Veritas Storage Foundation and High Availability Solutions Release Notes

Solaris

5.1 Rolling Patch 2

Symantec

Storage Foundation and High Availability Solutions Release Notes 5.1 Rolling Patch 2

Product version: 5.1 RP2

Document version: 5.1RP2.0

Legal Notice

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/334829

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.
`[ -jumpstart jumpstart_path`	The `-jumpstart` option is used to generate finish scripts which can be used by Solaris JumpStart Server for automated installation of all packages and patches for every product, an available location to store the finish scripts should be specified as a complete path. The `-jumpstart` option is supported on Solaris only.
`[ -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.

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

Command Line Option	Function
`[ -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`.
`[ -rootpath root_path ]`	The `-rootpath` option is used to re-root the install of all packages to the given path. On Solaris, -rootpath passes -R <root_path> to pkgadd.
`[ -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.

Enhancements to the Zone agent

- Enhanced the Zone agent to run fsck on VxFS file system configured as part of Zone boot configuration file.
  Added attribute RunFsck, which if enabled zone agent scans **/etc/zones/*zone_name*.xml** file for any VxFS file systems configured and runs fsck before booting the Zone. By default this attribute is disabled.
- Enhanced the Zone agent to support Live Upgrade in Solaris non-global zone environments.
  Added attribute DetachZonePath, which if disabled zone agent skips detaching the zone root leaving zone in installed state that enables Live Upgrade to succeed. By default this attribute is enabled.

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.

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 a SFCFS or SF Oracle RAC 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 50.

System requirements

This section describes the system requirements for this release

Supported Solaris operating systems

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

- Solaris 9 9/05 Release (SPARC 64-bit and 32-bit) with Update 7, 8 and 9
- Solaris 10 (64 bit, SPARC and x86) with Update 6, 7 and 8

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

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

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

http://entsupport.symantec.com/docs/334829

Required Solaris patches

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

See http://sunsolve.sun.com for the latest Solaris patch updates.

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

Table 1-2 Solaris SPARC patches

Operating system	Oracle patch number
Solaris 9	114477-04 122300-29 (required for Live Upgrade) Patches required for FS, SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC: ■ For Solaris 9 Update 7 (to bring OS kernel to FS patch 122300 level): 117171-17 113073-14 ■ For Solaris 9 Update 7 & Update 9: 112908-33 (required for 122300-10) 118558-39 (required for 122300-10) 122300-10
Solaris 10	119254-06 119042-02 125731-02 128306-05 127111-01

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

Table 1-3 Solaris x64 patches

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

Database requirements

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

http://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.

For latest information on support for Oracle database versions with SF Oracle RAC, see the Technical Support TechNote:

http://entsupport.symantec.com/docs/280186

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 for Oracle RAC (SF Oracle RAC)
- 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 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 Storage Foundation: Issues fixed in 5.1 RP2

Table 1-4 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.
1976928	dbed_clonedb of offline checkpoint fails with ORA-00600

Veritas Storage Foundation: Issues fixed in 5.1 RP1

Table 1-5 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 1-6 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
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.
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.

Table 1-6 Veritas File System fixed issues in 5.1 RP2 *(continued)*

Fixed issues	Description
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
2084071	Fixed an issue in fcladm(1m) where it used to generate core when no savefile was specified
2072165	Fixed an active level leak issue while fsadm resize operation.
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 1-7 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number)

Fixed issues	Description
1979429	mount usage error message is thrown while using from /opt/VRTS/bin
1978029	fsadm -R returns EFAULT in certain scenarios.
1976287	conform.backup test is skipped even if no local zone is available.
1973739	vxdisk reclaim fails on dg version 140 works only with ver 150-Cannot reclaim space on dg/vols created on 5.0MP3 after 5.1 upgrade.
1973539	Wrong boundary for reclamation on Hitachi AMS2000 Series
1972882	use separate structures for ioctl interfaces and CFS messages
1972207	full fsck is very slow on an fs with many ilist holes.
1969334	'vxupgrade' test failed.
1967027	LM-conform test odm hits an assert of "f:fdd_advreload:2"
1961790	LM.CMDS->fsck->full->scripts->fextop_12 fails

Table 1-7 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) *(continued)*

Fixed issues	Description
1960436	CFS.Comform->revnlookup hit "vx_msgprint" via "vx_cfs_iread" on the slave node
1958198	cfs odm stress/noise tests failed due to "bcmp error"
1957365	CFS-Conformance test failed
1957296	CFS-Conformance-Reconfig test hit assert "f:vx_validate_cistat:3"
1957043	LM -conformance/fcl/fcl_fsetquota.3 is failing
1957035	CFS cmds:fsck is failing
1957032	fsqa lm vxmssnap.9 test fails
1956926	cfs-cmds aborting due to fsck, mount, fsted, libtst 64-bit binaries
1954897	CFS-Conformance test hit assert "f:vx_mark_fset_clean:2"
1953913	LM-Command "vxedquota" test failed.
1952827	LM / CFS - Cmds-> alerts test failed.
1952818	LM / CFS -Cmds-> vxtunefs test failed.
1949962	Fix the vxrsh and vxproxyrshd processes for cfs reconfig testing.
1949077	cfs.conform.dbed hits assert "..f:vx_imap_process_inode:4a". by "..vx_workitem_process".
1948451	kernel-conform "sunppriv" and "getattr" tests are missing
1947359	mkdstfs fails to add new volumes
1947356, 1883938	Due to incorrect Makefile 'make clobber' is removing mkdstfs
1946442	tot build setup machine running LM-cmds -> fsppadm got failures.
1946433	Enhance mkdstfs for explicitly selecting the purpose of volumes data vs metadata
1946431	mkdstfs only adds to existing volume tags
1946134	fsadm keeps relocating and copying already relocated and copied reorg-ed regions of a file in subsequent passes.

