Veritas Storage Foundation™ and High Availability Solutions Release Notes

Solaris

5.1 Service Pack 1 Rolling Patch 2

December 2011

Symantec™

Veritas Storage Foundation and High Availability Solutions Release Notes 5.1 Service Pack 1 Rolling Patch 2

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

Product version: 5.1 SP1 RP2

Document version: 5.1SP1RP2.2

Legal Notice

Copyright © 2011 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.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

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

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

Chapter 1

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Rolling Patch 2 release.

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

http://www.symantec.com/docs/TECH75362

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

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

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

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

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

This rolling patch is available as:

- 5.1 SP1 RP2
- 5.1 SP1 PR3 RP2

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

- 5.1 SP1 PR3 RP1

About the installrp and the uninstallrp scripts

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 provides an upgrade script.

See "Supported upgrade paths" on page 102.

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

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

Command Line Option	Function
`[ -jumpstart jumpstart_path]`	Use the `-jumpstart` option to generate finish scripts. You can use the finish scripts with the Solaris JumpStart Server to automate installation of all packages and patches for every product. You need to specify an available location to store the finish scripts as a complete path. The `-jumpstart` option is supported on Solaris only.
`[ -keyfile ssh_key_file ]`	Use the `-keyfile` option to specify a key file for SSH. When you use this option the `-i ssh_key_file` is passed to every SSH invocation.
`[ -patchpath patch_path ]`	Use the `-patchpath` option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by `installrp`.
`[ -rootpath root_path ]`	Use the `-rootpath` option to re-root the installation of all packages to the given path. On Solaris, -rootpath passes -R <root_path> to pkgadd.

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

Command Line Option	Function
`[ -rsh | -redirect | -listpatches | -makeresponsefile | -pkginfo | -serial | -version ]`	Use the `-rsh` option when rsh and rcp are to be forced for communication though ssh and scp is also setup between the systems. Use the `-redirect` option to display progress details without showing the progress bar. Use the `-listpatches` option to display product patches in the correct installation order. Use the `-makeresponsefile` option to generate a response file without doing an actual installation. The text that displays install, uninstall, start, and stop actions are simulations. These actions are not performed on the system. Use the `-pkginfo` option is used to display the correct installation order of packages and patches. This option is available with or without one of following options: -allpkgs, -minpkgs, and -recpkgs. Use the `-serial` option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion. Use the `-version` option to have the installer check and report 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.

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

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

The uninstallrp script options

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 provides a new uninstallation script.

See About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 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. Do not use the `uninstallrp` script for rolling back, because it removes the entire stack.

Table 1-2 The command line options for the product upgrade script

Command Line Option	Function
`[ system1 system2... ]`	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.
`[ -logpath log_path ]`	Use the `-logpath` option to select a directory other than `/opt/VRTS/install/logs` as the location where the uninstallrp log files, summary file, and response file are saved.

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

Command Line Option	Function
`[ -responsefile response_file ]`	Use the `-responsefile` option to perform automated installations or uninstallations using information stored in a file rather than prompting for information. *`response_file`* is the full path of the file that contains configuration definitions.
`[ -tmppath tmp_path ]`	Use the `-tmppath` option to select a directory other than `/var/tmp` as the working directory for `uninstallrp`. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
`[ -hostfile hostfile_path ]`	Use the `-hostfile` option to specify the location of a file containing the system names for installer.
`[ -keyfile ssh_key_file ]`	Use the `-keyfile` option to specify a key file for SSH. When you use this option the `-i ssh_key_file` is passed to every SSH invocation.
`[ -rootpath root_path ]`	Use the `-rootpath` option to re-root the installation of all packages to the given path. On Solaris, -rootpath passes -R <root_path> to pkgadd or pkgrm.

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

Command Line Option	Function
`[ -rsh | -redirect | -makeresponsefile | -serial | -version ]`	Use the `-rsh` option when rsh and rcp are to be forced for communication though ssh and scp is also setup between the systems. Use the `-redirect` option to display progress details without showing the progress bar. Use the `-makeresponsefile` option to generate a response file without doing an actual installation. Text displaying installation, uninstallation, start and stop operations are simulations. These actions are not being performed on the system. Use the `-serial` option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion. Use the `-version` option to have the installer check and report 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.

Changes introduced in 5.1 SP1 RP2

This section lists the changes in 5.1 SP1 RP2.

Changes related to Storage Foundation and High Availability

Storage Foundation and High Availability includes the following changes in 5.1 SP1 RP2:

New README files provide detailed information on the included patches

This release includes README_SYMC files that provide detailed information on the patches included in this release. These README_SYMC files provide the following information:

- Patch IDs, incidents fixed through the patch
- Symptom, description, and resolution for the addressed issues

The README_SYMC files are available in the following directory:

- `/patches`

The Rolling Patch Release Notes continue to provide a summary of fixed issues.

Oracle VM Server 2.0 for SPARC

Storage Foundation High Availability Solutions is now qualified with Oracle VM Server 2.0 for SPARC.

See the *Veritas Storage Foundation and High Availability Solutions 5.1 SP1 Virtualization Guide* for supported use cases.

Changes related to installing, upgrading and rolling back

The following changes are related to installing, upgrading and rolling back of the product in 5.1 SP1 RP2 release.

Use the installrp or uninstallrp script with the –version option to determine product versions

To determine a product's version, use the `-version` option with installrp or uninstallrp. After you install 5.1 SP1 RP2, only the installrp and uninstallrp scripts can detect 5.1 SP1 RP2 versions.

About rolling back the RP installation

Use the `uninstallrp` script when you want to remove this Veritas rolling patch (RP) from systems. Once you complete the roll back process, your systems revert to the previous version of the product. You can use the roll back feature with the following products: Storage Foundation, Storage Foundation and High Availability, Veritas Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, Veritas Cluster Server, Symantec VirtualStore, Veritas Dynamic Multi-Pathing, and Veritas Volume Manager.

Changes to the supported Solaris operating systems

5.1SP1RP2 is now supported on Solaris 10 Update 10. Before upgrading to Solaris 10 Update 10, install the VxVM patch 5.1SP1SP1RP2HF1. Contact Symantec customer support to download the patch.

Changes related to Veritas Volume Manager

Veritas Volume Manager includes the following changes in 5.1 SP1 RP2.

ASM co-existence now enabled by default

The VxVM and ASM co-existence is now enabled by default. VxVM now identifies ASM disks automatically.

support vxcdsconvert utility for EFI disks.

The vxcdsconvert utility is now supported for EFI disks.

Changes related to Veritas Cluster Server

Veritas Cluster Server includes the following changes in 5.1 SP1 RP2:

New attribute for Zone agent

Symantec has modified the Zone agent by adding a new attribute `DeleteVCSZoneUser` to the Zone type. If this attribute is enabled for non-secure clusters, the Zone agent deletes the zone user from the VCS configuration. This attribute is disabled by default.

The Zone agent no longer creates the VCS zone user. You need to run the `hazonesetup` to add the VCS zone user to the cluster configuration on each node of the cluster. After you add the VCS zone user to the configuration, the user remains persistent during the lifetime of the cluster. You should remove the user when the zone resources are no longer required in the cluster.

The full value displays when you use hares -display

For `hares -display`, the maximum number of characters was 20. Any attribute value that was more than 20 characters would be truncated.

Symantec has removed the 20 characters limit. Now the full value displays when you use `hares -display`.

Changes to the hazonesetup command

Symantec enhanced the `hazonesetup` command to support creation of parallel service group for local zones. This helps you to run zones parallely on different systems.

The new usage for hazonesetup command is as below:

Usage:`hazonesetup` *`sg_name res_name zone_name passwordd autostart parallel systems`*

Where

sg_name: Name of the zone service group to be configured

res_name: Name of the zone resource in VCS configuration

zone_name: Name of the zone configured on the system

password: Password for the VCS user you want to create for zone

autostart: (0/1) Populate AutoStartList for the group

parallel: Set to 1 if you want to configure parallel service group, 0 otherwise

systems: Names of systems on which service group can run

Changes to the Application agent

The following changes are made to the Application agent in 5.1 SP1 RP2.

- Symantec enhanced the Application agent to support use of shared disks for StartProgram, StopProgram, CleanProgram and MonitorProgram.
- The Application agent now supports MonitorProcesses attribute with strings more than 80 characters.
  You need to use the `/usr/ucb/ps -ww` *`pid`* command output to configure the MonitorProcesses attribute with the COMMAND name.
- The Application agent no longer uses StopProgram if CleanProgram not specified during clean operation.

Restrict the max value of FencingWeight to 9999

When Preferred Fencing is enabled, the max weight value that you can assign to the FencingWeight attribute is 9999. `had` adds 1 to every weight that you assign for each node. If the value of FencingWeight is set to 10000, VCS fails to set the node weight to 10001 because the maximum node weight accepted by vxfen driver is 10000.

Below message displays if the value of FencingWeight that you assign is equal to or more than 10000.

```
VCS WARNING V-16-1-50003 FencingWeight should be an \
integer between [0..9999] inclusive
```

Changes to Application Agent

The following are changes about Application Agent for shared disk support

Shared disk support

Symantec enhanced the Application agent to support use of shared disks for StartProgram, StopProgram, CleanProgram and MonitorProgram attributes.

Support for UNIX style return values in MonitorProgram

The Application agent handles the standard UNIX style return values, that is, "0" for success and "1" for failure, and now reports the resource state based on following set of values:

100 or 1 --> OFFLINE

101 to 110 or 0 --> ONLINE

Any other value --> UNKNOWN.

Validating the user that is specified in the User attribute

The Application agent is modified to validate the user name configured in the User attribute. If the user does not exist on the system, the agent reports the state of a configured resource on that system as UNKNOWN.

User home directory validation for the user that is specified in the User attribute

The Application agent checks for presence of home directory for the user only if UseSUDash is set to 1 and home directory for the user is set in the /etc/passwd.

System requirements

This section describes the system requirements for this release

Supported Solaris operating systems

This release of the Veritas products is supported on the following Solaris operating systems:

- SPARC
  - Solaris 10 Update 6, 7, 8, 9, 10

- Solaris 9 Update 7, 8 and 9 + latest recommend cluster patch
  Solaris 9 Update 9 is based on Update 8 but added hardware support for the V445, V245 and Ultra 45 SPARC platforms

- x64
  - Solaris 10 Update 6, 7, 8, 9,10

If necessary, upgrade Solaris before you install the Veritas products.

Install all the latest required Solaris patches listed in this *Release Notes*.

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

http://www.symantec.com/docs/TECH75362

Required Solaris patches

Before installing Veritas product, ensure that the correct Solaris patches are installed.

See http://support.oracle.com for the latest Solaris patch updates.

The following patches (or a later revision of those patches) are required for Solaris SPARC:

Table 1-3 Solaris SPARC patches

Operating system	Oracle patch number

Solaris 9	`114477-04`
	`122300-29 (required for Live Upgrade)`
	
	Patches required for FS, SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC:
	■ For Solaris 9 Update 7 (to bring OS kernel to FS patch 122300 level):
	`117171-17`
	`113073-14`
	■ For Solaris 9 Update 7, Update 8 and Update 9:
	`112908-33 (required for 122300-10)`
	`118558-39 (required for 122300-10)`
	`122300-10`

Table 1-3 Solaris SPARC patches *(continued)*

Operating system	Oracle patch number
Solaris 10	119254-06 119042-02 125731-02 128306-05 127111-01

The following patches (or a later revision of those patches) are required for Solaris x64:

Table 1-4 Solaris x64 patches

Operating system	Oracle patch number
Solaris 10	118344-14 118855-36 119043-11 119131-33 120012-14 125732-05 127128-11

Database requirements

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

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

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

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

Recommended memory and swap space

Symantec recommends the following memory and swap space sizes:

- On the system where you run the installation, use the following guidelines for memory minimums when you install on:

- One to eight nodes, use 1 GB of memory
- More than eight nodes, use 2 GB of memory or more

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

List of products

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

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

Fixed issues

This section describes the issues fixed in this release.

