually

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

```
# hagr -state
```

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

```
# hagr -offline -force ClusterService -sys system
```

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

```
# haconf -makerw
```

- 4 Freeze all service groups. On any node, type:

```
# hagrps -freeze service_group -persistent
```

where *service\_group* is the name of the service group. Note that the ClusterService group cannot be frozen.

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

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

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

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

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

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

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

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

- 9 Verify that VCS has shut down. On any node, type:

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

The output resembles: GAB Port Memberships Port a gen 23dc0001 membership 01 The output shows no membership for port h.

- 10 For Solaris 10, on nodes that run non-global zones, check if the non-global zones are in the running state. Boot the non-global zones that are not in the running state.

- Check the zone's state. On each node, type:

```
# zoneadm list -icv
```

- Boot the zone if it is not in the running state. On each node, type:

```
# zoneadm -z zone boot
```

where *zone* is the name of the non-global zone.

---

**Note:** Do not configure one or more Solaris zones to boot from the shared storage.

---

- 11 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:

```
# /sbin/vxfenconfig -U
```

- 12 Unload vxfen. On each node, perform the following steps:

- Identify the vxfen kernel module, for example:

```
# modinfo | grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 6.0 RP1)
```

- Unload vxfen using the module number.

```
# modunload -i 210
```

- 13 Unconfigure GAB. On each node, type:

```
# /sbin/gabconfig -U
```

- 14 Unload GAB. On each node, perform the following steps:

- Identify the GAB kernel module. For example:

```
# modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 6.0 RP1)
```

- Unload GAB using the module number:

```
# modunload -i 149
```

- 15 Unconfigure LLT. On each node, perform the following steps:

- Type:

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

- Type **y** on each node in response to the message.
- 16** Unload LLT. On each node, perform the following steps:
- Identify the LLT kernel module. For example:

```
# modinfo | grep llt
147 50ca4000 d6bc 110 1 llt (LLT 6.0 RP1)
```
  - Unload LLT using the module number:

```
# modunload -i 147
```
- 17** Remove the VCS 6.0 RP1 patches. On each node, perform the following steps:
- Get the list of 6.0 RP1 patches, type:

```
# ./installrp -listpatches
```
  - Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```
- 18** Verify that the patches have been removed. On each node, type:

```
# showrev -p | grep VRTS
```
- 19** If the LLT, GAB, or VXFEN modules cannot be stopped or unloaded following the patch removal, reboot all nodes in the cluster.
- 20** If you do not perform step 19, start the VCS components manually. On each node, type:

```
# /sbin/lltconfig -c
# /sbin/gabconfig -cx
# /sbin/vxfenconfig -c
# /opt/VRTSvcs/bin/hastart
```
- You do not have to start vxfen unless you use the fencing option.
- 21** After VCS has started, perform the following steps:
- Verify all resources have been probed. On any node, type:

```
# hastatus -summary
```

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

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

where *service\_group* is the name of the service group.

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

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

where *system* is the node name.

## Rolling back Symantec VirtualStore manually

Use the following procedure to roll back to 6.0 manually.

### To roll back SVS manually

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

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

- 4 Check if the root disk is under VxVM control by running this command:

```
# 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
```

---

**Note:** 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. There must be at least one other disk in the rootdg disk group in addition to the root disk for vxunroot to succeed.

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

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

- 6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the vxrvrg stop command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

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

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

---

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

---

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

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

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

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

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

- 9** Stop VCS along with all its resources. Then, stop the remaining resources manually:

```
# svcadm disable -t vcs
```

- 10** If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```

- 11** Unmount `/dev/odm`:

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

- 12** Unload the ODM module:

```
# svcadm disable -t odm  
# modinfo | grep odm  
# modunload -i odm_mod_id
```

- 13** Unload the cluster fencing (`vxfen`) module:

```
# svcadm disable -t vxfen  
# modinfo | grep vxfen  
# modunload -i vxfen_mod_id
```

- 14** Stop GAB and LLT in the following order:

```
# svcadm disable -t gab  
# svcadm disable -t llt
```

**15** 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
```

**16** Remove the SVS 6.0 RP1 patches.

- Get the list of 6.0 RP1 patches, type:

```
# ./installrp -listpatches
```

- Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```

## Rolling back Dynamic Multi-Pathing manually

Use the following procedure to roll back to 6.0 manually.

**To roll back DMP manually**

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

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

- 4 Check if the root disk is under VxVM control by running this command:

```
# 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
```

---

**Note:** 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. There must be at least one other disk in the rootdg disk group in addition to the root disk for vxunroot to succeed.

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

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

- 6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the vxrvrg stop command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

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

```
# vxlink -g diskgroup status rlink_name
```

---

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

---

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

**8** Stop all VxVM volumes by entering the following command for each disk group:

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

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

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

**9** Stop VCS along with all its resources. Then, stop the remaining resources manually:

```
# svcadm disable -t vcs
```

**10** If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```

**11** Unmount /dev/odm:

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

**12** Unload the ODM module:

```
# svcadm disable -t odm  
# modinfo | grep odm  
# modunload -i odm_mod_id
```

**13** Unload the cluster fencing (vxfen) module:

```
# svcadm disable -t vxfen  
# modinfo | grep vxfen  
# modunload -i vxfen_mod_id
```

**14** Stop GAB and LLT in the following order:

```
# svcadm disable -t gab  
# svcadm disable -t llt
```

**15** 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
```

**16** Remove the DMP 6.0 RP1 patches.

■ Get the list of 6.0 RP1 patches, type:

```
# ./installrp -listpatches
```

■ Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```