Table 1-7 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) *(continued)*

Fixed issues	Description
1944283	fcl close is not happening properly
1943116	mkdstfs uses wrong perl instead of /opt/VRTS/bin/perl
1943116	mkdstfs uses wrong perl instead of /opt/VRTS/bin/perl
1940870	Documentation/Test discrepancies
1940409	CFS support for cached ODM
1940390	Cached ODM needs improvements for async requests codm
1934107, 1891400	Incorrect ACL inheritance
1934103	[PRI-1] Performance issues with mmap VxFS 4.1MP4 RP2
1934101	cfs Test cfs-stress-enterprise hit the same assert :f:vx_cwfrz_wait:2
1934098, 1860701	clone removal can block resive ops
1934096, 1746491	core dump of fsvmap.
1934095, 1838468	Data page fault at vx_qiostats_update due to fiostats structure already free'd
1934094, 1846461	vxfsstat metrics to monitor UNHASHED entries in the dcache
1934085, 1871935	secondaries ias_ilist not updated fully.
1933975, 1844833	fsadm shrink fs looping in vx_reorg_emap due to VX_EBMAPMAX from vx_reorg_enter_zfod
1933844	bad mutex panic in VxFS
1933635, 1914625	[VxFS] Behavior of DST Access age-based file placement policy with preffered files
1931973	fsppadm gives spurious messages when run fron multiple CFS nodes found only from 5.1 onwards
1908776	UX:vxfs mount: ERROR: V-3-22168: Cannot open portal device...

Table 1-7 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) *(continued)*

Fixed issues	Description
1906521	CFS-conform/quotas test hit assert vx_populate_pnq via vx_detach_fset
1902241	9-15a driver regression observed on SFCFSORA TPCC test
1897458, 1805046	wrong alert generation from vxfs when file system usage threshold is set
1895454	Sol10x86 lm.conform->ts some TCs fail
1878583	CFS: getattr call optimization to speedup the case when binaries are being mmapped from many nodes on CFS.

Veritas Volume Manager fixed issues

This section lists fixed issues for Veritas Volume Manger.

Veritas Volume Manager: Issues fixed in 5.1 RP2

Table 1-8 Veritas Volume Manager 5.1 RP2 fixed issues

Fixed issues	Description
1831969	VxVM: ddl log files are created with world write permission.
1871447	Mirrored encapsulated disk panics on boot when the primary is removed & mpxio is enabled.
1874034	Race between modunload and an incoming IO leading to panic.
1880279	Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment.
1882789	Enabling a path to the EMC lun in LDOM environment fails.
1899943	CPS based fencing disks used along with CPS servers does not have coordinator flag set.
1901847	ISCSI devices not visible to VxVM after boot since VxVM discovers devices before ISCSI device discovery.
1911546	Vxrecover hung with layered volumes.

Table 1-8 Veritas Volume Manager 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
1920761	I/O hang observed after connecting the storage back to master node incase of local detach policy.
1923906	CVM: Master should not initiate detaches while leaving the cluster due to complete storage failure.
1929074	vxbootsetup not processing volumes in an ordered manner.
1929083	Vxattachd fails to reattach site in absence of vxnotify events.
1932023	vxdiskadm option 'Allow multipathing of all disks on a controller by VxVM' fails due to script errors.
1933375	Tunable value of 'voliomem_chunk_size' is not aligned to page-size granularity.
1933477	Disk group creation on EFI disks succeeds but with error messages.
1933528	During Dynamic reconfiguration vxvm disk ends up in error state after replacing physical LUN.
1936611	vxconfigd core dump while splitting a diskgroup.
1938708	With EBN naming, root (un)encapsulation is not handling dump/swap device properly.
1942985	Improve locking mechanism while updating mediatype on vxvm objects.
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.
1946941	vxsnap print shows incorrect year.
1952177	Machine panics after creating RVG.
1956777	CVR: Cluster reconfiguration in primary site caused master node to panic due to queue corruption.
1960341	Toggling of naming scheme is not properly updating the daname in the vxvm records.
1972755	TP/ETERNUS:No reclaim seen with Stripe-Mirror volume.
1974393	Avoiding cluster hang when the transaction client timed out.

Table 1-8 Veritas Volume Manager 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
1982715	vxclustadm dumping core while memory re-allocation.
1983768	IO hung on linked volumes while carrying out third mirror breakoff operation.
1989662	/opt/VRTSsfmh/bin/vxlist causes panic.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout.
1992872	Vxresize fails after DLE.
1996162	Bootdg not reset after unencapsulation.
1998447	Vxconfigd dumped core due to incorrect handling of signal.
1999004	I/Os hang in VxVM on linked-based snapshot.
2006454	AxRT5.1P1: vxsnap prepare is displaying vague error message.
2010426	Tag setting and removal do not handle wrong enclosure name.
2011316	VVR: After rebooting 4 nodes and try recovering RVG will panic all the slave nodes.
2012016	Slave node panics while vxrecovery is in progress on master.
2015570	File System read failure seen on space optimized snapshot after cache recovery.
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1.
2029480	Diskgroup join failure renders source diskgroup into inconsistent state.
2029735	System panic while trying to create snapshot.
2031462	Node idle events are generated every second for idle paths controlled by Third Party drivers.
2034564	I/Os hung in serialization after one of the disk which formed the raid5 volume was pulled out.
2036929	renaming a volume with link object attached causes inconsistencies in the disk group configuration.

Table 1-8 Veritas Volume Manager 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.
2040150	Existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration.
2049952	Vxrootadm shows incorrect messages in Japanese on VxVM5.1RP1 with L10N for Solaris.
2052459	CFS mount failed on slave node due to registration failure on one of the paths.
2053975	Snapback operation panicked the system.
2055609	Allocation specifications not being propagated for DCO during a grow operation.
2059046	FMR:TP: snap vol data gets corrupted if vxdisk reclaim is run while sync is in progress.
2060974	vxrootadm sets wrong prtvtoc on VxVM5.1RP1.
2061066	vxisforeign command fails on internal cciss devices.
2065669	After upgrading to 5.1, reinitalizing the disk makes public region size smaller than the actual size.
2078111	When the IOs are large and need to be split, DRL for linked volumes cause I/Os to hang.
2113831	vxconfigd core dumps while including the previously excluded controller.
2126731	VxVM 5.1: vxdisk -p list output is not consistent with previous versions.

Veritas Volume Manager: Issues fixed in 5.1 RP1

Table 1-9 Veritas Volume Manager 5.1 RP1 fixed issues

Fixed issues	Description
1972852, 1972848	vxconfigd dumped core in dg_config_compare() while upgrading to 5.1.