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

Veritas Volume Manager fixed issues

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

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP2

Table 1-5 describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP2.

Table 1-5 Veritas Volume Manager 5.1 SP1 RP2 fixed issues

Fixed issues	Description
2480600	I/O permanent hung on master node when IO size larger than 512K, and 32+ threads write in parallel
2440349	DCO volume may grow into any 'site' even when 'alloc=site:xxxx' is specified by a list of 'site' to be limited
2431470	vxpfto uses DM name when calling vxdisk, but vxdisk will match DA name first and thus cause corruption
2431423	CVR: Panic in vol_mv_commit_check after I/O error on DCM
2428875	I/O on both nodes (wait for the DCM flush started), and crash the slave node, lead to the master reconfiguration hang
2428631	Allow same fence key to be used for all Disk groups
2425722	vxsd move operation failed for disk size greater than or equal to 2 TB
2425551	IO hung for 6 mintues when reboot the slave node, if there is I/O on both master and slave
2424833	while autosync and deport is ongoing the primary logowner hits ted assert nmcom_send_msg_tcp
2421067	Vxconfigd hung in both nodes of primary
2419348	DMP panic: race between dmp reconfig and dmp pass through ioctl
2413904	Multiple issues are seen while performing Dynamic LUN reconfiguration
2411698	VVR:iohang: On I/O to both master and slave
2410845	Lots of 'reservation conflict' messages seen with XIV arrays
2408771	vxconfigd does not scan and discover all the storage device; some storage devices are skipped
2407192	Application I/O hangs because of race between CVM reconfiguration and Log-owner change protocol
2406292	Panic in vol_subdisksio_delete()

Table 1-5 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2400654	Stale array.info file can cause vxdmpadm commands to hang
2396293	I/Os loaded, sanboot failed with vxconfigd core dump
2387993	While testing including/excluding libvxpp.so vxconfigd goes into disabled mode
2386120	Enhancement request to add diagnostic logging to help triage a CVM master takeover failure situation
2385680	VVR: vxio panic in vol_rv_async_childdone+1147
2383158	VVR: vxio panic in vol_rv_mdship_srv_done+680
2379029	Changing of enclosure name is not working for all devices in enclosure
2369786	VVR:A deadloop about NM_ERR_HEADR_IO
2365951	Growto failing with error V-5-1-10128 Unexpected kernel error in configuration update
2364253	VVR: Kernel memory is leaked on VVR secondary while using SO snapshots
2359814	vxconfigbackup doesn't handle errors well
2357798	CVR:Memory leak due to unfreed vol_ru_update structure
2357507	In presence of large number of NR (Not-Ready) devices, server panics due to NMI triggered and when DMP continuously generates large no of path disable/enable events
2356744	VxVM script daemons should not allow its duplication instance in itself
2349352	During LUN provisioning in single path IO mode environment a data corruption is observed
2346470	Excluding and including a LUN in a loop triggers a huge memory leak
2337694	TP: "vxdisk -o thin list" showing size 0 for over 2TB LUNs
2337353	vxdmpadm include vxvm dmpnodename= includes all excluded dmpnodes along with the requested one
2334534	In CVM environment, vxconfigd level join is hung when Master returns error "VE_NO_JOINERS" to a joining node and cluster nidmap is changed in new reconfiguration

Table 1-5 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2323925	If rootdisk is encapsulated and if install-db is present, clear warning should be displayed on system boot
2322752	Duplicate DA records seen for NR devices upon restart of vxconfigd
2320917	vxconfigd core dump and lost dg config after removing volume and disk on thin reclaim LUN
2317703	Vxesd/Vxconfigd leaks file descriptors
2316297	Error message "Device is in use" appears during boot time
2299670	Disk Groups created on EFI LUNs do not auto import at boot time using VxVM version 5.1SP1 and later
2286559	kernel heap corruption detected panic after array controller reboot
2263317	CLONE: Diskgroup import with dgid needs to be clearly documented in manual for the case in which original dg was destroyed and cloned disks are present
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated
2255182	Handling misconfiguration of CLARiiON array reporting one failovermode value through one HBA and different from other HBA
2253970	Support per-disk maxiosize for private region I/Os
2253552	Leak in vxsfdefault_parse.y at function vxsf_getdefault (*val)
2249113	vol_ru_recover_primlog_done return the same start address to be read from SRL, if the dummy update is greater than MAX_WRITE
2248730	vxdg import command hangs as vxrecover daemon (spawned by vxdg) doesn't close standard error stream
2242268	panic in voldrl_unlog
2240056	vxdg move' transaction not completing and backups fail
2237089	vxrecover might start the recovery of data volumes before the recovery of the associated cache volume is recovered
2232789	supporting NetApp Metro Cluster
2228531	cvm master vxconfigd process hung in vol_klog_lock()

Table 1-5 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2205108	SVS: vxconfigd clubbing all luns in a single dmpnode
2204752	Multiple VM commands succeed but throw "GPT entries checksum mismatch" error message for hpdisk format
2200670	vxattachd does not recover disks if disk group is not imported
2197254	While creating volumes on thinrclm disks, the option "logtype=none" does not work with vxassist command
2196918	Snapshot creation with cachesize fails, as it doesn't take into account diskgroup alignment
2196480	The disk initialization failed due to wrong number of cylinders reported in devintf_disk_geom_raw() gotten from raw geometry
2194685	vxconfigd daemon core dump during array side switch ports disable and re-enable
2193429	IO policy not getting preserved when vold is restarted and migration from one devlist to other is taking place
2190020	dmp_deamon applying 1m continuous memory paging which is too large
2165394	CLONE: dg imported by selecting wrong disks. After destroying original dg, when try to import clone devices without useclonedev option with dgname, then it import dg with original disks
2154287	Improve handling of Not-Ready(NR)devices which are triggering "VxVM vxdmp V-5-3-1062 dmp_restore_node: Unstable path" messages
2152830	In multilevel clone disks environment, regular DG import should be handled properly and in case of DG import failure, it should report correct error message
2144775	Failoverpolicy "local" is not getting preserved after upgrade from 5.1RP1
2139179	SSB check invalid when lun copy
2094672	CVR: vxconfigd on master hangs while reconfig is running in cvr stress with 8 users
2033909	In SF-RAC configuration, IO hung after disable secondary path of A/PG array Fujitsu ETERNUS3000

Table 1-5 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2484685	Race between two vol_subdisk sios while doing 'done' processing which causes one thread to free sio_fsvm_priv before other thread accesses it
2369177	DDL: do_diskio function should be able to handle offset greater than 2TB
2346470	Excluding and including a LUN in a loop using vxdmpadm, triggers a huge memory leak
2339251	Newfs fails for volumes greater than 2Tb due to DKIOCGMEDIAINFO and DKIOCGETEFI ioctls failing
2320669	SDS to VxVM conversion failed on Solaris
2291176	vxrootadm does not set dump device correctly with LANG=ja
2270880	On Solaris 10 SPARC, disk capacity will be truncated to 2TB for EFI labeled disk with size greater than 2TB
2268408	suppressing a powerpath disk's path using vxdiskadm 17-2 causes the disk to go in error state
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated
2244210	vxdisksetup failed with default parameters passed through CLI on EFI setup
2230716	SVM migration to VxVM fails to convert due to /etc/lvm/md.cf not being cleared
2226304	Cannot create 1TB+ ufs file system with Solaris 9
2216515	vxunreloc may corrupt boot disk, if original offsets are used
2215262	Netapp iSCSI LUN goes into error state while initializing via VEA GUI
2215256	Support of Oracle F5100 flash array on an Oracle X4470 server with new HBA
2196480	Initialization of VxVM cdsdisk layout fails on a disk of size less than 1 TB
2179259	DMP SCSI bypass needs to be enhanced to handle I/O greater than 2TB
2146833	vxmirror operation fails with error "Device has UFS FS on it"
2108152	vxconfigd goes to DISABLED state upon system reboot with NR (Not Ready) devices in configuration
2064490	Ensure vxcdsconvert works for greater than 1 TB CDS disks

Table 1-5 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2063152	vxbootadm failed with xml related error
1791397	VVR:RU thread keeps spinning sending START_UPDATE message repeatedly to the secondary
1675599	Memory leaks in DDL and ASLs

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP1

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

Table 1-6 Veritas Volume Manager 5.1 SP1 RP1 fixed issues

Fixed issues	Description
2160199	Master takeover fails as the upcoming Master could not import shared DG. SUN Bug ID is 6988420.
2148682	while shipping a command node hangs in master selection on slave nodes and master update on master node
2181631	Striped-mirror volume cannot be grown across sites with -oallowspansites with DRL
2133503	Renaming enclosure results in dmpevents.log reporting Mode for Enclosure has changed from Private to Private
2200670	vxattachd does not recover disks if disk group is not imported
2191693	'vxdmpadm native list' command is not displaying any output nor error
2080730	vxvm/vxdmp exclude file contents after updation should be consistent via vxdiskadm and vxdmpadm
2129989	EVA ASL should report an error message if pref_bit is not set for a LUN
2129477	vxdisk reclaim command fails after resize operation.
2194492	VxVM-ASM co-existence enablement
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core

Table 1-6 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2166682	checks needed to make sure that a plex is active before reading from it during fsvm mirror read interface
2088007	Possibility of reviving only secondary paths in DMP
2188590	An ilock acquired by a SLAVE node for a read on a DCL object can lead to IO hang when the node becomes MASTER before reading is done
2015467	Performance improvement work for NetBackup 6.5.5 on SF 5.1 VxVM mapping provider
1829285	vxconfigd coredumps while assigning unique native name to a disk
2226813	VVR: rlinks remain disconnected with UDP protocol if data ports are specified
2227923	renaming of enclosure name is not persistent
2172488	FMR: with dco version 0 restore operation doesn't sync the existing snapshot mirrors
2183984	System panics due to race condition while updating DMP I/O statistics
2038928	creation of pre 5.1SP1 (older) version diskgroup fails
1970560	When vxconfigd is idle (which is not shipping the command) slave dies and command shipping is in progress, vxconfigd core dumped on Master
2082450	In case of failure, vxdisk resize should display more meaningful error message
1869002	Introduction of Circular buffer at vold level for master-slave communication.
1426480	VOLCVM_CLEAR_PR ioctl does not propogate the error returned by DMP to the caller
2105547	tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault
1959513	Propagate -o noreonline option of disk group import to slave nodes
1940052	vxconfigd hung on Master after removing the hba alias from zone and node leave followed by join
2199496	Data Corruption seen with "site mirror" Campus Cluster feature

Table 1-6 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2234844	asm2vxfs conversion fails
2215216	vxkprint does not report TP related values
2141820	vxrootadm:boot entry not created after second split on Solaris opetron
2211283	system goes in maintainence mode after removal of 51SP1 patch with boot disk encapsulated
2148738	vxdmpadm invoked by vxvm-sysboot is killed during system boot
2062190	vxrootadm split/join operation fails when there is a rvg present in the rootdg/backupdg
2130744	1TB disks are seen in error state on sol9u7 and vxdisksetup also fails.
2121183	installmp/rp/patchadd results in a partially installed state/unusable
2164906	Need to enhance vxlustart script to support new option to luupgrade command in Sol10U9
2209391	'vxdisk init' failure modifies disk label from sliced to EFI type for devices >1TB on Solaris 10

Veritas File System fixed issues

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

Veritas File System: Issues fixed in 5.1 SP1 RP2

Table 1-7 describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP2.

