Veritas Storage Foundation and High Availability Solutions Release Notes

Linux

5.1 Rolling Patch 2

Storage Foundation and High Availability Solutions Release Notes 5.1 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 RP2

Document version: 5.1RP2.0

Legal Notice

Symantec, the Symantec Logo, Veritas, Veritas Storage Foundation 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

Documentation feedback

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:

sfha_docs@symantec.com

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

Contents

Chapter 1

About Veritas Storage Foundation and High Availability Solutions 5.1 RP2

This chapter includes the following topics:

Introduction

This document provides information about the Storage Foundation and High Availability Solutions 5.1 Rolling Patch 2.

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

http://entsupport.symantec.com/docs/335001

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

http://entsupport.symantec.com/docs/330441

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

About the installrp script

To upgrade from Veritas Storage Foundation and High Availability Solutions version 5.1 or later, the recommended method is to use the `installrp` script, which allows you to upgrade all the patches associated with the packages installed. The `installrp` script will automatically restart all the processes after the upgrade. But if it failed to do so, you will be asked to reboot the system.

installrp script options

Table 1-1 shows 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 ]`	The -precheck option is used to confirm that systems meet the products install requirements before installing.
`[ -logpath log_path ]`	The `-logpath` option is used to select a directory other than `/opt/VRTS/install/logs` as the location where installrp log files, summary file, and response file are saved.

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

Command Line Option	Function
`[ -responsefile response_file ]`	The `-responsefile` option is used 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.
`[ -makeresponsefile ]`	The `-makeresponsefile` option generates a response file without doing an actual installation. The text displaying Install, uninstall, start, and stop actions are simulations. These actions are not performed on the system.
`[ -tmppath tmp_path ]`	The `-tmppath` option is used 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 ]`	The `-hostfile` option specifies the location of a file containing the system names for installer.
`[ -keyfile ssh_key_file ]`	The `-keyfile` option specifies a key file for SSH. When this option is used, `-i <ssh_key_file>` is passed to every SSH invocation.
`[ -patchpath patch_path ]`	The `-patchpath` option is used to define the complete path of a directory available to all install systems (usually NFS mounted) that contains all patches to be installed by `installrp`.

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

Command Line Option	Function
`[-kickstart ]`	The `-kickstart` option is used to generate kickstart scripts which can be used by Redhat Linux Kickstart for automated installation of all rpms for every product, an available location to store the kickstart scripts should be specified as a complete path. The `-kickstart` option is supported on Redhat Linux only.
`[ -rsh | -redirect | -listpatches | -pkginfo | -serial | -upgrade_kernelpkgs | -upgrade_nonkernelpkgs ]`	The `-rsh` option is used when rsh and rcp are to be forced for communication though ssh and scp is also setup between the systems. The `-redirect` option is used to display progress details without showing the progress bar. The `-listpatches` option is used to display product patches in the correct installation order. The `-pkginfo` option is used to display the correct installation order of packages and patches. The `-serial` option is used to perform installation, uninstallation, start, and stop operations, typically performed simultaneously on all systems, in a serial fashion. The `-upgrade_kernelpkgs` option is used for the rolling upgrade's upgrade of kernel packages to the latest version The `-upgrade_nonkernelpkgs` option is used 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.

Changes introduced in 5.1 RP2

The following sections describe changes in product behavior in this release.

RemoteGroup is supported in local parallel service groups

RemoteGroup resource can be configured inside local parallel Service Groups. It comes online on all the cluster nodes if the remote service group is online. When configured inside the parallel service group:

- The RemoteGroup resource will continue to monitor the remote Service Group even when the resource is offline.
- The RemoteGroup resource will not send the remote Service Group offline if the RemoteGroup resource is online anywhere in the cluster.
- After agent restart, the RemoteGroup resource will not go offline if the RemoteGroup resource is online on any other cluster node.
- The RemoteGroup resource will send the remote Service Group offline if it is the only instance of RemoteGroup resource online in the local cluster.
- RemoteGroup resource online will not perform any operation if the same resource instance is online on other cluster node.

The DNS agent supports more flexible naming rules

The DNS agent has been enforcing RFC 1035, which invalidates host names with the underscore (“_”) character. Symantec has modified the DNS agent to allow the use of underscore character in host names.

Note: You must make sure that the DNS server supports the underscore character before you configure any DNS resource records to have underscores in their host names.

Enhanced IP and IPMultiNIC agents

Symantec has modified the IP/IPMultiNIC agent to select an interface alias based on the first available virtual interface index for the device.

New cluster attribute - AutoAddSystemToCSG

The new AutoAddSystemToCSG attribute has been added at the cluster level to determine if the SystemList of the ClusterService group should be populated automatically. By default, its value is 1 (true) to retain current behavior. If this attribute is disabled, when a new node joins or is added, it will not be added to the SystemList of the ClusterService group automatically. The user will have to add it manually.

SourceIP for Notifier

Notifier provides option to bind to specific source IP.

Changes introduced in 5.1 RP1

The following sections describe changes in product behavior in this release.

CVM master node needs to assume the logowner role for VCS managed VVR resources

If you use VCS to manage VVR resources in an SFCFS or an SFCFSRAC environment, Symantec strongly recommends that you perform the steps in the section “Using the preonline_vvr trigger for RVGLogowner resources.” These steps ensure that the CVM master node always assumes the logowner role. Not doing this can result in unexpected issues. These issues are due to a CVM slave node that assumes the logowner role.

See “Using the preonline_vvr trigger for RVGLogowner resources” on page 45.

System requirements

This section describes the system requirements for this release

Supported Linux operating systems

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

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

http://entsupport.symantec.com/docs/335001

The Veritas 5.1 RP2 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)

- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel) or SP3 (2.6.16.60-0.54.5) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)

  Note: SFCFSRAC is not supported on SLES 11.

- Oracle Enterprise Linux (OEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)

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

Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

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

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

http://entsupport.symantec.com/docs/335001

Xen platform for Linux

The Veritas 5.1 RP2 release is supported on the Xen platform for Linux with some restrictions.

Note: Veritas Cluster Server is not supported on the Xen platform for Linux.

VMware Environment

For information about the use of this product in a VMware Environment, refer to http://entsupport.symantec.com/docs/289033

Note: This technote is being updated to include information specific to 5.1 versions. Please check this technote for the latest information.

Database requirements

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

http://entsupport.symantec.com/docs/331625

Note: SF and SFCFS support running Oracle, DB2, and Sybase on VxFS and VxVM.

SF and SFCFS do not support running SFDB tools with DB2 and Sybase.

Recommended memory and swap space

Symantec recommends the following memory and swap space sizes:

- On the system where you run the installation commands, 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 commands, 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
  - More than eight nodes, 1 GB of free swap space

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 Cluster File System for Oracle RAC (SFCFSRAC)
- Veritas Volume Manager (VM)
- Veritas File System (FS)
- Veritas Cluster Server (VCS)

Fixed issues

This section describes the issues fixed in this release.

- Veritas Storage Foundation fixed issues
- Veritas File System fixed issues
- Veritas Volume Manager fixed issues
- Veritas Storage Foundation Cluster File System fixed issues
- Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues
- Veritas Cluster Server fixed issues
- Veritas Storage Foundation Manager fixed issues
- Veritas Enterprise Administrator fixed issues

Veritas Storage Foundation fixed issues

This section lists the fixed issues for Veritas Storage Foundation.

See also:

- Veritas Volume Manager fixed issues
- Veritas File System fixed issues

Veritas Storage Foundation: Issues fixed in 5.1 RP2

Table 1-2 Veritas Storage Foundation fixed issues in 5.1 RP2

Fixed issues	Description
2088355	dbed_ckptrollback fails for –F datafile option for 11gr2
2080633	Fixed the issue with vxdbd dumping core during system reboot.
2080565	vxdbd fails to start if the ipv6 kernel module is not loaded
1976928	dbed_clonedb of offline checkpoint fails with ORA-00600

Veritas Storage Foundation: Issues fixed in 5.1 RP1

Table 1-3 Veritas Storage Foundation fixed issues in 5.1 RP1

Fixed issues	Description
1974086	reverse_resync_begin fails after successfully unmount of clone database on same node when primary and secondary host names do not exactly match.
1940409, 471276	Enhanced support for cached ODM
1901367, 1902312	dbed_vmclonedb failed to umount on secondary server after a successful VM cloning in RAC when the primary SID string is part of the snapplan name.
1896097	5.1 GA Patch:dbed_vmclonedb -o recoverdb for offhost get failed
1873738, 1874926	dbed_vmchecksnap fails on standby database, if not all redologs from primary db are present.
1810711, 1874931	dbed_vmsnap reverse_resync_begin failed with server errors.

Veritas File System fixed issues

This section lists fixed issues for Veritas File System.