Table 1-9 Veritas Volume Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1955693	VxVM 5.0MP3RP3 patch 122058-13 disables vxfsldlic service and prevents boot multi-user mode after jumpstart
1938484	EFI: Prevent multipathing don't work for EFI disk
1937841	VxVM: checkin the fmrshowmap utility
1915356	I/O stuck in vxvm caused cluster node panic
1935297	vxconfigd dumps core in get_prop()
1907796	Corrupted Blocks in Oracle after Dynamic LUN expansion and vxconfigd core dump
1901827	vxdg move failed silently and drops disks.
1899688	[VVR] Every I/O on smartsync enabled volume under VVR leaks memory
1889747	vxlustart customer is unable to do live upgrade with Solaris Zone on vxfs
1886007	vxconfigd lose license information, vxesd leaking File descriptors
1884070	When running iotest on volume, primary node runs out of memory
1881336	VVR: Primary Panic in vol_ru_replica_sent()
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1870049	Dump device changed to none after boot disk encapsulation
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
1850166	vxvm vxdisk error v-5-1-8643 device 0_bpcs001_fra: resize failed:
1846165	Data corruption seen on cdsdisks on Solaris-x86 in several customer cases
1840832	vxrootadm does not update the partition table while doing a grow operation
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
1834848	TP:Solaris:reclamation causes data corruption

Table 1-9 Veritas Volume Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1826088	After pulling out FC cables of local site array, plex became DETACHED/ACTIVE
1825516	Unable to initialize and use ramdisk for VxVM use
1825270	Need for dmp_revive_paths(in dmp reconfiguration/restore_demon code path.
1792795	supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
1766452	VVR: VRAS: AIX: vradmind dumps core during collection of memory stats.
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1528160	An ioctl interrupted with EINTR causes frequent vxconfigd exit()'s on 4.1MP4RP3
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 1-10 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`

Table 1-10 Veritas Storage Foundation Cluster File System 5.1 RP2 fixed issues (listed incident number, parent number) *(continued)*

Fixed issues	Description
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 1-11 Veritas Storage Foundation Cluster File System 5.1 RP1 fixed issues (listed incident number, parent number)

Fixed issues	Description
1980842, 1983222	Fixed issue in cfsadmin command for RVG volumes.
1878583, 1544221	getattr call optimization to speedup the case when binaries are being mmapped from many nodes on CFS.

Veritas Storage Foundation for Oracle RAC fixed issues

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

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

There are no fixed issues for Veritas Storage Foundation for Oracle RAC in 5.1 RP2.

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

Table 1-12 Veritas Storage Foundation for Oracle RAC 5.1 RP1 fixed issues

Fixed issues	Description
1980842, 1983222	Fixed issue in cfsadmin command for RVG volumes.
1938797	LMX should register with NULL canput for performance.
1934892	Issue: PrivNIC Agent support for Sun 10GbE NICs (nxge interfaces) with native 64k MTU default value. Resolution: Change the mtu of the active link (where the ip address is currently plumbed). Then change the mtu of the other links (other than active link) to high value and bring the active link down.
1908916	Issue: Panic lmx buffer modified after being freed. Resolution: Fix the manipulation of the work queue tail pointer/done queue tail pointer whenever the request is removed.
1853839	Issue: MultiPrivNIC resource state change to UNKNOWN once member node shutdown Resolution: The sum of the number nodes that are visible from all the devices would be zero if there is no valid LLT device. The code has been changed to handle this case.

Veritas Cluster Server fixed issues

This section lists fixed issues for Veritas Cluster Server.

Veritas Cluster Server: Issues fixed in 5.1 RP2

Table 1-13 Veritas Cluster Server 5.1 RP2 fixed issues

Fixed issues	Description
1937834	Fixed issues with partial parent group coming online after the child service group autostarts and issue in service group failover after node panic.
1952087	Fixed issue with IPMultiNICB resource going offline after mpathd failback

Table 1-13 Veritas Cluster Server 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
1967955	Fixed issue with MultiNICB agent setting base interface DEPRECATED and NOFAILOVER flags even if CheckConfig is set to 0.
1969999	Fixed issue with SNMP Traps receiver showing trimmed output
1970639	Fixed issue with Mount Agent in handling loopback mounts
1979662	Fixed issue with 'hastatus -sum' command for a non-root user with FSS
1980252	Fixed issue with memory leak during unloading in LLT
1982648	Enhanced NFS agent to clear /etc/rmtab before starting mountd
1986311	Fixed issue with accessing a freed pointer in vxfend
2009127	Fixed message ID errors in Samba agents
2011416	Fixed Memory leaks in llt_dlpi_setup and llt_closenode paths
2016490	Fixed issue with MultiNICB agent in showing Unknown status after changing existing NIC Device
2017344	Fixed MultiNICB agent issue with haping command for IPv6 addresses
2019202	Fixed issue with MultiNICB agent checking for interface out of its control which is configured in same subnet
2019899	Fixed issue with notifier in binding to specific source IP
2021018	Add new attribute to support fsck during zone boot for vxfs file systems for Zone agent
2025380	Fixed issue with Share agent which brings resource offline when had is restarted by hashadow.
2031922	Fixed dlpiping utility errors working in server mode
2031931	Fixed issue with 'lltconfig –a set 1 link1 <ip-address>' command
2035239	Fixed issue with Application agent flooding engine log with "Container is not up" messages
2059631	Fixed issue with llt links going off/on repeatedly after configuring llt over udp in ipv6
2064973	Fixed issue with Application agent in deleting an uninitialized pointer during clean

Table 1-13 Veritas Cluster Server 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
2066967	Improved heartbeating logic b/w engine and agfw during snapshotting and during steady cluster operation
2066986	Fixed issue with had getting killed in a heavily loaded system
2071549	Fixed issue with had getting killed by GAB when user runs ha command on heavily loaded system
2077022	Fixed unwanted Solaris patch messages logging into engine log while zone resource goes online
2077363	Fixed issue with duplicate entries in main.cf when a static attribute for a resource with empty type definition is overridden
2077375	Fixed issue with parallel group in PARTIAL state autostarting when engine restarts
2077396	Fixed VCS core dump issues under heavy load system which results in node panic
2078802	Fixed issue with halogin failures during zone resource online for branded zones.
2079596	Fixed issue with zone resource taking long time to go down when the zone root file system is forcefully unmounted
2079601	Fixed issue with Zone agent not faulting when its zone root no longer exists
2079617	Fixed issue with NFSRestart agent in remaining online when HAD is forcefully stopped and restarted.
2079664	Fixed issue with ClusterService Group autopopulating system list
2079854	Added fencing support for non SCSI-3 capable environments
2079868	Fixed mismatch in displaying ContainerInfo at group level & resource level
2079977	Fixed Zone agent to remove UserNames added during online to allow successful _had snapshot of remote node(s)
2080703	Fixed Memory leak in 'hareg -group' command
2081674	Fixed problem with 'hastatus -sum' command in branded zone
2081741	Fixed hastart command issue which fails with error message GabHandle::open failed errno = 16