Table 1-7 Veritas File System fixed issues

Fixed issues	Description
2529356 (2340953)	cfs.stress.enterprise hit an assert f:vx_iget:1a.

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

Fixed issues	Description
2508164 (2481984)	file system will hang if customer creates 400 shares
2494464 (2247387)	LM stress.S3 test hit an assert "vx_ino_update:2"
2486597 (2486589)	threads blocked behind vx_ireuse_steal
2482337 (2431674)	panic in vx_common_msgprint() via vx_inactive()
2480949 (2480935)	fsppadm: ERROR: V-3-26626: File Change Log IOTEMP and ACCESSTEMP index creation failure for /vx/fsvm with message Argument list too long
2478237 (2384861)	CFS stress+reconfig test hit assert "f:vx_do_filesnap:1b".
2439526 (2333907)	LM Conformance->fsmig->migvops test get fail.
2427281 (2413172)	There is a priority 1 issue reported by AXA Rosenburg for Filestore replication and issue seems related to VxFS
2427269 (2399228)	TRuncate up size updates can be missed
2426039 (2412604)	it does not work when set homedir user softlimit numspace quota after generate data
2425429 (2422574)	Reboot one node and the node can't mount file system , after turn on the homedir quota on
2418819 (2283893)	Add functionality of free space defragmentation through fsadm.
2412181 (2372093)	new fsadm -C hung

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

Fixed issues	Description
2412179 (2387609)	User quota corruption
2412177 (2371710)	user quota information corrupts on 5.1SP1
2412029 (2384831)	vxfs panic in iput() from vx_softcnt_flush() ,after filesystem full fsck,and run reboot
2407896 (2407895)	System panic in atomic_add_32_nv because of bad address.
2406572 (2146573)	qdetails performance downgraded on Aug 16th.
2402643 (2399178)	fsck : pass2c needs performance enhancements
2386483 (2374887)	Accessing FS hung. FS marked full fsck after reboot of node.
2373565 (2283315)	cfs-stress_S5 hits assert of "f:vx_reorg_emap:10 via vx_extmap_reorg"
2368738 (2368737)	RCQ processing code should set FULLFSCK flag if it finds a corrupt indirect block.
2360821 (1956458)	fsckpt_fbmap for changed blocks failed with ENXIO due to inode mapped to hole in ILIST of down stream checkpoint
2360819 (2337470)	In the process of shrink fs, the fs out of inodes, fs version is 5.0MP4HF*
2360817 (2332460)	vxedquota slow on some systems
2341007 (2300682)	Question about IOTemp on fsppadm query

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

Fixed issues	Description
2340831 (2272072)	Threads stuck in vx_rwsleep_rec_lock_em
2340825 (2290800)	investigation on ilist HOLE
2340817 (2192895)	VxFS 5.0MP3RP4 Panic while set/get acls - possible race condition
2340811 (2172485)	Metadata was not updated correctly after write() with O_SYNC flag.
2340799 (2059611)	Panic in vx_unlockmap() due to NULL ml_tranp
2340741 (2282201)	vxdump core dumped whilst backing up layout 7 local VxFS file system
2329893 (2316094)	There was discrepancy between vxi_bcache_maxkbyte and vx_bc_bufhwm.
2320044 (2419989)	ncheck -i does not limit output to the specified inodes when using -o device/block/sector
2311490 (2074806)	dm_punch_hole request does not invalidate pages
2296277 (2296107)	Operation not applicable appear on fsppadm query result
2280552 (2246579)	Panic at getblk() when growing a full filesystem with fsadm
2280386 (2061177)	fsadm -de' command erroring with 'bad file number' on filesystem(s) on 5.0MP3RP1
2275543 (1475345)	write() system call hangs for over 10 seconds on VxFS 3.5 on 11.23

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

Fixed issues	Description
2271886 (2271878)	WARNING message with fsmigadm start command
2257904 (2251223)	df -h after removing files takes 10 seconds
2255786 (2253617)	LM stress aborted due to "run_fsck : Failed to full fsck cleanly".
2249658 (2220300)	vx_sched' is hogging CPU resources.
2248489 (2226762)	vx_ntran overflow causes a lot of buffer flushes
2243063 (1949445)	hang due to large number of files in a directory
2243061 (1296491)	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2240442 (2230351)	After installation of SF 5.1SP1, Solaris 10 system PANICed 8 times due to \"BAD TRAP: type=31 rp=2a1001c6e40 addr=28 mmu_fsr=0 occurred in module \"vxfs\" due to a NULL pointer dereference\"
2169326 (2169324)	5.1SP1 sol_sprac Test LM-stress_S5 hits an assert of "f:vx_idelxwri_off:5a vai vx_trunc_tran"

Veritas File System: Issues fixed in 5.1 SP1 RP1

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

Table 1-8 Veritas File System fixed issues

Fixed issues	Description
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.

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

Fixed issues	Description
2030119	fsppadm core dumps when analysing a badly formatted XML file, is resolved
2162822	During online migration from ufs to vxfs, df command returns a non-zero return value.
2169273	During online migration, nfs export of the migrating file system leads to system panic
2177253	A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for vxupgrade purposes
2178147	Linking a IFSOC file now properly calls `vx_dotdot_op`(), which fixes the cause of a corrupted inode.
2184528	`fsck` no longer fails to repair corrupt directory blocks that have duplicate directory entries.
2194618	Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system
2221623	Fixed a performance loss due to a delxwri_ilist spin lock with the default values for `vx_idelxwri_timelag`.
2255786	Failure to full fsck cleanly due to incorrect string comparison of file name.
2228311	Support for mostly online migration from ZFS to VxFS

Veritas Storage Foundation Cluster File System fixed issues

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

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

Table 1-9 describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP2.

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

Fixed issues	Description
2406572 (2146573)	qdetails performance downgraded
2407896 (2407895)	System panic in atomic_add_32_nv because of bad address.

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

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

Table 1-10 Veritas Storage Foundation Cluster File System fixed issues

Fixed issues	Description
1296491	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2149659	In the presence of file system checkpoints containing snapshot files, truncate operation of a file with shared extent results in hitting f:xted_validate_cuttran:10 or vx_te_mklbtran:1b
2169538	The cfsmntadm add command fails, if one host name is a substring of another host name in the list
2180905	fsadm -S shared mountpoint gives wrong error message when mount points of veritas filesystem version is other than 8.
2181833	"vxfilesnap" gives wrong error message on checkpoint filesystem on cluster
2184114	In a large filesystem, stat operation on cluster mount leaves the file system frozen for too long leading to CVMVoldg and CFSMount agent timeouts.
2203917	ODM I/O performance bottleneck due to threads waiting in odm_rwsleep_lock() to get lock on odm_iop_table is resolved
2232554	System panic in vx_iupdat_clustblks() due to an unprotected inode getting corrupted.
2241123	glmdump failure with the error “/opt/VRTSglm/sbin/glmdump: syntax error at line 340: `count=$' unexpected” is resolved.

Veritas Storage Foundation for Oracle RAC fixed issues

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

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

Table 1-11 describes the incidents that are fixed in Veritas Storage Foundation for Oracle RAC in 5.1 SP1 RP2.

Table 1-11 Veritas Storage Foundation for Oracle RAC fixed issues

Fixed issues	Description
2374977	Oracle instance crashed; failure occurred at: vcsipc_dosnd
2390892	memory leak in vcsmm_set_cluster_proto
2374970	Update the CRSResource agent to avoid using -n option for srvctl on 11gR2
2380469	SFRAC support inside Solaris Local zones
2429449	The cssd agent explicitly uses hard-coded string "cssd" as resource name.

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

There are no fixed issues in this release.

Veritas Cluster Server fixed issues

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

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP2

Table 1-12 describes the incidents that are fixed in Veritas Cluster Server in 5.1 SP1 RP2.

Table 1-12 Veritas Cluster Server 5.1 SP1 RP2 fixed issues

Fixed Issues	Description
2423740	Offline registration for a process with AMF succeeds even when the process is running.
2423819	System cannot unload the AMF driver.
2533303	Fixed Oracle agent so that Oracle resource configured inside locale zone goes into offline state when the Zone is in 'DOWN' state.
2528475	Fixed preonline_ipc to support IPMultiNIC/ IPMultiNICB type in preonline_ipc
2528317	Fixed issue with agent framework where Application agent logs the message "Resource(app_res): Output of the completed operation (monitor)" continuously.
2516807	Fixed issue with Application agent to support physical to virtual failover.
2514514	Added support for benchmarking tool ab2 as well as ab to Apache.
2513928	Fixed issue with Mount agent while supporting physical to virtual failover.
2512840	Fixed the issue with Oracle agent to support physical to virtual failover.
2511385	Fixed issue with Sybase where online script marks the database as online before Database has recovered
2508114	Fixed issue with Sybase agents to support physical to virtual failover.
2491635	Fixed issue with DiskGroup which throws ERROR message when 5.1 SF is used.
2485202	Fixed issue with Application agent where monitor fails when 80th character of psargs is whitespace
2483044	Fixed issue with VCS where 'had' crashed with SIGSEGV when asserting against gp->activecount()->gets32GL(nodename) == 0\", in "Resource.C" in check_failover function
2477372	Reduced lltd CPU consumption by reducing the wakeup calls
2477296	Fixed issue with VCS where Application service group did not fail over when the node panic
2477280	Fixed issue with VCS where Application resource failed to failover when system reboot after concurrency violation

Table 1-12 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2477268	Fixed Zone agent to support physical to virtual failover.
2476897	Fixed issue with IP agent while supporting physical to virtual failover.
2439895	Fixed issue with lltconfig which reports its own cluster node as part of duplicate cluster
2439772	Fixed issue with VCS where wac resource offline failed after network interruption
2439695	Fixed issue with VCS where VXFEN module gets loaded even though user chooses not to enable VXFEN.
2438261	Fixed issue with vxfen where it fails to perform online migration from scsi raw to scsi dmp policy.
2434782	Fixed issue with VCS to allow ContainerInfo attribute to be updated even when Group is not completely offline
2433347	Fixed issue with Application Agent where it looks for Home Directory in Global Zone for a particular user even if it configured for local zones.
2427464	Enhanced Multinic Agent to support for dual stack IPv4 and IPv6 (issue with IPv6)
2426663	Fixed issue with vxfen where vxfend does not terminate on OCPR from customized mode to scsi3 mode
2426572	Fixed issue with VCS where persistent resource is reported OFFLINE (not FAULTED) when system is added to the group using hagrp -modify command
2423990	Fixed issue with Application Agent which fails to work correctly when nonexistent user is configured.
2423838	Enhanced IP Agent to log ERROR message for a failure instead of DEBUG message.
2416956	Fixed Application agent not to call StopProgram when CleanProgram is not defined.
2416842	Fixed issue with VCS where had consuming over 99% CPU time. Multiple ha commands are hung in pollsys()
2411860	Fixed issue with VCS where service group switch fails