Veritas File System: Issues fixed in 5.1 RP2

Table 1-4 Veritas File System fixed issues in 5.1 RP2

Fixed issues	Description
1995399	Fixed a panic due to null i_fsext pointer de-reference in vx_inode structure
2016373	Fixed a warning message V-3-26685 during freeze operation without nested mount points
2036841	Fixed a panic in vx_set_tunefs
2081441	Fixed an issue in vxedquota regarding setting quota more than 1TB
2018481	Fixed an issue in fsppadm(1M) when volume did not have placement tags
2066175	Fixed panic in vx_inode_mem_deinit

Table 1-4 Veritas File System fixed issues in 5.1 RP2 *(continued)*

Fixed issues	Description
1939772	Fixed an issue in vxrepquota(1m) where username and groupname were truncated to 8 characters
2025155	Fixed an issue in fsck(1m) which was trying to free memory which was not allocated.
2043634	Fixed an issue in quotas API
1933844	Fixed a panic due to race condition in vx_logbuf_clean()
1960836	Fixed an issue in Thin Reclaim Operation
2026570	Fixed a hang issue in vx_dopreamble () due to ENOSPC error.
2026622	Fixed a runqueue contention issue for vx_worklists_thr threads
2030889	Fixed a hang issue during fsppadm(1m) enforce operation with FCL
2036214	Fixed a core dump issue in ncheck(1m) in function printname().
2076284	Optimized some VxMS api for contiguous extents.
2085395	Fixed a hang issue in vxfsckd.
2072162	Fixed the issue of writing zero length to null buffer
2059621	Fixed a panic due to null pointer de-reference in vx_unlockmap()
2016345	Fixed an error EINVAL issue with O_CREATE while creating more than 1 million files.
1976402	Fixed the issue in fsck replay where it used to double fault for 2TB luns.
1954692	Fixed a panic due to NULL pointer de-reference in vx_free()
2026599	Fixed a corruption issue when Direct IO write was used with buffered read.
2072161	Fixed a hang issue in vx_traninit()
2030773	Fixed issue with fsppadm(1m) where it used to generate core when an incorrectly formatted XML file was used.
2026524	Fixed a panic in vx_mkimtran()
2080413	Fixed an issue with storage quotas

Table 1-4 Veritas File System fixed issues in 5.1 RP2 *(continued)*

Fixed issues	Description
2084071	Fixed an issue in fcladm(1m) where it used to generate core when no savefile was specified
2026637	Support for kernel extended attributes
2072165	Fixed an active level leak issue while fsadm resize operation.
2059008	Fixed an issue with quotas where hard limit was not enforced in CFS environment
1959374	Fixed a resize issue when IFDEV is corrupt
2098385	Fixed a performance issue related to ‘nodatainlog’ mount option.
2112358	Fixed an issue with file-system I/O statistics.

Veritas File System: Issues fixed in 5.1 RP1

Table 1-5 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number)

Fixed issues	Description
1897458, 1805046	Fixed issue in alert generation from vxfs when file system usage threshold is set.
1933635, 1914625	Fixed issues in fs pattern assignment policy of the file system.
1933975, 1844833	Fixed VX_EBMAPMAX error during filesystem shrinking using fsadm..
1934085, 1871935	We now update ilist on secondary even if error received from primary for a VX_GETIAS_MSG is EIO.
1934095, 1838468	Fixed a race in qiostat update which was resulting in data page fault.
1934096, 1746491	Fix to avoid core dump while running fsvmap by initializing a local pointer.
1934098, 1860701	Moved drop of active level and reaquire to top of loop to stop resize from being locked out during clone removal.

Table 1-5 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) *(continued)*

Fixed issues	Description
1934107, 1891400	Fixed incorrect ACL inheritance issue by changing the way it cached permission data.
1947356, 1883938	Added utility mkdstfs to create DST policies.
1934094, 1846461	Fixed an issue with vxfsstat(1M) counters.

Veritas Volume Manager fixed issues

This section lists fixed issues for Veritas Volume Manger.

Veritas Volume Manager: Issues fixed in 5.1 RP2

Table 1-6 Veritas Volume Manager 5.1 RP2 fixed issues

Fixed issues	Description
1973367	VxVM Support for Virtio Virtual Disks in KVM virtual Machines
1938907	RHEL5 U3: WWN information is not displayed due to incorrect device information returned by HBA APIs
2069022	Booting between Linux kernels results in stale APM key links.
2067568	EqualLogic iSCSI - Disabling array switch port leads to disk failure and disabling of path.
2015570	File System read failure seen on space optimized snapshot after cache recovery
1665094	Snapshot refresh causing the snapshot plex to be detached.
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout
1946936	CVM: IO hangs during master takeover waiting for a cache object to quiesce
1946939	CVM: Panic during master takeover, when there are cache object I/Os being started on the new master

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

Fixed issues	Description
2053975	Snapback operation panicked the system
1983768	IO hung on linked volumes while carrying out third mirror breakoff operation.
1513385	VVR:Primary panic during autosync or dcm replay.
2052459	CFS mount failed on slave node due to registration failure on one of the paths
1936611	vxconfigd core dump while splitting a diskgroup
1992872	Vxresize fails after DLE.
1960341	Toggling of naming scheme is not properly updating the daname in the vxvm records.
1933528	During Dynamic reconfiguration vxvm disk ends up in error state after replacing physical LUN.
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1
1933375	Tunable value of 'voliomem_chunk_size' is not aligned to page-size granularity
2040150	Existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration
1441406	'vxdisk -x list' displays wrong DGID
1956777	CVR: Cluster reconfiguration in primary site caused master node to panic due to queue corruption
1942985	Improve locking mechanism while updating mediatype on vxvm objects
1911546	Vxrecover hung with layered volumes
2012016	Slave node panics while vxrecovery is in progress on master
2078111	When the IOs are large and need to be split, DRL for linked volumes cause I/Os to hang
1880279	Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment.
1952177	Machine panics after creating RVG

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

Fixed issues	Description
1929083	Vxattachd fails to reattach site in absence of vxnotify events
1097258	vxconfigd hung when an array is disconnected
1972755	TP/ETERNUS:No reclaim seen with Stripe-Mirror volume.
2061066	vxisforeign command fails on internal cciss devices
2021737	vxdisk list shows HDB TrueCopy S-VOL read only devices in error state.
2065669	After upgrading to 5.1, reinitalizing the disk makes public region size smaller than the actual size.
1974393	Avoiding cluster hang when the transaction client timed out
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.
2031462	Node idle events are generated every second for idle paths controlled by Third Party drivers.
1982715	vxclustadm dumping core while memory re-allocation.
1998447	Vxconfigd dumped core due to incorrect handling of signal
1999004	I/Os hang in VxVM on linked-based snapshot
1899943	CPS based fencing disks used along with CPS servers does not have coordinator flag set
1923906	CVM: Master should not initiate detaches while leaving the cluster due to complete storage failure
2006454	AxRT5.1P1: vxsnap prepare is displaying vague error message
1989662	/opt/VRTSsfmh/bin/vxlist causes panic.
2059046	FMR:TP: snap vol data gets corrupted if vxdisk reclaim is run while sync is in progress
2011316	VVR: After rebooting 4 nodes and try recovering RVG will panic all the slave nodes.
1485075	DMP sending I/O on an unopened path causing I/O to hang
1874034	Race between modunload and an incoming IO leading to panic

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

Fixed issues	Description
2055609	Allocation specifications not being propagated for DCO during a grow operation
2029480	Diskgroup join failure renders source diskgroup into inconsistent state
2029735	System panic while trying to create snapshot
1897007	vxesd coredumps on startup when the system is connected to a switch which has more than 64 ports
1831969	VxVM: ddl log files are created with world write permission
2010426	Tag setting and removal do not handle wrong enclosure name
2036929	renaming a volume with link object attached causes inconsistencies in the disk group configuration
1920894	vxcheckhbaapi can loop forever
1920761	I/O hang observed after connecting the storage back to master node incase of local detach policy
2034104	Unable to initialize a disk using vxdiskadm
1946941	vxsnap print shows incorrect year
1829337	Array firmware reversal led to disk failure and offlined all VCS resources
2034564	I/Os hung in serialization after one of the disk which formed the raid5 volume was pulled out
2113831	vxconfigd core dumps while including the previously excluded controller
2112568	System panics while attaching back two Campus Cluster sites due to incorrect DCO offset calculation
2126731	VxVM 5.1: vxdisk -p list output is not consistent with previous versions

Veritas Volume Manager: Issues fixed in 5.1 RP1

Table 1-7 Veritas Volume Manager 5.1 RP1 fixed issues

Fixed issues	Description
1938484	EFI: Prevent multipathing don't work for EFI disk
1915356	I/O stuck in vxvm caused cluster node panic
1899688	[VVR] Every I/O on smartsync enabled volume under VVR leaks memory
1884070	When running iotest on volume, primary node runs out of memory
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1860892	Cache Object corruption when replaying the CRECs during recovery
1857729	CVM master in the VVR Primary cluster panic when rebooting the slave during VVR testing
1857558	[CVM] Need to ignore jeopardy notification from GAB for SFCFS/RAC, since oracle CRS takes care of fencing in this stack
1840673	After adding new luns one of the nodes in 3 node CFS cluster hangs
1835139	CERT : pnate test hang I/O greater than 200 seconds during the filer giveback
1826088	After pulling out FC cables of local site array, plex became DETACHED/ACTIVE
1792795	supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.

Veritas Storage Foundation Cluster File System fixed issues

This section lists fixed issues for Veritas Storage Foundation Cluster File System.

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

Table 1-8 Veritas Storage Foundation Cluster File System 5.1 RP2 fixed issues (listed incident number, parent number)

Fixed issues	Description
1982730, 1952484	Fixed a panic in `vx_recv_getemapmsg`() due to an alignment fault.
2043651, 1991446	Changing the nodes in a cluster from `largefiles` to `nolargefiles` with the `fsadm` command no longer results in the following error when you re-mount the nodes: `UX:vxfs mount: ERROR: V-3-21272: mount option(s) incompatible with file system`
1933839, 1807536	Added support for `VX_FREEZE_ALL` ioctl in a cluster environment.
2069672, 2069059	Fixed a hang issue in a cluster environment.
2049381, 2049378	Fixed an issue that caused database checkpoint rollback to fail on a non-English locale setup.
2092088, 2030289	Fixed a file system corruption issue in a cluster environment that occurred when mounting the file system on the secondary node while the primary node was 100% full.

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

Table 1-9 Veritas Storage Foundation Cluster File System 5.1 RP1 fixed issues (listed incident number, parent number)

Fixed issues	Description
1878583, 1544221	getattr call optimization to speedup the case when binaries are being mmapped from many nodes on CFS.