Table 1-13 Veritas Cluster Server 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
2083268	Enhanced DNS Agent to disable RFC check
2084977	Fixed stack overflow issue in hacf when a very large filename types definition file is included in main.cf
2086251	Enhanced GAB driver to prevent 'gabconfig -c' while port 'a' is in the middle of iofence processing which could cause system panic.
2086273	Fixed issue with gabfd static array which contains stale entry if open succeeds and registration fails
2086280	Fixed node panics after running the GTX tc.
2089182	Fixed issue with vxcpserv where memory keeps growing when cpsadm commands are executed
2089193	Fixed issue with CP Server which gets killed when CP Client of higher version attempts to communicate with CPServer running at 5.1 version
2092199	Fixed a packaging issue in VCS where shuting down a branded zone fails.
2092201	Fixed Zone agent not to run "zlogin halogin" command if the user is already created or if /.vcspwd file is already present.
2096212	Fixed Netlsnr agent issue inside a zone where the agent fails to monitor listener if the oracle home directory is NFS/NAS mounted inside the zone
2097029	Enhanced Zone agent not to detach zone root by adding attribute DetachZonePath which defaults to 1
2102776	Enhanced Zone agent to handle Sun bug 6757506 where zoneadm list fails causing Zone resource to fault
2102777	Fixed RemoteGroup resource to support remote groups that are a parallel SG.
2083268	Disabled RFC check with DNS Agent.
2111294	Fixed problem in which some VCS nodes in a group are online and some are not.
2116161	Fixed problem that prevented LVMVG from going online.
2119570	Fixed problem that prevented Oracle detailed monitoring from sending SNMP notification on WARN signal.

Table 1-13 Veritas Cluster Server 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
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.
2126871	Fixed a memory leak issue in MultiNICB agent.
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
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 1-14 Veritas Cluster Server 5.1 RP1 fixed issues

Fixed issues	Description
1975424	[IPMultiNICB][410-647-713][AIG] In a zone, if there are multiple IPMultiNICB resources on same subnet, issues with source address of IP pkts.
1972789	[VCS ASMInstAgent]ASMInstAgent Monitor problem inside a zone where the oracle home directory is NFS / NAS mounted inside the zone
1972770	[VCSOR]ASMInstAgent does not detect the state of the instance correctly inside local zones.
1968572	Postpatch script throws an error while installing in the non-global zone though llt and gab are hollow packages
1962548	ASMInstAgent dumpipng core.
1954723	[VCS Oracle Agent] [410-989-573] Oracle Agent Monitor problem inside a zone where the oracle home directory is NFS / NAS mounted inside the zone
1950427	[VCSOR] ASMDGAgent should disable and enable diskgroups in offline and online EPs for 11g R2.
1941647	haalert CLI hangs if engine is not in running state.
1922411	vxfentsthdw should detect storage arrays which interpret NULL keys as valid for registrations/reservations
1916022	[VCSOR][240-998-619] Changes made to Oracle agent via e1722109 do not honour ContainerName attribute
1916004	ASMagent connecting as sysdba instead of sysasm for 11gR2
1915909	[VCS][281-889-442] hares allows to create resources which has "." special character
1911287	group stuck at OFFLINE|STOPPING state when there is no ip to be cleaned in IPMultiNICB.
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
1848114	SxRT5.1:Oakmont:IPMultiNICB:Resource does not come online when failover from paniced node

Table 1-14 Veritas Cluster Server 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1504123	[SFW-HA 5.1 GCO] Symantec SE - GCO failover does not work when user account has "!" in name.

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 1-15 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
1893699	Unable to add a host to the management server. V-39-4095-903 401 Unauthorized User Error

Table 1-15 Storage Foundation Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
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
1831711	Volume Migration fails because it cannot find a target enclosure

Table 1-15 Storage Foundation Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
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 1-16 VEA 5.1 RP1 fixed issues

Fixed issues	Description
1961540	vmprov does not calculate disk nolabel state correctly.
1961519	vxsvc running as a daemon shows stderr and stdout printf's

Table 1-16 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.

- 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 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.

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.

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

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

Workaround: Set the environment variable LANG=C systemwide, for example, by setting the environment variable in /etc/profile.

Post Upgrade:

Note: These post upgrade steps are relevant only if CC storage server is installed on the machine.

---

1. If you are using dbed, run sfua_rept_migrate without the -r option. This ensures the `VRTSdbms3` package will not be uninstalled.
2. Copy the /sbin/rc3.d/NO_S*vxdbms3 to /sbin/rc3.d/S*vxdbms3

Veritas Volume Manager known issues in 5.1 RP2 release

Changing naming scheme fails (1958711)

Changing naming scheme fails for devices controlled by MPxIO driver on Solaris.

There is no workaround at this time.

EMC CLARiiON MirrorView read-only devices are reported as "error" by vxdisk list (2054201)

When running these commands:

```
vxconfigd
vxdctl enable
vxdisk online
```

the following SCSI errors occur on non-readable EMC MirrorView secondary devices:

```
May 21 00:02:37 thor393 scsi: WARNING:
/pci@1d,700000/fibre-channel@1/fp@0,0/ssd@w500601601060249a,0 (ssd0):
May 21 00:02:37 thor393      Error for Command: read     Error
Level: Informational
May 21 00:02:37 thor393 scsi:Requested Block: 0          Error
Block: 0
May 21 00:02:37 thor393 scsi:Vendor: DGC
Serial Number: 680000C856CL
May 21 00:02:37 thor393 scsi:Sense Key: Illegal Request
May 21 00:02:37 thor393 scsi:ASC: 0x25 ([vendor unique code 0x25]), ASCQ:
```

There is no workaround at this time.

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.

Reclamation is not triggered at the user specified time (2076423)

The free space reclamation does not happen at the time specified by *reclaim_on_delete_start_time*.

Workaround: Start the reclamation manually.

Missing patch error occurs during installation when the SUNWcfcl package is absent (2138029)

On Solaris SPARC 9, if you try to install patch 142629-05, and if the package `SUNWcfcl` and the patch 114477-04 are missing, an error will occur and terminate the installation.

Workarond: Install the package `SUNWcfcl` and the patch 114477-04 (or higher version), and then re-install the patch 142629-05.

Veritas File System known issues in 5.1 RP2 release

The following are new additional Veritas File System known issues in this 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.