Table 1-12 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2411858	Fixed issue with agent framework where various VCS service groups fails to switch
2411653	Added check for MAX message size in GAB
2407872	Fixed shell errors in Application agent "sh: setenv: not found" error during Application online or offline
2407755	Fixed issue with agent framework where async-signal function call removed from signal handler and between fork and exec, which is casuing agents to fail.
2407653	Fixed issue with AMF where in case of forceful unload of AMF module, module reference count for 'vxfs'/'ext3' handled correctly.
2407617	Fixed VCS SMF dependency to be clean after install
2407612	Fixed AMF SMF dependency to be clean after install
2406748	Fixed issue with AMF where AMF allowing to register already online process for offline monitor with AMF.
2405391	Fixed LLT to include the nodename in the arp ack packet
2403851	Fixed issue with AMF status where it is showing module loaded but not configured.
2403782	Enhanced Sybase agent scripts to open the file and read instead of using cat command.
2403633	Enhanced agent framework to allow ContainerInfo attribute to be updated even when Group is not completely offline if not already set.
2400485	Fixed issue with vxfen vxfenconfig -c where with mode A has returned EFAULT, all subsequent runs of vxfenconfig -c with mode B fail with error EBADMSG
2400330	Fixed issue with preonline trigger where whyonlining does not behave as documented
2399898	Fixed issue with VCS where hagrp -switch of child group fails if two or more parent groups online on other nodes with 'online-global-soft' dependency.
2398807	Fixed VCS to set soft limits for file descriptors in /opt/VRTSvcs/bin/vcsenv

Table 1-12 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2394176	Fixed issue with vxfenswap process hangs, "ps -ef" shows "vxfenconfig -o modify" on one node but not on other.
2386326	Fixed issue with vxfen where vxfenadm prints same Serial Number for all LUNs which have more than 96 bytes of SCSI Inquiry data in page 0x83
2382592	Fixed issue with displaying "ResourceInfo" Attribute of SRDF Resource using hares -display
2382583	Fixed issue with CP Agent where it does not show coordination point information in engine log when CP server is not accessible.
2382559	Fixed issue with vxfen where online migration fails while migrating coordination points from CP server to disks.
2382493	Fixed issue with VCS where parent service group does not failover in case of online local firm dependency with child service groups
2382463	Fixed issue with VCS where had weight(1) is not added if we reach the boundary condition(10000) in System policy in CPS preffered fencing.
2382460	Fixed issuw with vxfen where the message 'configuring fencing is successful with 3 disks even when single_cp=1'o displays.
2382452	Fixed issue syntax errors with CPS while unconfiguring CP server using configure_cps.pl script
2382384	Fixed VXFEN SMF dependency to be clean after install
2382380	Fixed GAB SMF dependency to be clean after install
2382335	Fixed issue with vxfentsthdw which fails to choose the same fencing disk on two nodes.
2373431	Fixed LLT SMF dependency to be clean after install
2372072	Fixed issue with 'hacf' which dumps core if it could not get current working directory
2371652	Fixed issue with IPMultiNIC agent where it's resources unable to probe when underlying MultiNICA resource is in a global service group
2367721	Fixed owner.vfd file for Oracle to compare only the uid and gid of the user, instead of the complete output string of id command.

Table 1-12 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2366701	Enhanced Zone agent to support shared disks for Start/Stop/Monitor/Clean programs.
2366201	Enhanced Fencing to start when a majority of the coordination points are available.
2358616	Fixed issue with MultiNICB for nxge interface in base mode when IgnoreLinkStatus is disabled.
2354932	Fixed issue with 'hacli -cmd' which triggers had coredump on a system
2330980	Fixed issue with VCS where HAD sends a resource snapshot to agents running on existing nodes when a node is added to or removed from SystemList.
2330045	Fixed issue with RemoteGroup agent where it's resource during network failure cannot be recovered without bouncing entire cluster.
2330041	Fixed issue with VCS where the group dependencies do not online parallel parent group
2323312	Fixed issue with Application Agent where agent fails to monitor when MonitorProcesses attribute is configured with strings more than 79 characters.
2318334	Fixed issue with Oracle where it needs database's $Oracle_home/lib library to be first in LD_LIBRARY_PATH before /usr/lib
2317067	Fixed issue with Application Agent where the application resource unable to come online inside zones if pid files are specified.
2301731	Fixed panic issue in amf_lock() due to bad mutex during system shutdown.
2298775	Fixed issue with hazonesetup command to set localized zone name for parallel zone service group configuration.
2296172	Fixed issue with VCS where Failover Policy does not take into consideration AutoFailover = 2 and SystemZones when the nodes within a SystemZone are brought down/rebooted.
2276622	Fixed issue with vxfen to configure SCSI-3 fencing using RamSan DMP devices.
2275376	Enhanced Zone agent to used default BootState (multi-user) if not specified.
2271882	Fixed issue with Netlsnr where MonitorMethod attribute does not reflect IMF value without setting Listener attribute on Netlsnr Resource.

Table 1-12 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2253349	Enhanced IP agent to log a warning message when netmask changed outside of VCS
2393939	Enhanced Apache agent version parsing to accommodate IBM HTTP server 7.0.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP1

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

Table 1-13 Veritas Cluster Server 5.1 SP1 RP1 fixed issues

Fixed Issues	Description
1949294	`fdsetup` can now correctly parse disk names containing characters such as "-".
1949303	`fdsetup` no longer allows volume that are not part of the RVG, which fixes a possible cause of the RVGSnapshot agent failing.
2011536	Added IMF support for the db2udb agent.
2159991	Fixed an issue with messages in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2172181	Fixed an issue with AMF-related messages for the CAVF agent in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2175599	Fixed an issue with fencing configuration on a 64-node cluster.
2179652	The monitor script of the db2udb agent can now handle empty attribute values.
2184205	Fixed an issue with HAD in which the parent service group did not fail over if the parent service group had an online local firm dependency with a child service group.
2194473	HAD no longer dumps core while overriding the static attribute to the resource level.
2203430	Fixed an issue with haping.

Table 1-13 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Fixed Issues	Description
2205556	Fixed an issue with the offline EP of the DNS agent, which did not remove all A/AAAA records if OffDelRR=1 for multi-home records.
2205563	A clean EP now properly removes resource records when OffDelRR=1.
2205567	Fixed an issue in which having an attribute set to `master.vfd` caused the DNS agent to fail to query the DNS server.
2208416	The zone agent now properly starts the network services when the `BootState` attribute is set to multi-user-server.
2208901	Fixed an issue with the RVGSnapshot agent.
2209337	Fixed an issue with VCSAPI where the RemoteGroup agent crashed if the VCSAPI log level was set to a non-zero value.
2210281	Fixed an issue with the agent framework in which MonitorTimeStats incorrectly showed 303 seconds intermittently.
2213725	Fixed an issue with the zone agent, in which some of the system services did not start if BootState was configured.
2214539	Fixed an issue in which rebooting a node sometimes set the intentonline of a group to 2, even if the group was online somewhere else. This caused the group to use the autostartlist and not perform a failover.
2218556	Fixed an issue in the `cpsadm` command in which it sometimes failed if LLT was not installed or configured on a single node cluster.
2230869	Fixed an issue with the `IPMultiNICB.xml` file in which the VCS GUI showed the `UseMpathd` attribute for IPMutliNICB resources.
2232633	Fixed an issue with the zone agent in which offline zone resource removed all VCS zone users.
2241419	Fixed an issue in which halogin did not work in a secure environment where the root broker was not a VCS node.

Storage Foundation Manager fixed issues

This section describes Storage Foundation Manager fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Storage Foundation Manager: Issues fixed in 5.1 SP1 RP2

There are no Storage Foundation Manager fixed issues in this release.

Storage Foundation Manager: Issues fixed in 5.1 SP1 RP1

There are no Storage Foundation Manager fixed issues in this release.

Veritas Storage Foundation for Databases (SFDB) tools fixed issues

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

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

Table 1-14 describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in 5.1 SP1 RP2.

Table 1-14 Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Fixed issues	Description
2395192	vxdbed looping doing read/write to IDLE sockets

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

There are no SFDB fixed issues in 5.1 SP1 RP1.

Known issues

This section covers the known issues in this release.

Issues related to installation

This section describes the known issues during installation and upgrade.

Upgrade or uninstallation of Veritas High Availability product may encounter module unload failures (2159652)

When you upgrade or uninstall Veritas High Availability product, 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.

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

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

Workaround: Run the installer with `-nospacecheck` option if the disk space is less than that installer claims but more than actually required.

Live Upgrade may fail on Solaris 9 if packages and patches are not current (2052544)

Live Upgrade may fail on a Solaris 9 host if a VxFS file system is in `/etc/vfstab`. This is a known Solaris 9 issue fixed in Solaris Live Upgrade patch 121430-23 and later. Install the latest version of the Live Upgrade patch and the required Solaris 9 Live Upgrade packages.

For information on the required packages and patches, visit the following site and search on "Live Upgrade requirements":

http://wikis.sun.com

Incorrect version prompt when SFRAC5.1SP1RP1 upgrade to RP2 (2530303)

When you upgrade from SFRAC 5.1 SP1 RP1 to SFRAC 5.1 SP1 RP2, the previous SFRAC version is incorrectly listed as 5.1.100.000.

Workaround: This message can be safely ignored.

installrp fails to install 5.1SP1RP2 when the root user shell is set to csh (2523643)

VCS installation fails if super user (root) logged-in is using C shell (csh). Currently the installer does not support c-shell (/usr/bin/csh).

Workaround: Change your super-user (root) shell to shell (/usr/bin/sh) and retry the installation.

Veritas Storage Foundation known issues

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

Some dbed DST commands do not work correctly in non-POSIX locales (2138030)

Some dbed DST commands do not work correctly in non-POSIX locale settings.

Workaround: Set the environment variable `LANG=C` systemwide in the `/etc/profile` file.

db2exp may frequently dump core (1854459)

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

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

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

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

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

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

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

DBI1070I  Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

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

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

To communicate in a dual-stack environment

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

  127.0.0.1 swlx20-v6

  Or

  127.0.0.1 swlx20-v6.punipv6.com

  where *127.0.0.1* is the IPv4 loopback address.

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

Boot fails after installing or removing Veritas product packages from a Solaris 9 system to a remote Solaris 10 system (1747640)

The following issue occurs if you install or remove a Storage Foundation package or patch from a Sparc Solaris 9 system to a remote Solaris 10 system, using the `-R` *rootpath* option of the `pkgadd`, `patchadd`, `pkgrm` or `patchrm` commands.

Generally, when you install or remove a Veritas product package on a Solaris 10 system, the package scripts update the boot archive. However if the local system is Solaris 9 and the remote system is Solaris 10, the scripts fail to update the boot archive on the Solaris 10 system.

Note: The boot archive is synchronized correctly when you upgrade Veritas product using Solaris Live Upgrade.

Workaround: The workaround is to manually clear the boot archive when you boot the alternate. The Solaris boot process detects that the boot archive is out sync and displays instructions for how to correct the situation.

For example:

```
WARNING: The following files in / differ from the boot archive:

     stale //kernel/drv/sparcv9/vxportal
```

```
    stale //kernel/drv/vxportal.conf
    stale //kernel/fs/sparcv9/vxfs
    ...
    new     /kernel/drv/vxlo.SunOS_5.10
    new     /kernel/drv/vxlo.conf
    changed /kernel/drv/vxspec.SunOS_5.9
    changed /kernel/drv/vxspec.conf

The recommended action is to reboot to the failsafe archive to correct
the above inconsistency. To accomplish this, on a GRUB-based platform,
reboot and select the "Solaris failsafe" option from the boot menu.
On an OBP-based platform, reboot then type "boot -F failsafe". Then
follow the prompts to update the boot archive. Alternately, to continue
booting at your own risk, you may clear the service by running:
"svcadm clear system/boot-archive"
```

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

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

Tools like `dbca` may hang during database creation.