Veritas Cluster Server fixed issues

This section lists fixed issues for Veritas Cluster Server.

Veritas Cluster Server: Issues fixed in 5.1 RP2

Table 1-10 Veritas Cluster Server 5.1 RP2 fixed issues

Fixed issues	Description
2019899	Fixed issue with notifier in binding to specific source IP
2102777	Fixed RemoteGroup resource to support remote groups that are a parallel SG.
2083268	Disabled RFC check with DNS Agent.
2080858	Fixed the issue with IP resource can't online when both eth0:0/eth0:255 were used
1992561	Fixed vxfentsthdw from connecting through SCP
2102798	VCS errors out with V-16-1-13027 (rplnfs01) because it did not handle NFS services (RPLNFS_nfs resource) correctly. Monitor procedure did not complete within expected time causing assertion failure on AgentFramework line 5015
2102802	NASGW: NFS: VCS share agent hostname comparison is case sensitive.
2107658	LXRT51RP2: hagrp -state oradb_grp Segmentation fault with gco config
2115132	hastop -local -force or HAD restart causes Share resource offline after restart.
2111294	Fixed problem in which some VCS nodes in a group are online and some are not.
2119570	Fixed problem that prevented Oracle detailed monitoring from sending SNMP notification on WARN signal.
2120020	Fixed problem in DB2 resource that generated excessive logging in engine_A.log.
2125442	Fixed an issue with had to online the failover service group only if it is offline on all nodes after the service group is unfrozen.
1937834	Make sure that 1) Partial parent group can auto-start after child online. 2) Service group failover after node panic.
2066967	Improve heartbeating logic b/w engine and agfw during snapshotting and during steady cluster operation.
2071549	Under heavy load, _had may wait indefinetly in waitpid() when user requests to run some command through "hacli", leading GAB killing _had

Table 1-10 Veritas Cluster Server 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
2077363	Overriding static attribute for any resource with empty type definition creates duplicate entries in main.cf
2077375	Parallel group in PARTIAL state autostarts in case of engine restart.
2079617	NFSRestart agent should remain online if HAD is forcefully stopped and restarted.
2079622	nfs splitbrain test: Delay forced import of disk group agent to avoid errors to client nodes.
2079664	ClusterService Group autopopulates system list
2080703	Memory leak in command hareg -group
2084977	stack overflow seen for a very large filename
2089182	vxcpserv process memory keeps growing incase cpsadm commmands are executed
2102764	RemoteGroup resource detected offline even with patch for e1834858
2110465	vxfenclearpre cannot clear keys from coordinator disks and data disks when there's a preexisting split brain.

Veritas Cluster Server: Issues fixed in 5.1 RP1

Table 1-11 Veritas Cluster Server 5.1 RP1 fixed issues

Fixed issues	Description
1973340	LLT: In llt_send_port() increment the wrenable_clue if we can't dupmsg
1971269	[VCS] Prerequisites/Limits attributes not being honored if service group faults during switch
1969999	[VCS51RP1][ENGINE][Notifier] SNMP Traps receiver shows trim output.
1962548	lxrt5.0mp4:sfora:sles10sp3_ppc_only ASM agent coring.
1961026	[VCS][411-279-558] Host unable to reach after MultiNICA (PM) failover
1960735	[SFHA][5.1RP1] big service group online need a long time

Table 1-11 Veritas Cluster Server 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1958122	[VCS5.0.1RP1][OracleASM]- Oracle ASM resource does not come to ONLINE state upon reboot in 11gR2 setup
1950427	[VCSOR] ASMDGAgent should disable and enable diskgroups in offline and online EPs for 11g R2.
1948627	[Oracle Agent] Add check for ohasd daemon for 11g R2 in ASMInst agent.
1941647	haalert CLI hangs if engine is not in running state.
1937672	ASM agent not detecting cssd process for 11gR2
1922411	vxfentsthdw should detect storage arrays which interpret NULL keys as valid for registrations/reservations
1916004	ASMagent connecting as sysdba instead of sysasm for 11gR2
1915909	[VCS][281-889-442] hares allows to create resources which has "." special character
1915627	had dumped core with "ASSERTION FAILED: file Group.C, line 14468, expression (!retval || !STRCMP(retval, snvelemp->name()))"
1885710	remove reference to VERITAS from message id 53021
1874267	[ENGINE] Don't set MonitorOnly to 0 if ExternalStateChange does not have "OfflineGroup" value
1870424	LLT should give error if an attempt is made to configure more than 8 links (LLT_MAX_LINK) under LLT
1504123	[SFW-HA 5.1 GCO] Symantec SE - GCO failover does not work when user account has "!" in name.

Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues

This section lists fixed issues for Veritas Storage Foundation Cluster File System for Oracle RAC.

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

The 5.1 RP2 release does not include any fixed issues for Veritas Storage Foundation Cluster File System for Oracle RAC.

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

The 5.1 RP1 release does not include any fixed issues for Veritas Storage Foundation Cluster File System for Oracle RAC.

Veritas Storage Foundation Manager fixed issues

This section lists fixed issues for Veritas Storage Foundation Manager.

Storage Foundation Manager: Issues fixed in 5.1 RP2

The 5.1 RP2 release does not include any fixed issues for Storage Foundation Manager.

Storage Foundation Manager: Issues fixed in 5.1 RP1

Table 1-12 Storage Foundation Manager 5.1 RP1 fixed issues

Fixed issues	Description
1934914	Configuration fails if 2.1 CS is not configured and directly upgraded to 2.1RP1 CS
1931017	Copyright year for Windows, Solaris and HP-UX patches are 2009
1918582	Licenses not getting discovered in case default locale is non-English
1917308	when had is stopped/started vcs based monitoring should continue to function
1910997	Checkpoint size showing zero in Webgui
1904090	LDR fails to display deployment summary
1897156	Paths are not shown for one of the array ports whereas Luns information is shown
1894441	'Refresh host' needed to populate the MHs info, after upgrading package/patch through sysaddon

Table 1-12 Storage Foundation Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1893699	Unable to add a host to the management server. V-39-4095-903 401 Unauthorized User Error
1893244	Unable to add a host to the management server. V-39-4095-803 401 Unauthorized User Error
1889739	LoP hosts get list out in 'Not Installed Hosts', when deployed the sysaddon for Linux x86 MH
1888082	After deploying sysaddon patch the operation status pop up is not having host details
1887241	remove use of threads in Perl discovery
1878876	vxlist core dumping after server firmware upgrade
1878266	too many hareg processes seen on a machine where sfmh is installed
1873461	DCLI does not properly handle 2 vdids for one OShandle
1872805	prtdiag and psrinfo -v not supported in Solaris 8, causing LDR not to display correct results
1869752	Add support for DB2 9.x support
1865225	IPv6 address not discovered in SFM gui for AIX hosts
1861664	Fix the library path for gvdid to work in case of HP 11.11
1858963	SFMH is uninstalled even if it was installed prior to install of SFW/SFWHA
1857468	VEA/vxpal continuously generate errors 0xc1000039 in vm_vxisis.log with no apparent reason
1855466	When a VVR RVG goes offline it is reported as at risk, however when it goes online again the state does not change in the UI
1855087	vxlist incorrectly shows nolabel flag for labeled disks
1854459	db2exp process is frequently core dumping on cluster node
1853081	vxship missing in VRTSsfmh for Linux
1850797	DMP Connectivity Summary view slow and causes high db CPU
1839795	Path type is empty on HP for SF 5.0 on 11.31-IA/PA

Table 1-12 Storage Foundation Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1831711	Volume Migration fails because it cannot find a target enclosure
1831697	Managing Storage Enclosure Summary reports 1 enclosure when actually 3 exist
1827451	Addhost log information is off by one month
1826556	dcli vdid can fail on HPUX LVM disks
1826409	SFM needs vxsvc service running to administer but service is not started
1825858	CS showing wrong gab port information
1809918	Servlet Exception error after adding Opteron MH to CS
1804496	postremove error messages on SFM uninstall
1797382	SFM is reporting numerous could not set locale correctly messages in error.log
1791528	VRTSsfmh error log reporting numerous errors from managed hosts
1791063	dclisetup.sh needs to be run again after upgrade to VxVM 5.1
1712298	WEBUI shows MH status as "Faulted - VEA: vxsvc or StorageAgent is not running" though all services running

Veritas Enterprise Administrator fixed issues

This section lists fixed issues for Veritas Enterprise Administrator.

VEA: Issues fixed in 5.1 RP2

The 5.1 RP2 release does not include any fixed issues for VEA.

VEA: Issues fixed in 5.1 RP1

Table 1-13 VEA 5.1 RP1 fixed issues

Fixed issues	Description
1961519	vxsvc running as a daemon shows stderr and stdout printf's

Table 1-13 VEA 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1958763	isisd wont start, core file generated.
1958351	VEA gui fails to show controller-enclosures mapping.
1954150	Appropriate message should be display while creating Multiple Volume when size is incorrect
1954118	Not able to edit Log Settings for Alert/Task log.
1954101	While launching Gui, VEA Error message thrown "creating an instance of a class vrts.vvr.ce.REntryPoint failed"
1954047	Incorrect host version in VEA gui for 5.1RP1.
1953701	vxsvc does not start after installing RP1.
1925365	the replicated data size is showing with a negative value in VEA. (>TB)
1879928	Finish button for Break-off Snapshot for a Vset does nothing
1873583	VVR event notification sending 2 messages per event
1857207	Enabling FastResync has no effect when creating a RAID-5 volume
1846581	Core generated while downloading extension using client utility.
1840050	Core got generated while performing Volume Set operation.
1635720	Need to support volume tagging related operations of GUI in VMPROVIDER

Known issues

The following are new additional Storage Foundation and High Availability known issues in this release.