Installation of VRTSvxfs 5.1RP2 patch on zones may fail with the following error messages (2086894)]

package `VRTSvxfs` failed to install - interrupted:

pkgadd: ERROR: duplicate pathname ***zone_path*/root/etc/fs/vxfs/qioadmin**

pkgadd: ERROR: duplicate pathname ***zone_path*/root/etc/vx/cdslimitstab**

pkgadd: ERROR: duplicate pathname ***zone_path*/root/kernel/drv/vxportal.conf**

pkgadd: ERROR: duplicate pathname***zone_path*/root/opt/VRTSvxfs/etc/access_age_based.xml**

pkgadd: ERROR: duplicate pathname*zone_path/root/opt/VRTSvxfs/etc/access_age_based_2tier.xml*

...

Workaround: Remove 5.1RP1 package Re-install the 5.1 Install 5.1RP2

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.

There is no workaround at this time.

installrp recognizes SFCFSHA 5.1RP1 as SFCFS after installing SFCFSHA 5.1RP1 using JumpStart (1991079)

The installrp script recognizes SFCFSHA 5.1RP1 as SFCFS after installing SFCFSHA 5.1RP1 using JumpStart.

Workaround:

After you finish running JumpStart, use the `installcfs -license -ha` command to license SFCFSHA. Run installrp to finish the configuration.

Mounting a file system as seconly by using the cfsmount command may fail (2104499)

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 for Oracle RAC known issues in 5.1 RP2

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

Message about mmpl_reconfig_ioctl in system log

If the Veritas Cluster Server Membership Module (VCSMM) calls the mmpl_reconfig_ioctl function of the fencing module (VxFEN) at the time of system startup, the call fails displaying the following error message on the console and the `/var/adm/messages` file:

```
mmpl_reconfig_ioctl: dev_ioctl failed, vxfen may not be configured
```

You may ignore this message.

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.

Live Upgrade may fail while removing VRTSvcsea package from the alternate disk (2100295)

While removing the VRTSvcsea package from an alternate disk, the package removal can fail. The failure occurs due to the pre-remove script of VRTSvcsea, when it checks if any of Oracle, Netlsnr, ASMInst, ASMDG, Db2udb, Sybase or SybaseBk agents are running. So even if agents are running from the first disk, you cannot remove the package from the second disk.

Workaround:

Stop the agent if any of the agents Oracle, Netlsnr, ASMInst, ASMDG, Db2udb, Sybase or SybaseBk running on the system before live upgrade.

Use the following command to stop agent forcibly so that the application does not need to go offline during the upgrade.

```
# haagent -stop Agent_Name -force -sys system_name
```

CFS nfs filesystem (setup by cfsshare command) isn't accessible after VIP switch to another node in cluster (2117482)

On solaris 9, CFS nfs filesystem (setup by cfsshare command) is not accessible after VIP swith to another node.

patchadd for 143268-07 may output unwanted message in local zones

patchadd for 143268-07 may output unwanted messages on the screen while installing in the local zones as shown below.

`svcadm`: Pattern 'system/llt' doesn't match any instances

`svcadm`: Pattern 'system/gab' doesn't match any instances

Workaround: These messages can be ignored.

patchrm for 143268-07 may fail to back out in global zone

While removing the patch 143268-07 from Solaris 10 x64 with zones the patch may fail to back out with the following message.

"Pkgadd failed. See /var/tmp/143268-07.log for details Patchrm is terminating. WARNING: patchrm returned [7] for global zone"

Workaround: These messages can be ignored.

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 for Oracle RAC 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

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
```

Veritas Storage Foundation for Oracle RAC software limitations in 5.1 RP2 release

CRSResource agent

CRSResource agent is not supported for Oracle 11g Release 2.

Documentation addendum

The following sections contain additions to current documents.

Disk agent

Monitors a physical disk or a partition. You can use the Disk agent to monitor a physical disk or a slice that is exported to LDoms (available using LDoms 1.2 or later). For LDoms with a physical disk or slice based boot image, a dependency must exist between the guest domain and primary domain. You configure the primary domain as the master of the guest domain.

Perform the following:

- Set the failure-policy of primary (control) domain to stop. For example, in the primary domain enter the following command to set the dependent domain to stop when the primary domain faults:

  ```
  # ldm set-domain failure-policy=stop primary
  ```

- Set the primary domain as the master for the guest domain.

  ```
  # ldm set-domain master=primary guestldom
  ```

Figure 1-1 Sample service group that includes a Disk resource on Solaris

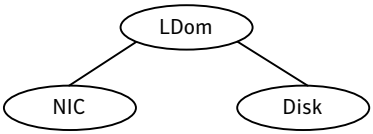

Using the preonline_vvr trigger for RVGLogowner resources

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.

Attribute

The Disk agent has the following required attribute.

Partition—Indicates which partition to monitor. Specify the partition with the full path beginning with a slash (/).

If this path is not specified, the name is assumed to reside in `/dev/rdsk/`.

Example: "`/dev/rdsk/c2t0d0s2`"

Type and dimension: string-scalar

NotifierSourceIP—Provide this IP if notifier needs to send packet from a particular IP.

Type and dimension: string-scalar

Resource type definition

The following is the agent's resource type definition.