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

Sybase ASE version 15.0.3 causes segmentation fault on some Solaris version (1819595)

Sybase ASE 15.0.3 produces segmentation fault on Solaris SPARC 10 Update 6 in a pure IPv6 environment. However, Sybase ASE 15.0.3 works on Solaris SPARC 10 Update 5.

When running Sybase ASE 15.0.3 GA on a pure IPv6 environment on Solaris SPARC 10 Update 6, you may receive a segmentation fault message. For example:

```
Building Adaptive Server 'CDGV240AIPV6':
Writing entry into directory services...
Directory services entry complete.
Building master device...
Segmentation Fault - core dumped
Task failed
Server 'CDGV240AIPV6' was not created.
```

This is a Sybase known issue. You should use Sybase Adaptive Server Enterprise Suite version 15.0.3 ESD 1 that supports Solaris 10 Update 6 or later. For details, refer to the Sybase Product Download Center regarding ESD 1.

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

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

Workaround:

To resolve this known issue

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

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

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

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

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

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

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

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

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

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

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

To resolve this issue

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

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

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

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

Veritas Volume Manager known issues

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

Dynamic LUN Expansion may fail on Solaris for EMC Clariion LUNs (2148851)

For EMC Clariion LUNs, if you perform Dynamic LUN Expansion operation using the `vxdisk resize` command while the I/O is in progress, the `vxdisk resize` command may fail with the following error:

```
VxVM vxdisk ERROR V-5-1-8643 Device device_name: resize failed:
New geometry makes partition unaligned
```

Work-around:

To resolve the issue, perform the following steps.

To recover from the error

1. Stop the I/O.
2. Reboot the system with the following command:

   ```
   # reboot -- -r
   ```

3. Retry the operation.

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 32Mb, then the public region data will be overridden. The work around is to specify explicitly the length of privoffset, puboffset, publen, privlen while initializing the disk.

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 for this issue.

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
```

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

If a path loses connectivity to the array, the path is marked with the NODE_SUSPECT flag. After the connectivity is restored, the restore daemon detects that the path is restored when the restore daemon probes the paths. The restore daemon clears the NODE_SUSPECT flag and makes the path is 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 `eeprom boot-device` command to set the boot device sequencing.

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

Changes in enclosure attributes are not persistent after an upgrade to VxVM 5.1 SP1 RP2 (2082414)

The Veritas Volume Manager (VxVM) 5.1 SP1 RP2 includes several array names that differ from the array names in previous releases. Therefore, if you upgrade from a previous release to VxVM 5.1 SP1 RP2, 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 5.1 SP1 RP2. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-15 shows the Hitachi arrays that have new array names.

Table 1-15 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. The persistence of the enclosure

name is achieved with the `/etc/vx/array.info` file, which stores the mapping between cabinet serial number and array name. 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 5.1 SP1 RP2. 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

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

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

Veritas File System known issues

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

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

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

Workaround: There is no workaround for this issue.

Shrinking a file system that is larger than 1 TB takes a long time (2097673)

Shrinking a file system shrink via either the `fsadm` command or `vxresize` command can take a long time to complete in some cases, such as if the shrink size is large and some large extent of a file is overlapping with the area to be shrunk.

Workaround: One possible workaround is to use the `vxtunefs` command and set`write_pref_io` and `write_nstream` to high values, such that `write_pref_io` multiplied by`write_nstream` is around 8 MB.

Storage Checkpoints can exceed the quota limit (2102201)

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

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

Workaround: There is no workaround for this issue.

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

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

Running fsppadmn enforce twice results in the "Too many open files" error (2118911)

If you run the `fsppadmn enforce` command twice, with the second instantiation running before the first instantiation completes, one of the instantiations displays the "Too many open files" error. This error only displays if the maximum open file limit on the system is too low.

Workaround: Set the maximum open file limit with the `ulimit` command to higher than the current limit.

cfsmount with the seconly option fails on Solaris 10 SPARC (2104499)

On Solaris 10 SPARC, the `cfsmount` command fails if you specify the `seconly` option.

Workaround: There is no workaround for this issue.

Installing the VRTSvxfs 5.1 SP1 RP2 package on non-global zones can fail (2406662)

Installing the VRTSvxfs 5.1 SP1 RP2 patch over 5.1 SP1 RP1 on non-global zones can fail with the following error messages:

```
package VRTSvxfs failed to install - interrupted:
pkgadd: ERROR: duplicate pathname zone_path/root/etc/fs/vxfs/qioadmin
pkgadd: ERROR: duplicate pathname zone_path/root/kernel/drv/vxportal.conf
pkgadd: ERROR: duplicate pathname zone_path/root/etc/vx/cdslimitstab
pkgadd: ERROR: duplicate pathname
  zone_path/root/opt/VRTSvxfs/etc/access_age_based.xml
pkgadd: ERROR: duplicate pathname
  zone_path/root/opt/VRTSvxfs/etc/access_age_based_2tier.xml
...
```

Workaround: The following procedure installs the VRTSvxfs5.1 SP1 RP2 over VRTSvxfs 5.1 SP1 RP1 on a non-global zone.

To install VRTSvxfs 5.1 SP1 RP2 on a non-global zone

1. Remove the `VRTSvxfs` 5.1 SP1 RP1 package.
2. Reinstall the `VRTSvxfs` 5.1 SP1 package.
3. Install the `VRTSvxfs` 5.1 SP1 RP2 package.

Unable to grow the File System on svm volume (2340820 (2213282))

This issue will be seen only on Solaris10 update8 > > (sol10u8) and above.

If the underlying volume is Solaris Volume Manger (SVM) and the size of the volume is > 1 TB, then operations like resize of file system through fsadm will fail with the following error:

```
UX:vxfs fsadm: ERROR: V-3-20058: read_vtoc failed with return value -7
```

Workaround:

A ONEOFF patch 5.1SP1RP2ONEOFF[Sol_sparc 2.10=> 146907-70, Sol_x86 2.10=> 146908-70] is available which contains the fix for this issue.

Possible error during an upgrade and when there is a local zone located on a VxFS file system(1675714)

During an upgrade and when there is local zone located on VxFS, you may receive an error message similar to the following:

```
Storage Foundation Uninstall did not complete successfully
VRTSvxvm package failed to uninstall on pilotv240-1
```

Workaround: You must reboot after the upgrade completes.

Possible write performance degradation with VxFS local mounts

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

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

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

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

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

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

To restore the benefits of the higher tunable value

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

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

NFS issues with VxFS checkpoint (1974020)

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

```
Stale NFS file handle
```

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

There is no workaround at this time.

Veritas Cluster Server known issues

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

imf_register reports error while doing registeration (2011536)

After a DB2 process dies, it can take some time for all other DB2 processes to die. So, even after IMF notifies of a DB2 process death, it's possible that the instantaneously invoked DB2 agent monitor reports DB2 resource state as online. In a corner case, it is possible that the subsequent online registration with IMF of the DB2 processes fails. This can happen when all DB2 processes have already

died when the actual event registration command is fired. Once all DB2 processes die, the monitor starts to report the DB2 resource offline.

Workaround: The error is safe to ignore.

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

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

Workaround:

1 Set `VCS_REMOTE_BROKER` to the remote AB:

```
# export VCS_REMOTE_BROKER=remote_broker
```

2 Set `VCS_DOMAIN` and `VCS_DOMAINTYPE`:

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

3 Run `halogin`:

```
# halogin ldap_user
```

Provide password when prompted.

4 Unset `VCS_DOMAIN` and `VCS_DOMAINTYPE`:

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

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

Parent service groups fail to restart after a child service group that has recovered from a fault restarts (2330038)

A child service group that has recovered from a fault will not be able to restart its faulted parent service group, if the parent service group's fault is detected before the child service group's fault.

Workaround:

Set the child service group's `OnlineClearParent` attribute to 1. When the child service group recovers from a fault and comes online, VCS clears the fault of the parent service group. This allows the VCS to bring the parent service group online.

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

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

VCS agent for Oracle: Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 (1985055)

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

Workaround: Set MonitorOption attribute for Oracle resource to 0.

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

The Application Agent uses the `system` command to execute the Start/Stop/Monitor/Clean Programs for root user. This executes Start/Stop/Monitor/Clean Programs in sh shell, due to which there is an error when root user has csh shell and `EnvFile` is written accordingly.

Workaround:

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

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

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

Workaround: Restarting the agent will resolve the issue.

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

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

Workaround: Restarting the agent will resolve the issue.

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

If multiple agent processes are running because an agent process is restarted multiple times by _had, then only one of the agent process survives and other

agent processes go offline on its own. Even though the agent process is running, _had does not recognize it and hence does not perform any resource operations.

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

pkgchk returns errors after installing VCS patches (2556541)

The `pkgchk` command returns errors for VCS patches as some of the files get modified during VCS configuration.

- Package: VRTSllt

```
# pkgchk -n VRTSllt
ERROR: /etc/default/llt
permissions <0744> expected <0644> actual
group name <sys> expected <root> actual
modtime <02/05/09 11:29:32 AM> expected <08/29/11 09:18:57 PM> actual
```

- Package: VRTSgab

```
# pkgchk -n VRTSgab
ERROR: /etc/default/gab
permissions <0744> expected <0644> actual
group name <sys> expected <root> actual
modtime <02/05/09 11:30:37 AM> expected <08/29/11 09:18:59 PM> actual
```

- Package: VRTSvxfen

```
# pkgchk -n VRTSvxfen
ERROR: /etc/default/vxfen
permissions <0544> expected <0644> actual
group name <sys> expected <root> actual
modtime <02/05/09 11:28:11 AM> expected <08/30/11 04:52:11 AM> actual
file cksum <21904> expected <21902> actual
```

- Package : VRTSvcs

```
# pkgchk -n VRTSvcs
ERROR: /etc/default/vcs
permissions <0744> expected <0644> actual
group name <sys> expected <root> actual
modtime <02/05/09 03:01:55 PM> expected <08/29/11 09:19:01 PM> actual
```

- Package: pkgchk -n VRTSvcsea

```
# pkgchk -n VRTSvcsea
ERROR: /etc/VRTSvcs/conf/config/Db2udbTypes.cf
File type <s> expected <f> actual
```

Workaround: There is no workaround for this issue. These errors are safe to ignore.

halting zone outside of VCS takes zone into down state (2557287)

If zone is halted outside of VCS control, zone goes into down state.

Workaround:

When zone goes to down state, an administrator intervention is required.

Unmount any mounted file systems inside of zone and bring the zone to steady state (configured/installed/running) using the **zoneadm –z *zonename* halt** command outside of VCS.

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

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

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

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

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

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

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

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

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

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

Zone agent does not report correct state of the local zone (2558234)

When the local zone goes into maintenance state, Zone agent reports zone state as offline when zone is not completely into offline state.

the zone from outside VCS using following command:

```
# zoneadm -z zone_name halt
```

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

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

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

the VRTSllt, VRTSgab and VRTSvxfen patches fails to install in non global zone (2562424)

The VRTSllt, VRTSgab and VRTSvxfen patches fails to install during the Live Upgrade from 5.1SP1RP1 to 5.1SP1RP2 on the system where non global zone is running.

Workaround: This message is safe to ignore.

warning for patch "143290-04" during the Live Upgrade from 5.1SP1RP1 to 5.1SP1RP2 (2562430)

If the local zone is not in running state, the following message displays for patch "143290-04" during the Live Upgrade from 5.1SP1RP1 to 5.1SP1RP2:

```
Executing postpatch script...