- Installation and upgrade known issues
- Veritas Storage Foundation known issues in 5.1 RP2 release
- Veritas Volume Manager known issues in 5.1 RP2 release
- Veritas File System known issues in 5.1 RP2 release
- Veritas Storage Foundation Cluster File System known issues in 5.1 RP2 release
- Veritas Storage Foundation Cluster File System for Oracle RAC known issues in 5.1 RP2

- Veritas Cluster Server known issues in 5.1 RP2

For the 5.1 known issues, see the 5.1 Release Notes for your Veritas product.

Installation and upgrade known issues

The following are new additional installation and upgrade known issues in this 5.1 RP2 release.

Installing the latest Support RPM (VRTSspt)

If you plan to upgrade from version 5.1 to version 5.1 RP2, and you have not installed the 5.1 P1 patch, you will not get the latest Support RPM.

Workaround:

You can get the latest VRTSspt RPM by following this link http://entsupport.symantec.com/docs/261451 and performing the instructions to connect to the FTP server and download the RPM.

Veritas Storage Foundation known issues in 5.1 RP2 release

The following are new additional Storage Foundation known issues in this 5.1 RP2 release.

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

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

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

Or

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

Workaround:

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

Veritas Volume Manager known issues in 5.1 RP2 release

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

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

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

Or

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

Workaround:

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

LVM volumes cannot be converted by vxvmconvert utility (1809789)

Because of changes in the LVM package, the `vxvmconvert` utility cannot convert LVM diskgroups to VxVM diskgroups after LVM version 2.02.32. LVM version 2.02.32 is the last known working version. Attempting to convert LVM volumes later than version 2.02.32 fails and the data becomes corrupted, nor can the failed conversion be reverted.

Workaround: There is no workaround for this issue.

Reclamation with fault injection does not reclaim any space (1984696)

This issue is found in some customer's array when paths to the LUNs are disabled during reclamation.

There is no workaround at this time.

Veritas File System known issues in 5.1 RP2 release

No additional known issues exist for Veritas File System in the 5.1 RP2 release.

Panic due to null pointer de-reference in vx_unlockmap(2059611)

The above mentioned issue partially fixes the panic problem by avoiding the access to the pointer which is NULL. However the reason for the this pointer to be NULL is yet to be RCA'd and the complete fix for the issue will be released in the next patch.

There is no workaround for this issue.

Veritas Storage Foundation Cluster File System known issues in 5.1 RP2 release

The following are new additional Veritas Storage Foundation Cluster File System known issues in this 5.1 RP2 release.

NFS issues with VxFS checkpoint (1974020)

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

```
Stale NFS file handle
```

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

Workaround

For SFCFS:

- You can specify the `fsid share` option during `cfsshare share` to force the `fsid` of an NFS-exported VxFS checkpoint to remain the same on all cluster nodes.

  For example:

  To NFS-export a VxFS checkpoint of a VxFS file system that has already been added to VCS configuration and mounted at `/ckpt1`, run the following command:

  ```
  # cfsshare share /ckpt1 "fsid=num"
  ```

  where *num* is any 32-bit number that is unique amongst all the exported file systems.

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

For SFHA:

◆ You can modify the Options attribute of the Share resource corresponding to the VxFS checkpoint and add the fsid share option to force the fsid of an NFS-exported VxFS checkpoint to remain the same on all cluster nodes.

See the exports(5) manual page for more information.

For example:

If share1 is the VCS Share resource corresponding to an NFS-exported VxFS checkpoint, then run the following commands:

```
# haconf -makerw
# hares -modify share1 Options "fsid=num"
# haconf -dump -makero
```

where *num* is any 32-bit number that is unique amongst all the exported filesystems.

The installrp displays SFCFSRAC is installed instead of SFCFS (1956921)

If SFCFS 5.1 RP1 is already installed on all the systems and you rerun installrp to install SFCFS 5.1 RP1, it displays the following message:

```
SFCFS Oracle RAC version 5.1.00.100 is already installed on redhat92157
SFCFS Oracle RAC version 5.1.00.100 is already installed on redhat95241
```

The message displays that SFCFSRAC is installed instead of SFCFS.

You can safely ignore this message.

Possible issue with CVM master takeover following the detection of a network split brain condition between cluster nodes. (2110589)

In a Veritas Storage Foundation Cluster File System (SFCFS) environment, if the cluster volume manager (CVM) master node is fenced to avoid a split-brain condition, the node sometimes increment the serial (SSB) ID on disks inconsistently before being fenced. This can cause shared disk group imports on the node that try to take over as master to fail, which in turn cause the takeover to fail.

Workaround: Use the vxsplitlines output to determine the vxdg commands that you run to import the shared disk groups using the available configuration copies.

Mounting a file system as seconly by using the cfsmount command may fail (2108603)

If you try to mount a file system with the seconly mount option by using the cfsmount command, the mount operation may fail with the following error:

Error: V-35-50: Could not mount *volume_name* at mount_point on *node name*

Look at VCS engine_A.log on *node_name* for possible errors for resource cfsmount1

The mount operation fails because mounting of the seconly file system is attempted before the primary mount operation is complete. This occurs because of a timing issue in the cfsmount script.

Workaround: There is no workaround for this issue.

Veritas Storage Foundation Cluster File System for Oracle RAC known issues in 5.1 RP2

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

The installer (installrp) may recognize Storage Foundation CFS High Availability as Storage Foundation CFS for Oracle RAC (1956921)

When using the installer (installrp), make sure to explicitly select the appropriate product for installation or upgrade. In some cases, the installer may mistake Storage Foundation CFS High Availability as Storage Foundation CFS for Oracle RAC.

Veritas Cluster Server known issues in 5.1 RP2

The following are new additional Veritas Cluster Server known issues in this 5.1 RP2 release.

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.

VCS agent for Oracle: Make sure that the ohasd has an entry in the init scripts (1985093)

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

VCS agent for Oracle: Intentional Offline does not work

Intentional Offline does not work for the VCS agent for Oracle.

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

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

Workaround:

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

Health check monitoring for Oracle Agent on SLES11 platform (1938167)

Oracle Agent does not support health check monitoring on SLES11 platform.

While upgrading the VCS stack, reconfiguration of MultiNICA IPv4RouteOptions attribute is required

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

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

Table 1-14 Whether attributes are configured and required actions that you need to perform during upgrade

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
Configured	May or may not be configured	May or may not be configured	In this case, the `ifconfig` command is used. If RouteOptions is set, attribute value is used to add/delete routes using command `route`. As the Options attribute is configured, IPv4RouteOptions values are ignored.	No need to configure IPv4RouteOptions .
Not configured	May or may not be configured	Must be configured	In this case the `ip` command is used. IPv4RouteOptions must be configured and are used to add/delete routes using the `ip route` command. As Options attribute is not configured, RouteOptions value is ignored.	Configure IPv4RouteOptions and set the IP of default gateway. The value of this attribute typically resembles: IPv4RouteOptions = “default via *gateway_ip*” For example: IPv4RouteOptions = “default via 192.168.1.1”

fdsetup cannot correctly parse disk names containing characters such as "-" (1949294)

The fdsetup cannot correctly parse disk names containing characters such as "-".

RVGPrimary online script does not function correctly (1949293)

The RVGPrimary online script does not function correctly.

Software limitations

The following are additional Veritas Storage Foundation and High Availability software limitations in this release.

- Veritas Storage Foundation software limitations in 5.1 RP2 release
- Veritas Volume Manager software limitations in 5.1 RP2 release

Veritas Storage Foundation software limitations in 5.1 RP2 release

Thin reclamation support limitations

The thin reclamation feature has the following limitations:

- Thin reclamation only supports VxFS file systems on VxVM volumes. Other file systems are not supported.
- Thin reclamation is only supported for mounted volumes.
  The file system map is not available to reclaim the unused storage space on unmounted file systems.
- Thin reclamation is not supported on raw VxVM volumes.
  VxVM has no knowledge of application usage on raw volumes. Therefore, VxVM cannot perform the reclamation on raw volumes. The application must perform the reclamation on raw volumes.
- Thin reclamation is not supported on the RAID-5 layout.
  The thin reclamation is storage dependent and the space underneath may or may not be reclaimed fully. Thin reclamation is not supported in a RAID-5 layout, because data consistency cannot be ensured.
- Thin Reclamation is not supported on volumes with snapshots or snapshots themselves. Any reclamation requests on such volumes or snapshots or their corresponding mount points will not result in any reclamation of their underlying storage.

Veritas Volume Manager software limitations in 5.1 RP2 release

Enable the mpt_disable_hotplug_remove tunable (1663167)

On Red Hat 5 (RHEL 5) systems with direct attached disks, enable the mpt_disable_hotplug_remove tunable so that path-level failover and failback work well.

Workaround

To enable the mpt_disable_hotplug_remove tunable

1 Check the version of the mptsas driver.

2 If the version is older than 4.00.43.00, then remove the old mptsas driver with tbe `rpm -e` command.

3 Install the latest (4.00.42.00 or above) mptsas driver.

4 Check that the system now has the desired version of the mptsas driver.

5 Edit the `/etc/modprobe.conf` file. At the end of the file, add the following line:

```
# options mptsas mpt_disable_hotplug_remove=0
```

6 Rebuild the initrd:

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

7 Reboot the system.

Cluster Volume Manager (CVM) fail back behavior for non-Active/Active arrays (1441769)

This describes the fail back behavior for non-Active/Active arrays in a CVM cluster. This behavior applies to A/P, A/PF, APG, A/A-A, and ALUA arrays.

When all of the Primary paths fail or are disabled in a non-Active/Active array in a CVM cluster, the cluster-wide failover is triggered. All hosts in the cluster start using the Secondary path to the array. When the Primary path is enabled, the hosts fail back to the Primary path. However, suppose that one of the hosts in the cluster is shut down or brought out of the cluster while the Primary path is disabled. If the Primary path is then enabled, it does not trigger failback. The remaining hosts in the cluster continue to use the Secondary path. When the disabled host is rebooted and rejoins the cluster, all of the hosts in the cluster will continue using the Secondary path. This is expected behavior.