```
type Disk (
static int OfflineMonitorInterval = 60
static str ArgList[] = { Partition }
static str Operations = None
str Partition
)
```

List of patches

This section lists the patches and packages for 5.1 RP2.

Table 1-17 Patches and packages for Solaris 9 on SPARC

Patch ID	5.1 Package Names	Products Affected	Patch Size
143260-03	VRTSllt	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	2148K

Table 1-17 Patches and packages for Solaris 9 on SPARC *(continued)*

Patch ID	5.1 Package Names	Products Affected	Patch Size
143262-03	VRTSgab	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	2496K
143706-03	VRTSvxfen	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	3104K
143264-07	VRTSvcs	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	32M
143265-07	VRTSvcsag	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	800K
143279-03	VRTScps	VCS, SFHA, SFCFSHA, SFCFS RAC	16M
143276-04	VRTSvcsea	VCS, SFHA, SFCFSHA, SFCFS RAC	15.64M
142629-05	VRTSvxvm	SFHA, SFCFS, SFCFSHA, SFCFS RAC	178M
141270-02	VRTSsfmh	SFHA, SFCFS, SFCFSHA, SFCFS RAC	25.72M
142633-03	VRTSvxfs	SFHA, SFCFS, SFCFSHA, SFCFS RAC	35.92M
142631-03	VRTSdbed	SFHA, SFCFS, SFCFSHA, SFCFS RAC	45.2M

Table 1-17 Patches and packages for Solaris 9 on SPARC *(continued)*

Patch ID	5.1 Package Names	Products Affected	Patch Size
143270-03	VRTSodm	SFHA, SFCFS, SFCFSHA, SFCFS RAC	1.45M
143273-02	VRTScavf	SFCFS, SFCFSHA, SFCFS RAC	1028K
143696-01	VRTSdbac	SFCFS RAC	7.4M
143687-01	VRTSob		457K

Table 1-18 Patches and packages for Solaris 10 on SPARC

Patch ID	Package Name	Products Affected	Patch Size
143261-03	VRTSllt	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	1936K
143263-03	VRTSgab	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	2028K
143707-03	VRTSvxfen	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	2564K
143264-07	VRTSvcs	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	115M
143265-07	VRTSvcsag	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	3936K

Table 1-18 Patches and packages for Solaris 10 on SPARC *(continued)*

Patch ID	Package Name	Products Affected	Patch Size
143279-03	VRTScps	VCS, SFHA, SFCFSHA, SFCFS RAC	40M
143276-04	VRTSvcsea	VCS, SFHA, SFCFSHA, SFCFS RAC	4.4M
142629-05	VRTSvxvm	SFHA, SFCFS, SFCFSHA, SFCFS RAC	178M
141270-02	VRTSsfmh	SFHA, SFCFS, SFCFSHA, SFCFS RAC	25.72M
142634-03	VRTSvxfs	SFHA, SFCFS, SFCFSHA, SFCFS RAC	46.45M
142631-03	VRTSdbed	SFHA, SFCFS, SFCFSHA, SFCFS RAC	45.2M
143271-03	VRTSodm	SFHA, SFCFS, SFCFSHA, SFCFS RAC	1.30M
143274-02	VRTScavf	SFCFS, SFCFSHA, SFCFS RAC	1028K
143697-01	VRTSdbac	SFCFS RAC	6.17M

Table 1-19 Patches and packages for Solaris 10 on x64

Patch ID	Package Name	Products Affected	Patch Size
143266-03	VRTSllt	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	1800K

Table 1-19 Patches and packages for Solaris 10 on x64 *(continued)*

Patch ID	Package Name	Products Affected	Patch Size
143267-03	VRTSgab	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	1804K
143708-03	VRTSvxfen	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	2520K
143268-07	VRTSvcs	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	119M
143269-07	VRTSvcsag	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	3.87M
143280-03	VRTScps	VCS, SFHA, SFCFS, SFCFSHA, SFCFS RAC	40.4M
143277-03	VRTSvcsea	VCS, SFHA, SFCFSHA, SFCFS RAC	19.76M
142630-05	VRTSvxvm	SFHA, SFCFS, SFCFSHA, SFCFS RAC	163.21M
141752-02	VRTSsfmh	SFHA, SFCFS, SFCFSHA, SFCFS RAC	27.01M
142635-03	VRTSvxfs	SFHA, SFCFS, SFCFSHA, SFCFS RAC	28.67M

Table 1-19 Patches and packages for Solaris 10 on x64 *(continued)*

Patch ID	Package Name	Products Affected	Patch Size
142632-03	VRTSdbed	SFHA, SFCFS, SFCFSHA, SFCFS RAC	9.98M
143272-03	VRTSodm	SFHA, SFCFS, SFCFSHA, SFCFS RAC	1152K
143275-02	VRTScavf	SFCFS, SFCFSHA, SFCFS RAC	1032K
143698-01	VRTSdbac	SFCFS RAC	5.54M

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/334829

Chapter 2

Installing the products for the first time

This chapter includes the following topics:

Installing the Veritas software using the script-based installer

This section describes how to install a 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 for Oracle RAC:

```
# ./installsfrac node1 node2 ... nodeN
```

- 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 73.

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.

```
# ./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 for Oracle RAC:

```
# ./installsfrac -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.

Note: Installing SF Oracle RAC using the Web-based installer is not supported in this release.

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 73.

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

Installing the products using JumpStart

This chapter includes the following topics:

Installing the products using JumpStart

These JumpStart instructions assume a working knowledge of JumpStart. See the JumpStart documentation that came with your operating system for details on using JumpStart. Only fresh installations of Veritas product are supported using JumpStart. Upgrading is not supported. The following procedure assumes a stand-alone configuration.

For the language pack, you can use JumpStart to install packages. You add the language packages in the script, and put those files in the JumpStart server directory.

Overview of JumpStart installation tasks

Review the summary of tasks before you perform the JumpStart installation.

Summary of tasks

1. Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.
2. Read the JumpStart installation instructions.
3. Generate the finish scripts.

   See "Generating the finish scripts" on page 64.

4 Prepare installation resources.

See "Preparing installation resources" on page 68.

5 Run JumpStart to install the Veritas product.

Note: JumpStart may reboot systems after product installation.

6 Run the `installer` command from the disc or from directory `/opt/VRTS/install` directory to configure the Veritas software.

```
# /opt/VRTS/install/installer installprod -configure
```

Where *installprod* is the product's installation command.

7 Modify the rules file for JumpStart.

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

Generating the finish scripts

Perform these steps to generate the finish script to install Veritas product.

To generate the script

1 Run the `installrp` program to generate the scripts.

```
# installrp -jumpstart directory_to_generate_scripts
```

Where the *directory_to_generate_scripts* is where you want to put the scripts.

For example:

2 When you are prompted to encapsulate the root disk automatically, choose **yes** to do so. If you do not want to encapsulate it automatically, choose **no** and go to step 6.

3 Specify a disk group name for the root disk.

```
Specify the disk group name of the root disk to be encapsulated:
rootdg
```

4 Specify private region length.

```
Specify the private region length of the root disk to be
encapsulated: (65536)
```

5 Specify the disk's media name of the root disk to encapsulate.

```
Specify the disk media name of the root disk to be encapsulated:
(rootdg_01)
```

6 JumpStart finish scripts, installer scripts, uninstaller scripts of Veritas products, and encapsulation scripts are generated in the directory you specified in step 1. Output resembles:

```
The finish scripts for AT50 is generated at
/js4/jumpstart_at50.fin
The installer script to configure AT is generated at
/js4/installat
The installer script to uninstall AT is generated at
/js4/uninstallat
The finish scripts for FS51 is generated at
/js4/jumpstart_fs51.fin
The installer script to configure FS is generated at
/js4/installfs
The installer script to uninstall FS is generated at
/js4/uninstallfs
The finish scripts for SF51 is generated at
/js4/jumpstart_sf51.fin
The installer script to configure SF is generated at
/js4/installsf
The installer script to uninstall SF is generated at
/js4/uninstallsf
The finish scripts for SFCFS51 is generated at
/js4/jumpstart_sfcfs51.fin
The installer script to configure SFCFS is generated at
/js4/installsfcfs
The installer script to uninstall SFCFS is generated at
/js4/uninstallsfcfs
The finish scripts for SFCFSHA51 is generated at
/js4/jumpstart_sfcfsha51.fin
The installer script to configure SFCFSHA is generated at
/js4/installsfcfsha
The installer script to uninstall SFCFSHA is generated at
/js4/uninstallsfcfsha
The finish scripts for SFHA51 is generated at
/js4/jumpstart_sfha51.fin
The installer script to configure SFHA is generated at
/js4/installsfha
The installer script to uninstall SFHA is generated at
/js4/uninstallsfha
The finish scripts for SFRAC51 is generated at
/js4/jumpstart_sfrac51.fin
The installer script to configure SF Oracle RAC is generated at
```

```
/js4/installsfrac
The installer script to uninstall SF Oracle RAC is generated at
/js4/uninstallsfrac
The finish scripts for VCS51 is generated at
/js4/jumpstart_vcs51.fin
The installer script to configure VCS is generated at
/js4/installvcs
The installer script to uninstall VCS is generated at
/js4/uninstallvcs
The finish scripts for VM51 is generated at
/js4/jumpstart_vm51.fin
The installer script to configure VM is generated at
/js4/installvm
The installer script to uninstall VM is generated at
/js4/uninstallvm
The encapsulation boot disk script for VM is generated at
/js4/encap_bootdisk_vm51001000.fin
```

List the js_scripts directory.

```
# ls /js_scripts
```

You could select scripts according to the products. For example.

For SF:

```
encap_bootdisk_vm51001000.fin installsf jumpstart_sf51.fin uninstallsf
```

For SF Oracle RAC:

```
encap_bootdisk_vm51001000.fin installsfrac jumpstart_sfrac51.fin
uninstallsfrac
```

For SFHA:

```
encap_bootdisk_vm51001000.fin installsfha jumpstart_sfha51.fin
uninstallsfha
```

For VCS:

```
encap_bootdisk_vm51001000.fin jumpstart_vcs51.fin installvcs
uninstallvcs
```

7 Modify the JumpStart script according to your requirements. You must modify the BUILDSRC and ENCAPSRC values. Keep the values aligned with the resource location values.

```
BUILDSRC="hostname_or_ip:/path_to_pkgs_patches_scripts"
// If you don't want to encapsulate the root disk automatically
// comment out the following line.
ENCAPSRC="hostname_or_ip:/path_to_encap_script"
```

8 ■ If you want to install other products packages in the Storage Foundation and High Availability suite, then use the product-specific install command with one of the following options to get a list of packages in the order to be installed:

- `minpkgs`
- `recpkgs`
- `allpkgs`

Use the list of packages that is generated to replace the package list in the finish scripts.

■ If you want to install other products patches in the Storage Foundation and High Availability suite, then use the product-specific install command with one of the following options to get a list of patches in the order to be installed:

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

See "About the installrp script" on page 12.

See "installrp script options" on page 12.

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1 Copy the contents of the installation disc to the shared storage.

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

Note: After you copied the patches, you must uncompress them using the unzip and untar commands.

2 Generate the response file for the package and patch list that you found when you generated the finish script.

See "Generating the finish scripts" on page 64.

In this example, the packages and patches are:

- For SF:
  Packages are:
  VRTSvlic VRTSperl VRTSspt VRTSob VRTSvxvm VRTSaslapm VRTSsfmh VRTSvxfs VRTSfssdk VRTSdbed VRTSodm VRTSat
  Patches are:
  142629-05 141270-02 142634-03 142633-03 142631-03 143271-03 143270-03
- For SF Oracle RAC:
  Packages are:
  VRTSvlic VRTSperl VRTSspt VRTSob VRTSvxvm VRTSaslapm VRTSsfmh VRTSvxfs VRTSfssdk VRTSllt VRTSgab VRTSvxfen VRTSvcs VRTScps VRTSvcsag VRTScutil VRTSat VRTSvcsea VRTSdbed VRTSglm VRTScavf VRTSgms VRTSodm VRTSdbac
  Patches are:
  142629-05 141270-02 142634-03 142633-03 143261-03 143260-03 143263-03 143262-03 143707-03 143706-03 143264-07 143279-03 143265-07 143276-04 142631-03 143274-02 143273-02 143271-03 143270-03 143697-01 143696-01
- For SFCFS:
  Packages are:
  VRTSvlic VRTSperl VRTSspt VRTSob VRTSvxvm VRTSaslapm VRTSsfmh VRTSvxfs VRTSfssdk VRTSllt VRTSgab VRTSvxfen VRTSvcs VRTScps VRTSvcsag VRTScutil VRTSat VRTSvcsea VRTSdbed VRTSglm VRTScavf VRTSgms VRTSodm
  Patches are:

142629-05 141270-02 142634-03 142633-03 143261-03 143260-03 143263-03 143262-03 143707-03 143706-03 143264-07 143279-03 143265-07 143276-04 142631-03 143274-02 143273-02 143271-03 143270-03

- For VCS:
  Packages are:
  VRTSvlic VRTSperl VRTSspt VRTSllt VRTSgab VRTSvxfen VRTSvcs VRTScps VRTSvcsag VRTScutil VRTSat VRTSvcsea
  Patches are:
  143261-03 143260-03 143263-03 143262-03 143707-03 143706-03 143264-07 143279-03 143265-07 143276-04

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

3 Create the adminfile file under *BUILDSRC*/pkgs/ directory. The adminfile file's contents follow:

```
mail=
instance=overwrite
partial=nocheck
runlevel=quit
idepend=quit
rdepend=nocheck
space=quit
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=default
```

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

See "Generating the finish scripts" on page 64.

5 Modify the rules file as required.

For example:

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

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

Adding language pack information to the finish file

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

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

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

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

do
...
done
```

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

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

do
.
.
.
done
```

Chapter 4

Upgrading to 5.1 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.

Use the procedures to perform a full upgrade to 5.1 RP2 on a cluster that has VCS, SFHA, SFCFS, or SF Oracle RAC 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.

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 SF Oracle 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.

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:

Solaris 10:

```
# svcadm disable -t vxfen
```

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

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

Solaris 9:

```
# /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 115.

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:

```
# 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 -all
```

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.

Solaris 10:

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

Solaris 9:

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

10 If required, apply the OS kernel patches.

See "System requirements" on page 16.