WARNING: Unable to find bmcmap
```

Workaround: This message is safe to ignore.

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

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

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

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

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

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

Veritas Volume Replicator known issues

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

vradmin syncvol command compatibility with IPv6 addresses (2075307)

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

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

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

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

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

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

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

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

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

Workaround: Destroy the instant snapshots manually using the `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
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Workaround:

To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

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

Workaround: There is no workaround for this issue.

vxassist relayout removes the DCM (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

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

3. Stop the affected RVG. Enter the following:

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

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

   ```
   # vxvol -g diskgroup dis vol
   ```

5. Resize the volumes. In this example, the volume is increased to 10 GB. Enter the following:

   ```
   # vxassist -g diskgroup growto vol 10G
   ```

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

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

7. Start the RVG. Enter the following:

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

8. Resume or start the applications.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

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

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

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

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

Workaround:

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

Database fails over during Flashsnap operations (1469310)

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

Workaround:

There is no workaround for this issue.

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

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

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

Workaround

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

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

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

   ```
   # vxdg join snapshop_disk_group_name
             primary_disk_group_name
   ```

2. Start all the volumes in primary disk group.

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

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

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

   Repeat this step for all the volumes.

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

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

Workaround:

There is no workaround for this issue.

Veritas Storage Foundation for Oracle RAC known issues

The following are new additional Veritas Storage Foundation for Oracle RAC known issues in this 5.1 SP1 RP2 release.

Incorrect ownership assigned to the parent directory of ORACLE_BASE causes Oracle Clusterware/Grid Infrastructure installations to fail

When you use the Veritas product installation program to install Oracle Clusterware/Grid Infrastructure, the ownership of the parent directory of ORACLE_BASE/GRID_BASE that is created by the installation program is incorrectly set to root. This causes the Oracle Universal Installer to display errors when it creates the `oraInventory` directory as the `oraInventory` directory must be created on behalf of the `oracle` user (Oracle RAC 10g Release 2/Oracle RAC 11g Release 1) or `grid` user (Oracle RAC 11g Release 2).

Workaround:

1. Log into each node in the cluster as the root user.
2. Perform the following operations:
   - If you have not yet installed Oracle Clusterware/Grid Infrastructure, create the directory and set the correct ownership as follows before you invoke the installation program:

     ```
     # mkdir -p oracle_base
     # chown user_name:oraInventory_group_name
                 oracle_base/..
     ```

     where:
     oracle_base is the name of the Oracle base directory.
     user_name is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).
     oraInventory_group_name is the name of the oraInventory group.
     Complete the other preparatory tasks before you invoke the installation program. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.
   - If you faced this issue during the installation of Oracle Clusterware/Grid Infrastructure, open another terminal session, and modify the ownership of the directory on all nodes in the cluster as follows:

     ```
     # chown user_name:oraInventory_group_name
                 oracle_base/..
     ```

where:
oracle_base is the name of the Oracle base directory.
user_name is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).
oraInventory_group_name is the name of the oraInventory group.
Return to the former session and proceed with the installation.

Failure to set the MTU (Maximum Transmission Unit) size in LLT over UDP environments causes issues with the PrivNIC/MultiPrivNIC agents (2557144)

If the MTU size field is not set explicitly when you configure the PrivNIC/MultiPrivNIC agents in an LLT over UDP environment, the agents may fail to plumb the private IP addresses during their operations or may configure incorrect MTU size on the LLT interfaces.

The agents use the `lltstat -l` command to retrieve MTU size information for LLT interfaces. In an LLT over UDP environment, the command retrieves 8192 as the MTU size. When the PrivNIC/MultiPrivNIC agents use this size information to plumb the IP addresses, the operation may fail causing the agents to fault. However, even if the plumbing operation succeeds, the incorrect MTU configuration may still cause issues in the cluster later.

Workaround:

To update the PrivNIC/MultiPrivNIC resource configuration in an LLT over UDP environment

1 Retrieve the MTU size of the network interfaces configured under PrivNIC/MultiPrivNIC agents:

For AIX: `# lsattr -El en1`

For HPUX: `# lanadmin -m 1`

For Linux: `# ifconfig eth1`

For Solaris: `# ifconfig ce0`

2 Set the MTU attribute for the PrivNIC/MultiPrivNIC resource:

```
# haconf -makerw
```

Run the following command for all the network interfaces configured under PrivNIC/MultiPrivNIC agents:

```
# hares -modify resource_name MTU -add interface_name mtu_size
```

Where:

resource_name is the name of the PrivNIC/MultiPrivNIC resource

interface_name is the name of the network interface for which the MTU size is set

mtu_size is the MTU size retrieved in step 1.

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

MultiPrivnic agent reports error messages in the VCS engine logs when zone goes down (2624304)

When a zone is down, the `MultiPrivNIC` agent fails to get active device information. The agent prints this message without specifying the message ID.

As a result, the following warning message displays:

```
VCS WARNING V-16-1-11328 Invalid message ID specified
```

Workaround: You may ignore this message.

Software limitations

This section covers the software limitations of this release.

Veritas Storage Foundation software limitations

There are no Veritas Storage Foundation software limitations in the 5.1 SP1 RP2 release.

Veritas Volume Manager software limitations

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

Converting a multi-pathed disk (2695660)

When converting a multi-pathed disk that is smaller than 1 TB from a VTOC label to an EFI label, you must issue the `format -e` command for each path. For example, if a node has two paths, `c1t2d0s2` and `c2tsd0s2`, you must run the `format -e` command on each of the two paths.

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, set the following DMP tunables:

Table 1-16

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

The change is persistent across reboots.

To change the tunable parameters

1 To change the tunable parameters, run the following commands:

```
# vxdmpadm settune dmp_restore_internal=60
# vxdmpadm settune dmp_path_age=120
```

2 To verify the new settings, run the following commands:

```
# vxdmpadm gettune dmp_restore_internal
# vxdmpadm gettune dmp_path_age
```

Dynamic LUN Expansion may fail on Solaris for EMC Clariion LUNs (2148851)

For EMC Clariion LUNs, if you perform Dynamic LUN Expansion operation using the `vxdisk resize` command while the I/O is in progress, the `vxdisk resize` command may fail with the following error:

```
VxVM vxdisk ERROR V-5-1-8643 Device device_name: resize failed:
New geometry makes partition unaligned
```

Work-around:

To resolve the issue, perform the following steps.

To recover from the error

1. Stop the I/O.
2. Reboot the system with the following command:

   ```
   # reboot -- -r
   ```

3. Retry the operation.

Veritas File System software limitations

There are no Veritas Storage Foundation software limitations in the 5.1 SP1 RP2 release.

Veritas Volume Replicator software limitations

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

Replication in a shared environment

Currently, replication support is limited to 4-node cluster applications.

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.

- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.
- A replication configuration in which an IPv6 address is specified for the `local_host` attribute of a primary RLINK and an IPv4 address is specified for the `remote_host` attribute of the same RLINK.
- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.
- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 5.1 SP1 RP2 and the prior major releases of Storage Foundation (5.0 MP3 and 5.1). Replication between versions is supported for disk group versions 140, 150, and 160 only. Both the Primary and Secondary hosts must be using a supported disk group version.

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 tools software limitations

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

Oracle Data Guard in an Oracle RAC environment

Database snapshots and Database Checkpoints are not supported in an environment where there are both Data Guard and Oracle RAC. But separately, either Data Guard or Oracle RAC supports Database snapshots and Database Checkpoints.

Upgrading if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 5.1 SP1 RP2: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 5.1 SP1 RP2.

Documentation errata

The following section provides documentation updates.

Veritas Cluster Server Administrator's Guide (2444653)

In the "VCS environment variables" section, the definition of the variable `VCS_GAB_RMTIMEOUT` should be "Timeout in milliseconds for HAD to register with GAB."

The value of `VCS_GAB_RMTIMEOUT` is specified in milliseconds. The minimum value is 200000 milliseconds or 200 seconds. If `VCS_GAB_RMTIMEOUT` is less than the minimum value then VCS overrides and sets it to 200000 milliseconds.

List of patches

This section lists the patches and packages for 5.1 SP1 RP2.

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

Table 1-17 Patches and packages for Solaris SPARC

Patch ID	5.1 Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
145451-03	VRTSdbac	SFRAC	5.3M		X
145450-03	VRTSdbac	SFRAC	6.7M	X	
142633-09	VRTSvxfs	FS, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	39 MB	X	
142634-09	VRTSvxfs	FS, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	50 MB		X
142629-12	VRTSvxvm	DMP, VM, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	199 MB	X	X
143687-03	VRTSob	SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	45 MB	X	X
144159-01	VRTSsfmh	FS, DMP, VM, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	32 MB	X	X
143287-07	VRTSvcs	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	116 MB	X	X
143288-12	VRTSvcsag	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	3.6 MB	X	X
145454-03	VRTSvxfen	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	2.9 MB	X	

Table 1-17 Patches and packages for Solaris SPARC *(continued)*

Patch ID	5.1 Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
145455-03	VRTSvxfen	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	2.3 MB		X
145471-02	VRTSamf	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.6 MB	X	
145473-02	VRTSamf	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	2.1 MB		X
143289-03	VRTScps	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	40 MB	X	X
143290-04	VRTSvcsea	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	15 MB	X	X
143281-03	VRTSllt	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.8 MB	X	
143282-03	VRTSllt	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.6 MB		X
143283-02	VRTSgab	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	2.1 MB	X	
143284-02	VRTSgab	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.6MB		X
143270-07	VRTSodm	SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.3 MB	X	
143271-07	VRTSodm	SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.2 MB		X
143680-03	VRTSglm	SFCFS, SFCFSHA, SFRAC, SVS	906 KB	X	
143681-03	VRTSglm	SFCFS, SFCFSHA, SFRAC, SVS	576 KB		X
143273-07	VRTScavf	SFCFS, SFCFSHA, SFRAC, SVS	816 KB	X	
143274-07	VRTScavf	SFCFS, SFCFSHA, SFRAC, SVS	818 KB		X

Table 1-17 Patches and packages for Solaris SPARC *(continued)*

Patch ID	5.1 Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
142631-05	VRTSdbed	SF, SFHA, SFCFS, SFCFSHA, SFRAC	11M	X	X

Table 1-18 Patches and packages for Solaris x64

Patch ID	5.1 Package Name	Products Affected	Patch Size	Solaris 10
145452-03	VRTSdbac	SFRAC	4.8 MB	X
143294-06	VRTSvcs	VCS, SFHA, SFCFSHA, SFRAC	204 MB	X
143295-12	VRTSvcsag	VCS, SFHA, SFCFSHA, SFRAC	5.8 MB	X
145456-03	VRTSvxfen	VCS, SFHA, SFCFSHA, SFRAC	2.3 MB	X
145472-02	VRTSamf	VCS, SFHA, SFCFSHA, SFRAC	1.8 MB	X
145458-01	VRTSsfmh	VCS, SFHA, SFCFSHA, SFRAC	37 MB	X
143296-03	VRTScps	VCS, SFHA, SFCFSHA, SFRAC	40 MB	X
143291-03	VRTSllt	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.5 MB	X
143292-02	VRTSgab	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.6 MB	X
143297-04	VRTSvcsea	VCS, SFHA, SFCFSHA, SFRAC	20 MB	X
142630-12	VRTSvxvm	DMP, VM, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	181 MB	X
142635-09	VRTSvxfs	FS, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	32 MB	X

Table 1-18 Patches and packages for Solaris x64 *(continued)*

Patch ID	5.1 Package Name	Products Affected	Patch Size	Solaris 10
143682-02	VRTSglm	SFCFS, SFCFSHA, SFRAC, SVS	611 KB	X
143275-06	VRTScavf	SFCFS, SFCFSHA, SFRAC, SVS	875 KB	X
143272-06	VRTSodm	SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.1 MB	X
142632-05	VRTSdbed	SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	50MB	X

Downloading the 5.1 SP1 RP2 archive

The patches that are included in the 5.1 SP1 RP2 release are available for download from the Symantec website. After downloading the 5.1 SP1 RP2 rolling patch, use gunzip and tar commands to uncompress and extract it.

For the 5.1 SP1 RP2 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

http://www.symantec.com/docs/TECH75362

Chapter 2

Installing the products for the first time

This chapter includes the following topics:

Installing the Veritas software using the script-based installer

This section describes how to install a 5.1 SP1 RP2 Veritas Storage Foundation and High Availability Solutions product for the first time on a host. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 5.1 SP1 *Installation Guide* and *Release Notes* for your product for more information.

See "Upgrading to 5.1 SP1 RP2" on page 102.

To install the Veritas software for the first time

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

4 Change the directory to /tmp/sfha51sp1:

```
# cd  /tmp/sfha51sp1
```

5 Run the installer to install SFHA 5.1SP1. See the Installation Guide for instructions on installing the 5.1 SP1 version of this product.