For A/P,APG, A/A-A, and ALUA arrays, if the disabled host is rebooted and rejoins the cluster before the Primary path is enabled, enabling the path does trigger the failback. In this case, all of the hosts in the cluster will fail back to the Primary path.

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 DMP restore daemon cycle to 60 seconds. The default value of this tunable is 300 seconds. The change is persistent across reboots.

Issue the following command at the prompt:

```
# vxdmpadm settune dmp_restore_internal=60
```

To verify the new setting, use the following command:

```
# vxdmpadm gettune dmp_restore_internal
```

Documentation addendum

The following sections contain additions to current documents.

Using the preonline_vvr trigger for RVGLogowner resources

For VCS configurations that use RVGLogowner resources, perform the following steps on each node of the cluster to enable VCS control of the RVGLogowner resources. For a service group that contains a RVGLogowner resource, change the value of its PreOnline trigger to 1 to enable it.

To enable the PreOnline trigger from the command line on a service group that has an RVGLogowner resource

- On each node in the cluster, perform the following command:

  ```
  # hagrp -modify RVGLogowner_resource_sg PreOnline 1 -sys system
  ```

  Where *RVGLogowner_resource_sg* is the service group that contains the RVGLogowner resource. The *system* is the name of the node where you want to enable the trigger.

On each node in the cluster, merge the preonline_vvr trigger into the default triggers directory.

To merge the preonline_vvr trigger

- On each node in the cluster, merge the preonline_vvr trigger to the /opt/VRTSvcs/bin/triggers directory.

  ```
  # cp /opt/VRTSvcs/bin/sample_triggers/preonline_vvr \
  /opt/VRTSvcs/bin/triggers
  ```

  Refer to the sample configurations directory for samples of how to enable these triggers (/opt/VRTSvcs/bin/sample_triggers.)

Agent functions

Monitor—Performs read I/O operations on the raw device to determine if a physical disk or a partition is accessible.

State definitions

ONLINE—Indicates that the disk is working normally

FAULTED—Indicates that the disk has stopped working or is inaccessible.

UNKNOWN—Indicates that a problem exists either with the configuration or the ability to determine the status of the resource.

List of patched RPMs

This section lists the RPMs for 5.1 RP2.

Table 1-15 Patches and RPMs for RHEL 5

5.1 RPM names	Products affected	RPM size
VRTScps-5.1.002.000-RP2_RHEL5.i686.rpm	VCS, SFHA, SFCFSHA, SFCFS RAC	32.06 MB
VRTSgab-5.1.002.000-RP2_RHEL5.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	3.04 MB
VRTSllt-5.1.002.000-RP2_RHEL5.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	2.71 MB

Table 1-15 Patches and RPMs for RHEL 5 *(continued)*

5.1 RPM names	Products affected	RPM size
VRTSvcs-5.1.002.000-RP2_RHEL5.i686.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	129.79 MB
VRTSvcsea-5.1.002.000-RP2_RHEL5.i686.rpm	VCS, SFHA, SFCFSHA, SFCFS RAC	11.06 MB
VRTSvcsag-5.1.002.000-RP2_RHEL5.i686.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	616 KB
VRTSvxfen-5.1.002.000-RP2_RHEL5.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	1.39 MB
VRTSvxvm-5.1.002.000-RP2_RHEL5.x86_64.rpm	VM, SF, SFHA SFCFS, SFCFSHA, SFCFSRAC	18.6 MB
VRTSaslapm-5.1.001.000-RP1_RHEL5.x86_64.rpm	VM, SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	200 KB
VRTSlvmconv-5.1.002.000-RP2_RHEL5.i686.rpm	VM, SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	70 KB
VRTSodm-5.1.002.000-RP2_RHEL5.x86_64.rpm	SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	690.9 KB
VRTSdbed-5.1.002.000-RP2_RHEL5.i686.rpm	SF, SFHA, SFCFS, SFCFSHA	16.69 MB
VRTSvxfs-5.1.002.000-RP2_RHEL5.x86_64.rpm	SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	21.23 MB

Table 1-16 Patches and RPMs for SLES 10

5.1 RPM names	Products affected	RPM size
VRTScps-5.1.002.000-RP2_SLES10.i686.rpm	VCS, SFHA, SFCFSHA, SFCFS RAC	32.07 MB
VRTSgab-5.1.002.000-RP2_SLES10.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	8.51 MB
VRTSllt-5.1.002.000-RP2_SLES10.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	5.8 MB
VRTSvcs-5.1.002.000-RP2_SLES10.i586.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	129.23 MB
VRTSvcsag-5.1.002.000-RP2_SLES10.i586.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	10.29 MB
VRTSvcsea-5.1.002.000-RP2_SLES10.i586.rpm	VCS, SFHA, SFCFSHA, SFCFS RAC	10.89 MB
VRTSvxfen-5.1.002.000-RP2_SLES10.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	9.46 MB
VRTSvxvm-5.1.002.000-RP2_SLES10.x86_64.rpm	VM, SF, SFHA SFCFS, SFCFSHA, SFCFSRAC	20 MB
VRTSaslapm-5.1.001.000-RP1_SLES10.x86_64.rpm	VM, SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	135 KB
VRTSlvmconv-5.1.002.000-RP2_SLES10.i686.rpm	VM, SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	60 KB
VRTSodm-5.1.002.000-RP2_SLES10.x86_64.rpm	SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	1.38 MB

Table 1-16 Patches and RPMs for SLES 10 *(continued)*

5.1 RPM names	Products affected	RPM size
VRTSdbed-5.1.002.000-RP2_SLES10.i586.rpm	SF, SFHA, SFCFS, SFCFSHA	16.87 MB
VRTSvxfs-5.1.002.000-RP2_SLES10.x86_64.rpm	SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	31.52 MB

Table 1-17 Patches and RPMs for SLES 11

5.1 RPM names	Products affected	RPM size
VRTScps-5.1.002.000-RP2_SLES11.i686.rpm	VCS, SFHA, SFCFSHA, SFCFS RAC	30.97 MB
VRTSgab-5.1.002.000-RP2_SLES11.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	3.98 MB
VRTSllt-5.1.002.000-RP2_SLES11.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	3.14 MB
VRTSvcs-5.1.002.000-RP2_SLES11.i686.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	109.84 MB
VRTSvcsag-5.1.002.000-RP2_SLES11.i686.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	10.87 MB
VRTSvcsea-5.1.002.000-RP2_SLES11.i686.rpm	VCS, SFHA, SFCFSHA, SFCFS RAC	815.37 KB
VRTSvxfen-5.1.002.000-RP2_SLES11.x86_64.rpm	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	4.54 MB
VRTSvxvm-5.1.002.000-RP2_SLES11.x86_64.rpm	VM, SF, SFHA SFCFS, SFCFSHA, SFCFSRAC	16 MB

Table 1-17 Patches and RPMs for SLES 11 *(continued)*

5.1 RPM names	Products affected	RPM size
VRTSaslapm-5.1.001.000-RP1_SLES11.x86_64.rpm	VM, SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	124 KB
VRTSlvmconv-5.1.002.000-RP2_SLES11.i686.rpm	VM, SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	60 KB
VRTSodm-5.1.002.000-RP2_SLES11.x86_64.rpm	SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	845.19 KB
VRTSvxfs-5.1.002.000-RP2_SLES11.x86_64.rpm	SF, SFHA, SFCFS, SFCFSHA, SFCFSRAC	24.46 MB

Downloading the 5.1 RP2 rolling patch archive

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

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

http://entsupport.symantec.com/docs/335001

Installing the products for the first time

This chapter includes the following topics:

- Installing the Veritas software using the script-based installer
- Installing Veritas software using the Web-based installer

Installing the Veritas software using the script-based installer

This section describes how to install a Storage Foundation and High Availability Solutions product for the first time on a host and install 5.1 RP2. Review the 5.1 Installation Guide and Release Notes for your product.

To install the Veritas software for the first time

1 Mount the 5.1 product disc and navigate to the folder that contains the installation program to install 5.1. Choose one of the following to start the installation:

- For Storage Foundation:

  ```
  # ./installsf node1 node2 ... nodeN
  ```

- For Storage Foundation High Availability:

  ```
  # ./installsf -ha node1 node2 ... nodeN
  ```

- For Storage Foundation Cluster File System:

  ```
  # ./installsfcfs node1 node2 ... nodeN
  ```

- For Storage Foundation Cluster File System High Availability:

  ```
  # ./installsfcfs -ha node1 node2 ... nodeN
  ```

- For Storage Foundation Cluster File System for Oracle RAC:

  ```
  # ./installsfcfsrac -ha node1 node2 ... nodeN
  ```

  Note: There is no SFCFS for Oracle RAC on SLES 11.

- For Veritas Cluster Server:

  ```
  # ./installvcs node1 node2 ... nodeN
  ```

2 Review the installation prerequisites for upgrading to 5.1 RP2.

  See "Prerequisites for upgrading to 5.1 RP2" on page 57.

3 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.

- If the 5.1 product is installed, run the `installrp` script to install 5.1 RP2 and configure the product.

  Note: On SUSE 10 SP3, do not configure 5.1 product until 5.1 RP2 product is installed. As support for SUSE 10 SP3 was just recently released.

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

  See "About the installrp script" on page 12.
  The `installrp` script will ask if you want to configure the product after the 5.1 RP2 installation. If you choose 'no', proceed to step 4.

4 Mount the 5.1 product disc and navigate to the folder that contains the installation program. Run the same 5.1 installation script that you used in step 1, this time specifying the `-configure` option to configure the software.