```
# ./installer -require complete_path_to_SP1_installer_patch
```

6 Download SFHA 5.1 SP1 RP2 from http://sort.symantec.com/patches.

7 Extract it to a directory called /tmp/sfha51sp1rp2.

8 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.

9 Change the directory to /tmp/sfha51sp1rp2:

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

10 Invoke the `installrp` script to install 5.1SP1 RP2:

```
# installrp -require complete_path_to_SP1RP2_installer_patch
```

11 If you did not configure the product after the 5.1 SP1 installation, the installer prompts you to configure the product during RP installation. If you do not want to configure the product now, answer **n** when prompted. To configure the product in the future, run the product installation script from the 5.1 SP1 installation media or from /opt/VRTS/install directory with the `-configure` option

Installing Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 5.1 SP1 RP2 using the Web-based installer. For detailed instructions on how to install 5.1 SP1 using the Web-based installer, follow the procedures in the 5.1 SP1 Installation Guide and Release Notes for your products.

See “Upgrading to 5.1 SP1 RP2” on page 102.

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

1. Start the Veritas XPortal Server process `xprtlwid`, on the installation server:

   ```
   # ./webinstaller start
   ```

   The webinstaller script displays a URL.

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

   ```
   Secure Connection Failed
   ```

   Obtain a security exception for your browser.
5. When prompted, enter `root` and root's password of the installation server.

Obtaining a security exception on Mozilla Firefox

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

To obtain a security exception

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

Installing 5.1 SP1 RP2 with the Veritas Web-based installer

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

To install Veritas product

1. The 5.1SP1 version of the Veritas product must be installed before upgrading to 5.1 SP1 RP2.

2. On the **Select a task and product** page, select **Install SP1 RP2** from the **Task** drop-down list, and click **Next**
3. Stop all applications accessing the filesystem. Unmount all mounted file systems before installation.
4. Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
5. After the validation completes successfully, click **Next** to install 5.1 SP1 RP2 patches on the selected system.
6. Select the checkbox to specify whether you want to send your installation information to Symantec.

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

   Click **Finish**.

Chapter 3

Installing the products using JumpStart

This chapter includes the following topics:

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 Veritas product 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 Veritas product.

To generate the finish script

1. Mount the 5.1 SP1 RP2 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 AT is generated at
/js_scripts/jumpstart_at.fin
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 SFCFS is generated at
/js_scripts/jumpstart_sfcfs.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 SFRAC 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_vm51102000.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_vm51102000.fin jumpstart_sf.fin
```

For SFHA:

```
encap_bootdisk_vm51102000.fin jumpstart_sfha.fin
```

For SFCFS:

```
encap_bootdisk_vm51102000.fin jumpstart_sfcfs.fin
```

For SF Oracle RAC:

```
encap_bootdisk_vm51102000.fin jumpstart_sfrac.fin
```

For VCS:

```
encap_bootdisk_vm51102000.fin jumpstart_vcs.fin
```

For DMP:

```
encap_bootdisk_vm51102000.fin jumpstart_dmp.fin
```

For SVS:

```
encap_bootdisk_vm51102000.fin jumpstart_svs.fin
```

For FS:

```
encap_bootdisk_vm51102000.fin jumpstart_fs.fin
```

For VM:

```
encap_bootdisk_vm51102000.fin jumpstart_vm.fin
```

For AT:

```
encap_bootdisk_vm51102000.fin jumpstart_at.fin
```

7 Copy the install and uninstall scripts of Veritas product that you are installing on to `BUILDSRC` shared location `/export/config` from 5.1 SP1 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 at/dmp/fs/sf/sfcfs/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
```

10 Once the installation is complete, refer to installation and configuration guide for the respective product from 5.1SP1 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.

4 Modify the rules file for JumpStart.

See the JumpStart documentation that came with your operating system for details.

5 Prepare installation resources.

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 `installer` 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 5.1 SP1 disk and 5.1 SP1 RP2 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.

To view the patches, packages and operating systems for your Veritas product use the `installrp -listpatches` command, type:

```
# ./installrp -listpatches
```

```
# cd BUILDSRC/pkgs/
# pkgask -r package_name.response -d /BUILDSRC/pkgs/packages_name.pkg
```

3 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
```

4 If you want to encapsulate the root disk automatically when perform the JumpStart installation, copy the scripts `encap_bootdisk_vm51001000.fin` created when you generated the finish script to *ENCAPSRC*.

5 Modify the rules file as required.

For example:

```
any - - profile_sf jumpstart_sf51.fin
```

For detailed instructions, see the *Oracle's JumpStart* documentation.

Adding language pack information to the finish file

For the language pack, copy the language packages from the language pack installation disc to the shared storage.

```
# cd /cdrom/cdrom0/pkgs
# cp -r * BUILDSRC/pkgs
```

Add lines for the language packages in the finish script. If the finish file resembles:

```
. . .
for PKG in VRTSperl VRTSvlic VRTSicsco . . .

do
...
done
```

Add the following lines for the language pack after the patch information for VCS. Copy the command syntax between the "do" and "done" lines and add that for the language pack lines as well. Note that the line that starts "for PKG" is on three lines in this guide, but should be on a single line in the file.

```
. . .
for PKG in VRTSmulic VRTSatJA VRTSjacav VRTSjacs VRTSjacse
VRTSjacsu VRTSjadba VRTSjafs VRTSjavm VRTSjadbe VRTSjaodm
VRTSatZH VRTSzhvm

do
.
.
.
done
```

Chapter 4

Upgrading to 5.1 SP1 RP2

This chapter includes the following topics:

Prerequisites for upgrading to 5.1 SP1 RP2

The following list describes prerequisites for upgrading to the 5.1 SP1 RP2 release:

- For any product in the Veritas Storage Foundation stack, you must have the 5.1 SP1 (or later) installed before you can upgrade that product to the 5.1 SP1 RP2 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 5.1 SP1 RP2 " on page 101.

Downloading required software to upgrade to 5.1 SP1 RP2

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

To download required software to upgrade to 5.1 SP1 RP2

1. Download SFHA 5.1 SP1 RP2 from http://sort.symantec.com/patches.
2. Extract it to a directory, say /tmp/sfha51sp1rp2.
3. Check http://sort.symantec.com/patches to see if there are patches available for the 5.1 SP1 RP2 installer. Download applicable P-patches and extract them to the /tmp directory.
4. When you run the `installrp` script, use the `-require` option and specify the location where you downloaded the patches.

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 5.1 SP1 to 5.1 SP1 RP2
- 5.1 SP1 RP1 to 5.1 SP1 RP2
- 5.1 SP1 P2 to 5.1 SP1 RP2

Upgrading to 5.1 SP1 RP2

This section describes how to upgrade from 5.1 SP1 (or later) to 5.1 SP1 RP2 on a cluster or a standalone system.

- Performing a full upgrade to 5.1 SP1 RP2 on a cluster
  Use the procedures to perform a full upgrade to 5.1 SP1 RP2 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System (SFCFS), Veritas Storage Foundation for Oracle RAC (SFRAC), or Symantec VirtualStore (SVS) installed and configured.
- Upgrading to 5.1 SP1 RP2 on a standalone system
  Use the procedure to upgrade to 5.1 SP1 RP2 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 89.

Performing a full upgrade to 5.1 SP1 RP2 on a cluster

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

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.1 SP1 RP2:

- Performing a full upgrade to 5.1 SP1 RP2 on a Veritas Cluster Server
- Performing a full upgrade to 5.1 SP1 RP2 on an SFHA cluster
- Performing a full upgrade to 5.1 SP1 RP2 on an SFCFS cluster
- Performing a full upgrade to 5.1 SP1 RP2 on an SF Oracle RAC cluster
  See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 101.

Performing a full upgrade to 5.1 SP1 RP2 on a Veritas Cluster Server

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

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

To upgrade VCS

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 101.

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 5.1 SP1 RP2 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 5.1 SP1 RP2 on an SFHA cluster

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

To perform a full upgrade to 5.1 SP1 RP2 on an SFHA cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 101.

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 5.1 SP1 RP2 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 5.1 SP1 RP2 on an SFCFS cluster

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

To perform a full upgrade to 5.1 SP1 RP2 on an SFCFS cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 101.

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:

```
# hagrp -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 `vxrvg stop` command to stop each RVG individually:

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

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

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

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

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:

- Solaris 9:

```
# /etc/init.d/odm stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- Solaris 10:

```
# 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/vxsvcctrl status
```

If the VEA service is running, stop it:

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

17 From the directory that contains the extracted and untarred 5.1 SP1 RP2 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:

```
# hagrp -unfreeze groupname -persistent
```

21 Make the configuration read-only:

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

22 Bring the CVM service group online on each node:

```
# hagrp -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:

```
# vxrvg -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 5.1 SP1 RP2 on an SF Oracle RAC cluster

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

To upgrade to 5.1 SP1 RP2 on an SF Oracle RAC cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 101.

2 Log in as superuser.

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

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

```
# haconf -makerw
```

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

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

6 Make the configuration read-only:

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

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

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

8 Stop VCS.

```
# hastop -all
```

9 If required, apply the OS kernel patches.

See *Oracle's* documentation for the procedures.

10 From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. If ssh key authentication is configured then enter:

```
# ./installrp  node1 node2
```

If ssh is not configured then enter:

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

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

11 After the entire cluster is upgraded, follow the installer instructions to proceed further.

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

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

```
# haconf -makerw
```

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

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

15 Make the configuration read-only:

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

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

```
# hagrp -online service_group -sys nodename
```

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

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

18 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

19 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
```

20 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 to 5.1 SP1 RP2 on a standalone system

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

To upgrade to 5.1 SP1 RP2 on a standalone system

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 101.

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 `vxrvg stop` command to stop each RVG individually:

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

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

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

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

8. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
9. Stop all VxVM volumes by entering the following command for each disk group:

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

   Verify that no volumes remain open:

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

10. Check if the VEA service is running:

    ```
    # /opt/VRTS/bin/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 6, restart each RVG:

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

15 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
# mount /checkpoint_name
```

16 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/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 5.1 SP1 to 5.1 SP1 RP2 using Live Upgrade.

Supported live upgrade paths:

- Upgrading Veritas Products without Solaris OS upgrade:
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 9 Update x 5.1 SP1 RP2
  - Upgrading Solaris 10 Update x 5.1 SP1 to Solaris 10 Update x 5.1 SP1 RP2
- Upgrading Veritas Products with Solaris OS upgrade
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 9 Update y 5.1 SP1 RP2
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 10 Update y 5.1 SP1 RP2
  - Upgrading Solaris 10 Update x 5.1 SP1 to Solaris 10 Update y 5.1 SP1 RP2

Prerequisites to upgrade to 5.1 SP1 RP2 using Live Upgrade:

- The node should have an alternate boot disk that is identical to the primary boot disk.
- Installation disc for 5.1 SP1 and 5.1 SP1 RP2 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 5.1 SP1 to 5.1 SP1 RP2 without OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, or SF for Oracle RAC from 5.1 SP1 to 5.1 SP1 RP2 using Live Upgrade where OS upgrade is not involved..

To upgrade your Veritas product using Live Upgrade

1. Ensure that 5.1 SP1 is installed and configured on PBE.

   See your Veritas product 5.1 SP1 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 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

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

7 Verify that the alternate boot environment is active.

```
# lustatus
```

8 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 5.1 SP1 to 5.1 SP1 RP2 with OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, or SF for Oracle RAC from 5.1 SP1 to 5.1 SP1 RP2 using Live Upgrade where OS upgrade is involved..

To upgrade your Veritas product using Live Upgrade

1 Ensure that 5.1 SP1 is installed and configured on PBE.

See your Veritas product 5.1 SP1 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 If you are upgrading from Solaris 9 Update x to Solaris 10 Update y, uninstall Veritas products on ABE, else go to Step 6.

```
# ./uninstallprod -rootpath /altroot_path
```

5 Install Veritas products again on ABE.

```
# ./installprod -rootpath /altroot_path
```

6 Run the `installrp` command to upgrade your Veritas product:

```
# ./installrp -rootpath  /altroot_path
```

7 In case of SFRAC, refer to the "Completing the Live upgrade" section in SFRAC 5.1 SP1 Installation and Configuration guide. For other products, go to the next step.

8 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
  ```

9 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

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

10 Verify that the alternate boot environment is active.

```
# lustatus
```

11 In a cluster environment, make sure that all the GAB ports are up. Note that different ports appear for different products.

```
# gabconfig -a
```

12 In case of SFRAC, refer to the “Performing post-upgrade Tasks” section to relink Oracle RAC libraries with SF Oracle RAC from 5.1SP1 Installation and Configuration guide.

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 non-kernel packages for VCS, SFHA, SFCFS, and SFCFSHA : phase 2

Note: The SF for Oracle RAC rolling upgrade from pre-5.1SP1RP2 to 5.1SP1RP2 and from 5.1SP1RP2 to future release is not supported. You could perform a full upgrade on an SF for Oracle RAC cluster. See “Performing a full upgrade to 5.1 SP1 RP2 on an SF Oracle RAC cluster” on page 108.

About rolling upgrades

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

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

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System
- Storage Foundation Cluster File System and High Availability

You can perform a rolling upgrade from 5.1 SP1 or later.

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.
- Make sure you have downloaded the latest software required for the upgrade. See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 101.

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 on kernel packages for VCS, SFHA, SFCFS, and SFCFSHA: phase 1

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

To perform the rolling upgrade on kernel packages: phase 1

1. Stop all applications that access volumes.
2. Unmount any file systems on the nodes that you plan to upgrade.

   You only need to unmount locally mounted file systems. The installer unmounts file systems that VxFS has mounted.
3. On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.

   ```
   # ./installrp -upgrade_kernelpkgs nodeA
   ```
4. Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
5. The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
6. The installer loads new kernel modules and starts all the relevant processes and brings all the service groups online.

7 If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step 8.

8 Before you proceed to phase 2, complete step 1 to step 6 on the second subcluster.

Performing a rolling upgrade on non-kernel packages for VCS, SFHA, SFCFS, 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 `-upgrade_nonkernelpkgs` option. Specify all the nodes in the cluster:

```
./installrp -upgrade_nonkernelpkgs nodeA nodeB nodeC nodeD
```

2 The installer checks system communications, fileset versions, product versions, and completes prechecks. It verifies completion of phase 1.

3 Installer will start HA daemon (had) on all nodes, HA will be available once HA daemon is up.

4 Verify the cluster's status:

```
# hastatus -sum
```

5 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

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

Verifying software versions

To verify the version of the software, enter the following command:

```
# pkginfo -l  pkgname
```

Chapter 5

Rolling back and removing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2

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

Roll back of version 5.1 SP1 RP2 to the 5.1 SP1 release is supported for the following products:

- Storage Foundation and High Availability (SFHA)
- Veritas Storage Foundation Cluster File System (SFCFS)
- 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.

```
# hagrp -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 or 11gR1:

```
# /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:

  ```
  # hastop -all
  ```

6 Stopping the applications that use VxVM or VxFS that are not under VCS control

- Using native application commands, stop the applications that use VxVM or VxFS.
- Verify that no processes use the VxFS mount point:

  ```
  # fuser -c mount_point
  ```

7 Unmounting VxFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

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

- Unmount each file system that is not controlled by VCS on each node:

  ```
  # umount mount_point
  ```

8 Run the `uninstallrp` command, type:

  ```
  # ./uninstallrp node A node Bnode C...
  ```

9 If you performed a roll back on a system that has an encapsualted 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 encapsualted 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.

Rolling back manually

Use one of the following procedures to roll back to 5.1 SP1 manually.

- Rolling back Storage Foundation or Storage Foundation and High Availability manually
- Rolling back Storage Foundation Cluster File System manually
- Rolling back Storage Foundation for Oracle RAC 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.1

Rolling back Storage Foundation or Storage Foundation and High Availability manually

Use the following procedure to roll back to 5.1 SP1 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 vxrvg stop command to stop each RVG individually:

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

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

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

  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:
   - For Solaris 9:

     ```
     # /etc/init.d/vcs stop
     ```

   - For Solaris 10:

     ```
     # 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:
    - For Solaris 9:

      ```
      # /etc/init.d/odm stop
      ```

- For Solaris 10:

```
# svcadm disable -t odm
# modinfo | grep odm
# modunload -i odm_mod_id
```

13 Unload the cluster fencing (`vxfen`) module:

- For Solaris 9:

```
# /etc/init.d/vxfen stop
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

- For Solaris 10:

```
# svcadm disable -t vxfen
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

14 Stop GAB and LLT in the following order:

For Solaris 9:

```
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SF 5.1 SP1 RP2 patches.

- Get the list of 5.1 SP1 RP2 patches, type:

```
# ./installrp -listpatches
```

- Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```

Rolling back Storage Foundation Cluster File System manually

Use the following procedure to roll back to 5.1 SP1 manually.

To roll back SFCFS or SFCFS HA 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 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:

- For Solaris 9:

```
# /etc/init.d/vcs stop
```

- For Solaris 10:

```
# 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:

- For Solaris 9:

```
# /etc/init.d/odm stop
```

- For Solaris 10:

```
# svcadm disable -t odm
# modinfo | grep odm
# modunload -i odm_mod_id
```

13 Unload the cluster fencing (`vxfen`) module:

- For Solaris 9:

```
# /etc/init.d/vxfen stop
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

- For Solaris 10:

```
# svcadm disable -t vxfen
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

14 Stop GAB and LLT in the following order:

For Solaris 9:

```
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SFCFS 5.1 SP1 RP2 patches.

- Get the list of 5.1 SP1 RP2 patches, type:

```
# ./installrp -listpatches
```

- Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```

Rolling back Storage Foundation for Oracle RAC manually

Use the following procedure to roll back to 5.1 SP1 manually.

To roll back SF for Oracle RAC manually

1 On each node, take the Oracle resources in the VCS configuration file (main.cf) offline.

```
# hagrp -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 or 11gR1:

  ```
  # /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:

   ```
   # hastop -all
   ```

6. Stopping the applications that use VxVM or VxFS that are not under VCS control
   - Using native application commands, stop the applications that use VxVM or VxFS.
   - Verify that no processes use the VxFS mount point:

     ```
     # fuser -c mount_point
     ```

7. Unmounting VxFS file systems that are not under VCS control.
   - Determine the file systems that need to be unmounted by checking the output of mount command.

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

- Unmount each file system that is not controlled by VCS on each node:

```
# umount mount_point
```

8 To stop the process, type:

```
# ./installsfrac -stop <node1> <node2> ... <nodeN>
```

9 Remove the SF for Oracle RAC 5.1 SP1 RP2 patches.

- Get the list of 5.1 SP1 RP2 patches, type:

```
# ./installrp -listpatches
```

- Remove each patch from the patch list. For example:

```
# patchrm 143287-07
```

10 Verify that the patches have been remove on all the nodes.

11 Reboot the nodes point.

```
# /usr/sbin/shutdown -g0 -y -i6
```

Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 5.1 SP1 RP2 to VCS 5.1 SP1 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:

```
# hagrp -state
```

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

```
# hagrp -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:

```
# hagrp -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 5.1SP1 RP2)
```

- 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 5.1SP1 RP2)
```

- 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 5.1SP1 RP2)
```

- Unload LLT using the module number:

```
# modunload -i 147
```

17 Remove the VCS 5.1 SP1 RP2 patches. On each node, perform the following steps:

- Get the list of 5.1 SP1 RP2 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
# hagrp -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:

```
# hagrp -online ClusterService -sys system
```

where *system* is the node name.

Rolling back Symantec VirtualStore manually

Use the following procedure to roll back to 5.1 SP1 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 vxrvg stop command to stop each RVG individually:

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

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

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

  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:

- For Solaris 9:

  ```
  # /etc/init.d/vcs stop
  ```

- For Solaris 10:

  ```
  # 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:

- For Solaris 9:

  ```
  # /etc/init.d/odm stop
  ```

- For Solaris 10:

```
# svcadm disable -t odm
# modinfo | grep odm
# modunload -i odm_mod_id
```

13 Unload the cluster fencing (`vxfen`) module:

- For Solaris 9:

```
# /etc/init.d/vxfen stop
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

- For Solaris 10:

```
# svcadm disable -t vxfen
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

14 Stop GAB and LLT in the following order:

For Solaris 9:

```
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SVS 5.1 SP1 RP2 patches.

- Get the list of 5.1 SP1 RP2 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 5.1 SP1 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 vxrvg stop command to stop each RVG individually:

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

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

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

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:

- For Solaris 9:

```
# /etc/init.d/vcs stop
```

- For Solaris 10:

```
# 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:

- For Solaris 9:

```
# /etc/init.d/odm stop
```

- For Solaris 10:

```
# svcadm disable -t odm
# modinfo | grep odm
# modunload -i odm_mod_id
```

13 Unload the cluster fencing (`vxfen`) module:

- For Solaris 9:

```
# /etc/init.d/vxfen stop
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

- For Solaris 10:

```
# svcadm disable -t vxfen
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

14 Stop GAB and LLT in the following order:

For Solaris 9:

```
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the DMP 5.1 SP1 RP2 patches.

- Get the list of 5.1 SP1 RP2 patches, type:

  ```
  # ./installrp -listpatches
  ```

- Remove each patch from the patch list. For example:

  ```
  # patchrm 143287-07
  ```