- For Storage Foundation:

  ```
  # ./installsf -configure node1 node2 ... nodeN
  ```

- For Storage Foundation High Availability:

  ```
  # ./installsf -ha -configure node1 node2 ... nodeN
  ```

- For Storage Foundation Cluster File System:

  ```
  # ./installsfcfs -configure node1 node2 ... nodeN
  ```

- For Storage Foundation Cluster File System High Availability:

  ```
  # ./installsfcfs -ha -configure node1 node2 ... nodeN
  ```

- For Storage Foundation Cluster File System for Oracle RAC:

  ```
  # ./installsfcfsrac -ha -configure node1 node2 ... nodeN
  ```

- For Veritas Cluster Server:

  ```
  # ./installvcs -configure node1 node2 ... nodeN
  ```

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

Installing Veritas software using the Web-based installer

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

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 products with the Veritas Web-based installer

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

To install Veritas product

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

See “Prerequisites for upgrading to 5.1 RP2” on page 57.

2 On the **Select a task and product** page, select **Install 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 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

Upgrading to 5.1 RP2

This chapter includes the following topics:

Prerequisites for upgrading to 5.1 RP2

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

- For any product in the Storage Foundation stack, you must have the 5.1 release installed before you can upgrade that product to the 5.1 RP2 release.
- Each system must have sufficient free space to accommodate patches.

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 5.1 to 5.1 RP2
- 5.1 P1 to 5.1 RP2
- 5.1 RP1 to 5.1 RP2

Upgrading from 5.1 to 5.1 RP2

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

- Performing a full upgrade to 5.1 RP2 on a cluster
  Use the procedures to perform a full upgrade to 5.1 RP2 on a cluster that has VCS, SFHA, SFCFS, or SFCFSRAC installed and configured.
- Upgrading to 5.1 RP2 on a standalone system
  Use the procedure to upgrade to 5.1 RP2 on a system that has SF and VCS installed.
- Performing a rolling upgrade using the installer
  Use the procedure to upgrade your Veritas product with a rolling upgrade.
- Performing a rolling upgrade manually
  Use the procedure to upgrade your Veritas product manually with the rolling upgrade.
- Upgrading to 5.1 RP2 on a system that has encapsulated boot disk
  Use the procedure to upgrade to 5.1 RP2 on a system that has encapsulated boot disk

Performing a full upgrade to 5.1 RP2 on a cluster

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

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

- Freeze service group operations and stop VCS on the cluster.
- Take the nodes offline and install the software patches.
- Bring the nodes online to restart cluster failover services.

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

- Performing a full upgrade to 5.1 RP2 for VCS
- Performing a full upgrade to 5.1 RP2 on an SFHA cluster
- Performing a full upgrade to 5.1 RP2 on an SFCFS cluster
- Performing a full upgrade to 5.1 RP2 on an SFCFS RAC cluster

Performing a full upgrade to 5.1 RP2 for VCS

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

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

See “While upgrading the VCS stack, reconfiguration of MultiNICA IPv4RouteOptions attribute is required” on page 41.

To upgrade VCS

1 Log in as superuser.

2 Verify that **/opt/VRTS/bin** is in your PATH so that you can execute all product commands.

3 Make the VCS configuration writable on a node that is being upgraded:

```
# haconf -makerw
```

4 Freeze the HA service group operations. Enter the following command on each node, if you selected a group of nodes on which to upgrade the operating system:

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

5 Make the VCS configuration read-only:

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

6 Close any instance of VCS GUI that is running on the node.

7 If you plan to upgrade Operating System at this time before upgrading VCS continue with step 8, otherwise go to 12.

8 Stop VCS:

```
# hastop -all
```

9 Stop the VCS command server:

```
# ps -ef | grep CmdServer
```

```
# kill -9 pid_of_CmdServer
```

where *pid_of_CmdServer* is the process ID of CmdServer.

10 Stop cluster fencing, GAB, and LLT. On each node, type:

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

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

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

11 Upgrade the Operating System and reboot the systems if required.

See "System requirements" on page 16.

12 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

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

13 Resolve any issues that the precheck finds.

14 Start the upgrade:

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

15 After the upgrade, review the log files for any issues.

16 After all of the nodes in the cluster are upgraded, shut down and reboot each of the nodes. After the nodes come up, application failover capability is available.

17 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

18 Unfreeze the service group operations on each node:

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

19 Make the VCS configuration read-only:

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

20 Verify the upgrade.

See "Verifying software versions" on page 91.

Performing a full upgrade to 5.1 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 RP2 on an SFHA cluster

1 Log in as superuser.

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

3 Make the VCS configuration writable on a node that is being upgraded:

```
# haconf -makerw
```

4 Freeze the HA service group operations. Enter the following command on each node, if you selected a group of nodes on which to upgrade the operating system:

```
# hagrp -freeze groupname -persistent
```

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

5 Make the VCS configuration read-only:

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

6 Close any instance of VCS GUI that is running on the node.

7 Stop VCS:

```
# hastop -local
```

8 Stop the VCS command server:

```
# ps -ef | grep CmdServer
# kill -9 pid_of_CmdServer
```

where `pid_of_CmdServer` is the process ID of CmdServer.

9 Stop cluster fencing, GAB, and LLT.

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

10 If required, upgrade the OS.

See “System requirements” on page 16.

See the appropriate OS documentation for the procedures.

11 Repeat step 7 through step 9 if the system reboots after upgrading the operating system. You need to perform this to stop the components that are started by the init scripts, if any.

12 From the directory that contains the extracted and untarred 5.1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

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

13 After all of the nodes in the cluster are upgraded, shut down and reboot each of the nodes. After the nodes come up, application failover capability is available.

14 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

15 Unfreeze the service group operations on each node:

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

```
# hagrp -freeze groupname -persistent
```

16 Make the VCS configuration read-only:

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

Performing a full upgrade to 5.1 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 RP2 on an SFCFS cluster

1 Log in as superuser.

2 Verify that /opt/VRTS/bin is in your PATH so that you can execute all product commands.

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

```
# haconf -makerw
```

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

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

5 Make the configuration read-only:

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

6 Determine if each node's root disk is under VxVM control and proceed as follows.

- Check if each node's root disk is under VxVM control:

  ```
  # df -v /
  ```

  The root disk is under VxVM control if /dev/vx/dsk/rootvol is listed as being mounted as the root (/) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the vxplex command to remove all the plexes of the volumes rootvol, swapvol, usr, var, opt and home that are on disks other than the root disk. For example, the following command removes the plexes mirrootvol-01, and mirswapvol-01 that are configured on a disk other than the root disk:

  ```
  # vxplex -o rm dis mirrootvol-01 mirswapvol-01
  ```

  Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

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

  Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

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

```
# mount | grep vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# mount | grep vxfs
```

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

```
# umount /filesystem
```

Note: If file system is CFS mounted then use `cfsumount` command.

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 Stop all VxVM volumes by entering the following command for each disk group on master node:

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

Verify that no volumes remain open:

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

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 If ODM is installed and port 'd' is up. Stop ODM service using the following command:

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

15 On each node, stop cluster fencing, GAB, and LLT.

```
# /etc/init.d/vxgms stop
# /etc/init.d/vxglm stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

16 If required, apply the OS kernel patches.

17 From the directory that contains the extracted and untarred 5.1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

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

18 After all nodes in the cluster are upgraded, the processes will be restarted automatically. Should there be any problem, the `installrp` script will ask you to reboot the system. Then the application failover capability will be available.

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

20 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

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

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

22 Make the configuration read-only:

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

23 Bring the CVM service group online on each node:

```
# hagrp -online cvm -sys nodename
```

24 Restart all the volumes by entering the following command on the master node, for each disk group:

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

25 If you stopped any RVGs in step 9, restart each RVG:

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

26 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

27 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

Performing a full upgrade to 5.1 RP2 on an SFCFS RAC cluster

To prepare for a full upgrade to 5.1 RP2 on an SFCFS RAC cluster

1 Log in as superuser.

2 Verify that /opt/VRTS/bin is in your PATH so you can execute all product commands.

3 Stop the applications that are not managed by VCS. Use native application commands to stop the application.

4 Stop Oracle Clusterware.

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

5 Stop high-availability cluster operations. This command can be executed from any node in the cluster, and stops cluster operations on all the nodes.

```
# hastop -all
```

6 Check if each node's root disk is under VxVM control by running this command:

```
# df -v /
```

7 The root disk is under VxVM control if /dev/vx/dsk/rootvol is listed as being mounted as the root (/) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the vxplex command to remove all the plexes of the volumes rootvol, swapvol, usr, var, opt and home that are on disks other than the root disk. For example, the following command removes the plexes mirrootvol-01, and mirswapvol-01 that are configured on a disk other than the root disk:

  ```
  # vxplex -o rm dis mirrootvol-01 mirswapvol-01
  ```

  Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

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

  Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

8 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

9 Unmount all file systems:

```
# umount /filesystem
```

10 Check if the VEA service is running:

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

11 If the VEA service is running, stop it:

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

12 If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.

13 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.

14 Stop all VxVM volumes by entering the following command for each disk group:

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

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

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

16 Disable the startup scripts before upgrading the operating system.

```
# insserv -r vcs
# insserv -r vxodm
# insserv -r vxfen
# insserv -r vxgms
# insserv -r vxglm
# insserv -r gab
# insserv -r llt
```

17 Upgrade the operating system. For instructions, see the operating system documentation.

To upgrade Storage Foundation Cluster File System for Oracle RAC

1 From the directory that contains the extracted and untarred 5.1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

2 After the initial system checks are complete, press Return to start the requirement checks.

3 When the Upgrade is complete, note the locations of the summary, log, and response files indicated by the installer.

4 Shut down and reboot the systems.

5 Upgrade Oracle RAC, if required.

6 Relink the Oracle's ODM library with Veritas ODM library.

- For Oracle RAC 10g:
  - Change to the $ORACLE_HOME/lib directory:

    ```
    # cd $ORACLE_HOME/lib
    ```

  - Back up libodm10.so file.

    ```
    # mv libodm10.so libodm10.so.oracle-`date '+%m_%d_%y-%H_%M_%S'`
    ```

  - Link libodm10.so file with the Veritas ODM library:

    ```
    # ln -s /opt/VRTSodm/lib64/libodm.so libodm10.so
    ```

- For Oracle 11g:
  - Change to the $ORACLE_HOME/lib directory:

    ```
    # cd $ORACLE_HOME/lib
    ```

  - Back up libodm11.so file.

    ```
    # mv libodm11.so libodm11.so.oracle-`date '+%m_%d_%y-%H_%M_%S'`
    ```

  - Link libodm11.so file with the Veritas ODM library:

    ```
    # ln -s /opt/VRTSodm/lib64/libodm.so libodm11.so
    ```

To bring the upgraded cluster online and restore components

1 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.

2 If necessary, reinstate any missing mount points in the /etc/fstab file on each node.

3 If any VCS configuration files need to be restored, stop the cluster, restore the files to the /etc/VRTSvcs/conf/config directory, and restart the cluster.

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

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

5 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

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

Upgrading to 5.1 RP2 on a system that has encapsulated boot disk

You can use this procedure to upgrade to 5.1 RP2 on a system that has encapsulated boot disk.

Note: Upgrading with encapsulated boot disk from 5.1 to 5.1 RP2 requires multiple reboots.

To upgrade to 5.1 RP2 on a system that has encapsulated boot disk

1 Manually unmount file systems and stop open volumes.

2 If required, manually break the mirror and un-encapsulate boot disk.

3 Reboot the system to have un-encapsulation take effect .

4 Upgrade to 5.1 RP2 using `installrp` command.

5 After upgrading, reboot the system to have the new VM drivers take effect.

6 Re-encapsulate the boot disk.

7 Again reboot to have the re-encapsulation take effect.

Upgrading to 5.1 RP2 on a standalone system

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

To upgrade to 5.1 RP2 on a standalone system

1 Log in as superuser.

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

3 If required, apply the OS kernel patches.

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

```
# df -T |  grep vxfs
```

5 Unmount all Storage Checkpoints and file systems:

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

6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

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

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

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

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

7 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

8 Stop all VxVM volumes by entering the following command for each disk group:

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

Verify that no volumes remain open:

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

9 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

10 Navigate to the folder that contains the installation program. Run the `installrp` script:

```
# ./installrp nodename
```

11 If necessary, reinstate any missing mount points in the `/etc/filesystems` file.

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

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

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

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

14 Remount all VxFS file systems and Storage Checkpoints:

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

15 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
```

Performing a rolling upgrade using the installer

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

- About rolling upgrades
- Prerequisites for a rolling upgrades
- Performing a rolling upgrade on kernel packages: phase 1
- Performing a rolling upgrade on non-kernel packages: phase 2

About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime. Rolling upgrades are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

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

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System
- Storage Foundation Cluster File System and High Availability
- Storage Foundation Cluster File System for Oracle RAC

You can perform a rolling upgrade from 5.1 to 5.1 RP2 or from 5.1 RP1 to 5.1 RP2.

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.
- For SFCFS for Oracle RAC, stop Oracle CRS before upgrading Kernel packages on any node.

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: 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 SFCFS 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.

7 The installer starts all the relevant processes and brings all the service groups online.

8 If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step 9.

9 Before you proceed to phase 2, complete step 3 to step 7 on the second subcluster.

Performing a rolling upgrade on non-kernel packages: phase 2

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

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

1 Stop all applications accessing volumes.

2 Unmount any filesystems on the nodes to be upgraded.

Note: Only locally mounted file systems need to be unmounted. File systems mounted by CFS will be unmounted by the installer.

3 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...
```

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

5 The installer upgrades non-kernel RPMs.

6 The installer reboots nodes that use encapsulated boot disks.

7 The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1. The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.

8 Verify the cluster's status:

```
# hastatus -sum
```

Performing a rolling upgrade manually

You can perform a split stack rolling upgrade–an independent upgrade of CFS/ODM/CVM/LLT/GAB/VXFEN and the VCS Engine (‘had’).

- Split stack rolling upgrade–an independent upgrade of CFS/ODM/CVM/LLT/GAB/VXFEN and the VCS Engine (‘had’)

Split stack rolling upgrade—an independent upgrade of CFS/ODM/CVM/LLT/GAB/VXFEN and the VCS Engine (‘had’)

Review the following notes:

- It is possible to conduct Rolling Upgrade of one node at a time.
- Recommended for clusters of any number of nodes and Service Group distributions, including N+1 configurations.
- Failover Service Groups will incur downtime 2 times, during failover/failback.

To perform a split stack rolling upgrade—an independent upgrade of CFS/ODM/CVM/LLT/GAB/VXFEN and the VCS engine (‘had’)

1 Consider a four node SFCFSRAC cluster. Identify sub-clusters to be upgraded together. A sub-cluster could even be just one of the nodes of the cluster.

2 Review cluster’s system list. Confirm that each Service Group will eventually have a target node to run on, when sub-clusters are upgraded in a rolling fashion.

3 Verify that `/opt/VRTS/bin` and `/opt/VRTSodm/bin` are added to `PATH` variable.

4 Display the system list:

```
# hagrp -display ServiceGroup -attribute SystemList
```

5 On the sub-cluster to be upgraded, run module specific commands as below for LLT, GAB, VXFEN, CVM, CFS, ODM on one of the nodes of the sub-cluster to be upgraded, to get the current protocol version. This version need not be same for all modules.

```
# lltconfig -W
# gabconfig -W
# vxfenconfig -W
# vxdctl protocolversion
# fsclustadm protoversion
# odmclustadm protoversion
```

6 On the sub-cluster to be upgraded, stop all the applications and resources that are not under VCS control but are still using CVM and CFS stack.

- Make the VCS configuration writable on a node that is being upgraded:

  ```
  # haconf -makerw
  ```

- Freeze the HA service group operations. Enter the following command on node:

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

- Close any instance of VCS GUI that is running on the node.

7 Switch the failover Service Groups from the sub-cluster to be upgraded, to the other sub-cluster. The following command needs to be run for each affected Service Group on each node where the Service Group is active, on the sub-cluster to be upgraded. You may also specify a target node for a given Service Group, as required. However there is a downtime to the failover Service Groups at this stage as part of the switch.

```
# hagrp -switch ServiceGroup -to target_system_name
```

8 Validate that the Service Groups are switched over as desired. In case the switch didn’t succeed for any of the Service Groups, the user still has a window available to make any changes to the impacted Service Groups at this stage.

9. Unmount all vxfs file systems on the sub-cluster to be upgraded. Enter the following command to check if any Storage Checkpoints are mounted:

   ```
   # df -F vxfs
   ```

   If any VxFS file systems are present, unmount all of the VxFS file systems:

   ```
   # cfsumount /filesystem nodename
   ```

10. Stop ‘had’ on the sub-cluster to be upgraded, and switch any remaining failover Service Groups that are online on this sub-cluster atomically.

    ```
    # hastop -local -evacuate
    ```

    Review the following notes:

    - If all the Service Groups had switched over in step 6 itself, the ‘evacuate’ operation for the above command is idempotent.
    - With the above step, it is ensured that if one of the nodes in the remaining sub-cluster goes down at this stage, the Service Groups that have already been moved to the remaining sub-cluster will not attempt to switch back to any of the nodes on the sub-cluster being upgraded. Any pending switches can also occur in this step.
    - The parallel Service Groups on the nodes of the sub-cluster to be upgraded are brought down at this stage. They will continue to be available on the remaining sub-cluster.
    - CVM, CFS will also be stopped by VCS on the nodes of the sub-cluster being upgraded. They will continue to be available on the remaining sub-cluster.

11. Stop applications/resources that are outside VCS control and use VxFS, VxVM.

12 Manually update the `/etc/vxfenmode` and `/etc/gabtab` files to indicate the protocol version that the corresponding module in the new stack should talk to that on the older stack on each of the nodes. This protocol version is the same as the one obtained in step 5. For CVM, CFS and ODM, run the following commands on each of the nodes, to set the protocol version.

```
# vxdctl setversion N
# fsclustadm protoset N
# odmclustadm protoset N
```

where *N* is the protocol version derived in step 5.

This step ensures that the sub-clusters consistently communicate at the older protocol version should there be any intermediate node joins/leaves until the entire cluster is explicitly rolled over to communicate at the new version.

For example, for /etc/vxfenmode:

```
# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# sybase     - use scsi3 disks in kernel but coordinate membership
# with Sybase ASE
# disabled   - run the driver but don't do any actual fencing
#

vxfen_mode=disabled
vxfen_protocol_version=10


# cat /etc/gabtab
/sbin/gabconfig -c -n4 -V33
```

13 Stop VXFEN, ODM, GMS, GLM, GAB and LLT in that order, on each of the nodes of the sub-cluster to be upgraded.

14 Simultaneously upgrade all components except the VCS Engine (‘had’) on the sub-cluster chosen for upgrade. VCS engine and agent related packages are not upgraded at this stage. CFS, ODM, CVM, GAB, LLT, VXFEN will be upgraded together.

- Upgrade all the packages with new product version, except VCS and agent related packages on the sub-cluster being upgraded.
- Re-link oracle in case of SFCFS RAC.
- Reboot all the nodes in the upgraded sub-cluster.
- After reboot, the VCS/SFHA or SFCFS stacks on the upgraded sub-cluster should come up automatically.
- Note that all components (except VCS engine) on the upgraded sub-cluster, will continue to communicate with the nodes of the remaining sub-cluster at the older protocol version at this stage.
- Switch back the failover Service Groups from the remaining sub-cluster to the upgraded sub-cluster. There is a downtime involved for failover Service Groups during the switch.

  ```
  # hagrp -switch ServiceGroup -to target_system_name
  ```

15 Upgrade the remaining sub-cluster(s) one by one, per above procedure from step 4 onwards.

16 After each of the nodes are upgraded to the new product version, initiate a cluster-wide and across-the-stack rollover of the kernel stack to the new protocol version.

- LLT are already at new protocol version at the end of step 14.
- Run `gabconfig -R` on one of the nodes of the cluster being upgraded. This command will block until roll over is complete cluster wide. GAB also quiesces I/Os, which will result in flow control.
- Run `vxfenconfig -R` on one of the nodes of the cluster being upgraded. Wait till the command returns.
- Run `vxdctl upgrade` on the CVM master node of the cluster being upgraded.
- Run `fsclustadm protoclear` to clear the set protocol version on all the nodes in the cluster.
- Run `fsclustadm protoupgrade` from any node of cluster to upgrade the protocol version across the cluster.
- Run `odmclustadm protoclear` to clear the set protocol version on all nodes.
- Run `odmclustadm protoupgrade` on one of the nodes of the sub-cluster being upgraded.

While upgrading odmcluster protocol version, you might see a message like:

```
"Protocol upgrade precheck fails:
        some nodes do not support multiple protocols"
```

You can ignore this message. The odm module is running on the latest version. You can verify this by using the following command on all the upgraded nodes:

```
# odmclustadm protoversion
Cluster Protocol Versions:
Node     #PROTOCOLS   CUR     PREF    FLAGS
local:   3            3       -
```

- Reverse the changes done to `/etc/vxfenmode` and `/etc/gabtab` files in step 12 above.

17 Upgrade VCS engine (‘had’) to the new version. Perform one of the following procedures:

- Force stop 'had' and install the new version.
  - Force stop ‘had’ on all the nodes. There is no HA from this point onwards.

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

  - Back up VCS configuration file including:
    /etc/sysconfig/vcs
    /etc/VRTSvcs/conf/config/types.cf
    /etc/VRTSvcs/conf/config/main.cf
  - Uninstall VRTSvcs and agent related packages.
  - Modify the VCS configuration to reflect version specific requirements if any.
  - Install new version of VRTSvcs and agent related patches.
  - Start VCS on all nodes. HA for the entire cluster is restored at this stage.
- Upgrade ‘had’ in a phased manner. This procedure will reduce the overall HA downtime during the upgrade.
  - Divide the cluster into two sub-clusters. Upgrade the first sub-cluster.

- Force stop VCS on the sub-cluster. On each node of the sub-cluster, run following command:

  ```
  # hastop -local -force
  ```

  There will be no HA for the sub-cluster to be upgraded from this step onwards.

- Uninstall VRTSvcs and agent related packages.
- Modify the VCS configuration to reflect version specific requirements if any.
- Install new version of VRTSvcs and agent related patches.
- Force stop VCS on the remaining sub-cluster. On each node of the remaning sub-cluster, run:

  ```
  # hastop -local -force
  ```

  There is no HA for the entire cluster from this point onwards.

- Start VCS on each of the nodes of the upgraded sub-cluster. VCS will not online the failover Service Groups at this time since they are autodisabled. Now HA is restored for the upgraded sub-cluster.

  ```
  # hastart
  ```

- Upgrade the remaining sub-cluster.
- Uninstall VRTSvcs and agent related packages.
- Install VRTSvcs and agent related patches.
- Start VCS on each of the nodes of the remaining sub-cluster. Now HA is restored for the entire cluster.

  ```
  # hastart
  ```

Upgrading the operating system and upgrading to 5.1 RP2

This section describes how to upgrade the operating system on a Storage Foundation node where you plan to upgrade to 5.1 RP2. This section includes the following topics

- Upgrading RHEL 5

- Upgrading OEL 5
- Upgrading SLES 10

Upgrading RHEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 RP2.

To upgrade to a later version of RHEL 5

1. Stop Storage Foundation.
2. Upgrade to the latest update version of RHEL 5.
3. Upgrade to 5.1 RP2.
4. Start Storage Foundation.

Upgrading OEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 RP2.

To upgrade to a later version of OEL 5

1. Stop Storage Foundation.
2. Upgrade to the latest update version of OEL 5.
3. Upgrade to 5.1 RP2.
4. Start Storage Foundation.

Upgrading SLES 10

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 RP2.

To upgrade to a later version of SLES 10

1. Stop Storage Foundation.
2. Upgrade to the latest update version of SLES 10.
3. Upgrade to 5.1 RP2.
4. Start Storage Foundation.

Chapter 4

Removing and rolling back

This chapter includes the following topics:

About removing Veritas Storage Foundation and High Availability Solutions 5.1 RP2

Roll back of the 5.1 RP2 to the 5.1 release is not supported. Symantec recommends that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the 5.1 release software.

For extensive uninstallation and reinstallation procedures, refer to the appropriate product's 5.1 Installation Guide.

- Uninstalling Veritas Storage Foundation Cluster File System 5.1 RP2
- Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 RP2

Uninstalling Veritas Storage Foundation Cluster File System 5.1 RP2

This section provides procedures for uninstalling Veritas Storage Foundation Cluster File System (SFCFS). You must complete the preparatory tasks before you uninstall SFCFS.

Preparing to uninstall Veritas Storage Foundation Cluster File System

The following procedure prepares your system for uninstalling Veritas Storage Foundation Cluster File System (SFCFS).

To prepare to uninstall Veritas Storage Foundation Cluster File System

1 Log in as the root user on any node in the cluster.

2 Verify that the following directories are set in your PATH environment variable:

```
/opt/VRTS/bin
/opt/VRTSvcs/bin
```

3 Back up the following configuration files:

```
# mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
```

4 Determine if each node's root disk is under VxVM control and proceed as follows.

- Check if each node's root disk is under VxVM control:

```
# df -v /
```

The root disk is under VxVM control if /dev/vx/dsk/rootvol is listed as being mounted as the root (/) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the vxplex command to remove all the plexes of the volumes rootvol, swapvol, usr, var, opt and home that are on disks other than the root disk. For example, the following command removes the plexes mirrootvol-01, and mirswapvol-01 that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

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

  Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `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
  ```

6 Check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

7 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint1
# umount /filesystem1
```

If file system is mounted in a cluster, then use `cfsumount` command.

8. Stop all VxVM volumes by entering the following command for each disk group:

   ```
   # vxvol -g dg1 stopall
   ```

   To verify that no volumes are open:

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

9. Stop VCS:

   ```
   # hastop -all
   ```

Uninstalling Veritas Storage Foundation Cluster File System

The following uninstalls Veritas Storage Foundation Cluster File System (SFCFS).

To uninstall Veritas Storage Foundation Cluster File System

1. Log in as the root user on any node in the cluster.
2. Navigate to the directory that contains the uninstallation program:

   ```
   # cd /opt/VRTS/install
   ```

3. Run the uninstallation program while specifying each node in the cluster:

   ```
   # ./uninstallsfcfs system1
           system2
   ```

4. Confirm the uninstallation:

   ```
   Are you sure you want to uninstall SFCFS [y,n,q] (y)
   ```

   The installer stops the Veritas Storage Foundation Cluster File System processes and uninstalls the packages.

5. Reboot the nodes:

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

After uninstalling the Veritas Storage Foundation Cluster File System, refer to the *Veritas Storage Foundation Cluster File System 5.1 Installation Guide* to reinstall the 5.1 software.

Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 RP2

To prepare to uninstall Veritas Storage Foundation Cluster File System for Oracle RAC from a cluster

1 Log in as the root user on any node in the cluster.

2 Verify that the following directories are set in your PATH environment variable in order to execute the necessary commands:

```
/opt/VRTS/bin
/opt/VRTSvcs/bin
```

3 Back up the following configuration files:

```
# mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
```

4 On all the nodes, stop the CFS-dependant applications that are not under VCS control using application specific commands.

For example, to stop Oracle Clusterware:

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

5 Stop VCS:

```
# hastop -all
```

6 Verify that port h is not open:

```
# gabconfig -a
```

7 Check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

8 Unmount all file systems:

```
# umount /filesystem
```

9 Stop all VxVM volumes by entering the following command for each disk group:

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

To verify that no volumes are open:

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

Perform the steps in the following procedure to uninstall Veritas Storage Foundation Cluster File System for Oracle RAC from a cluster.

To uninstall Veritas Storage Foundation Cluster File System for Oracle RAC from a cluster

1 Log in as the root user on any node in the cluster.

2 Navigate to the directory that contains the uninstallation program:

```
# cd /opt/VRTS/install
```

3 Start the uninstallation program:

```
# ./uninstallsfcfsrac node1
       node2 ... nodeN
```

4 Press Enter to uninstall Veritas Storage Foundation Cluster File System for Oracle RAC.

```
Do you want to uninstall SFCFSRAC from these systems [y,n,q] (y)
```

The installer checks the RPMs installed on the system.

5 Confirm the uninstallation at the following prompt:

```
Are you sure you want to uninstall SFCFSRAC [y,n,q] (y)
```

The installer stops the Veritas Storage Foundation Cluster File System for Oracle RAC processes and uninstalls the packages.

6 Reboot the nodes:

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

After uninstalling the Veritas Storage Foundation Cluster File System for Oracle RAC, refer to the *Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 Installation and Configuration Guide* to reinstall the 5.1 software.

Chapter 5

Verifying software versions

This chapter includes the following topics:

- Verifying software versions

Verifying software versions

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

```
# rpm -qa | egrep VRTS
```