See Oracle's 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 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
```

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` and `/opt/VRTSvcs/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 On each node, enter the following command to check if any Storage Checkpoints are mounted:

```
# df -F vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# df -F vxfs
```

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

  ```
  # umount /filesystem
  ```

8 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.

9 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.

10 On each node, stop all VxVM volumes by entering the following command for each disk group:

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

Verify that no volumes remain open:

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

11 Stop VCS:

```
# hastop -all
```

12 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.

13 On each node, stop ODM, cluster fencing, GAB, and LLT in the following order:

- Solaris 9:

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

- Solaris 10:

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

14 If required, apply the OS kernel patches.

See "System requirements" on page 16.

See Oracle's documentation for the procedures.

15 On each node, check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 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 [-rsh] node1 node2 ... nodeN
```

Review the output.

17 Start services for LLT, GAB, cluster fencing and ODM on all upgraded nodes:

For Solaris 9:

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

For Solaris 10:

```
# svcadm enable llt
# svcadm enable gab
# svcadm enable vxfen
# svcadm enable vxodm
```

18 Start vcs on all upgraded nodes:

```
# /opt/VRTSvcs/bin/hastart
```

19 If necessary, reinstate any missing mount points in the `/etc/vfstab` 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 for each disk group:

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

25 If you stopped any RVGs in step 10, 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 SF Oracle RAC cluster

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

To upgrade to 5.1 RP2 on an SF Oracle RAC cluster

1 Log in as superuser.

2 Verify that `/opt/VRTSvcs/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 If CRS is not controlled by VCS, enter the following command on each node of the cluster to stop CRS:

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

7 Stop VCS.

```
# hastop -all
```

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 VCSMM and LMX if they are running:

For Solaris 9:

```
# /etc/init.d/vcsmm stop
# /etc/init.d/lmx stop
```

For Solaris 10:

```
# svcadm disable -t vcsmm
# svcadm disable -t lmx
```

10 Stop cluster fencing, ODM, and GAB:

For Solaris 9:

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

For Solaris 10:

```
# svcadm disable -t vxfen
# svcadm disable -t vxodm
# svcadm disable -t gab
```

11 On each node, unload the vxfen, LMX, GAB, VCSMM, GMS, and GLM kernel modules if they are still loaded:

- Verify if the `vxfen` kernel module is loaded. For example:

```
# modinfo|grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.1)
```

If the vxfen kernel module is loaded then unload it. For example:

```
# modunload -i 210
```

- Verify if the LMX kernel module is loaded. For example:

```
# modinfo | grep lmx
239 ffffffff1253000  13a30 236   1  lmx (LLT Mux '5.1')
```

If the LMX kernel module is loaded then unload it. For example:

```
# modunload -i 239
```

- Verify if the VCSMM kernel module is loaded. For example:

```
# modinfo | grep vcsmm
   312 78bc0000  43ae8 293   1  vcsmm (VRTSvcsmm 5.1)
```

If the VCSMM kernel module is loaded then unload it. For example:

```
# modunload -i 312
```

- Verify if the GMS kernel module is loaded. For example:

```
# modinfo | grep gms
311 78289c91   4867 292   1  vxgms (VxGMS 5.1 (SunOS))
```

If the GMS kernel module is loaded then unload it. For example:

```
# modunload -i 311
```

- Verify if the GLM kernel module is loaded. For example:

```
# modinfo | grep glm
310 78b68000  24268 291   1  vxglm (VxGLM 5.1 (SunOS 5.10))
```

If the GLM kernel module is loaded then unload it. For example:

```
# modunload -i 310
```

- Verify if the GAB kernel module is loaded. For example:

```
# modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.1)
```

If the GAB kernel module is loaded then unload it. For example:

```
# modunload -i 149
```

12 Stop LLT:

For Solaris 9:

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

For Solaris 10:

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

- Verify if the LLT kernel module is loaded. For example:

  ```
  # modinfo|grep llt
  147 50ca4000 d6bc 110 1 llt (LLT 5.1)
  ```

  If the LLT kernel module is loaded then unload it. For example:

  ```
  # modunload -i 147
  ```

13 If required, apply the OS kernel patches.

See "System requirements" on page 16.

See *Oracle's* documentation for the procedures.

Note: If you are upgrading a SF Oracle RAC cluster, you must upgrade the nodes of the cluster at this stage to one of the operating system versions that this release supports.

---

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

```
# df -F vxfs
```

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

  ```
  # umount /filesystem
  ```

- On each node of the cluster, verify that all file systems have been cleanly unmounted:

  ```
  # echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean
  flags 0 mod 0 clean clean_value
  ```

  A *clean_value* value of 0x5a indicates the file system is clean, 0x3c indicates the file system is dirty, and 0x69 indicates the file system is dusty. A dusty file system has pending extended operations.

- If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem
# mount -F vxfs filesystemmountpoint
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly. There may be a pending large fileset clone removal extended operation if the umount command fails with the following error:

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system
file system still in progress.
```

- If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
- Repeat the following command to verify that the unclean file system is now clean:

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean
flags 0 mod 0 clean clean_value
```

15 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.

16 On each node of the cluster, 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
```

17 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
```

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

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

If ssh is not configured then enter:

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

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

19 After the entire cluster is upgraded, reboot all of the nodes of the cluster.

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

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

21 Run the following commands to start the Storage Foundation for Oracle RAC processes:

For Solaris 9:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/odm start
# /etc/init.d/vxfen start
# /etc/init.d/vcsmm start
# /etc/init.d/lmx start
# /opt/VRTSvcs/bin/hastart
```

For Solaris 10:

```
# svadm enable llt
# svadm enable gab
# svadm enable vxodm
# svadm enable vxfen
# svadm enable vcsmm
# svadm enable lmx
# /opt/VRTSvcs/bin/hastart
```

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

```
# haconf -makerw
```

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

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

24 Make the configuration read-only:

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

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

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

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

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

27 If CRS is not controlled by VCS, enter the following command on each node to start CRS.

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

28 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

29 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

Upgrading to 5.1 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.

See "System requirements" on page 16.

See Oracle's documentation for the procedures.

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

```
# df -F 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 Copy the patch archive downloaded from the patch central to temporary location, untar the archive and browse to the directory containing the installrp installer script. Enter the `installrp` script:

```
# ./installrp nodename
```

11 If necessary, reinstate any missing mount points in the `/etc/vfstab` 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

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 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.

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 2 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 packages 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 packages.

6 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.

7 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/LMX/VCSMM and the VCS Engine (‘had’).

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

1 Consider a four node SFRAC 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, LMX, VCSMM, 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
# lmxconfig -W
# vcsmmconfig -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`, `/etc/gabtab`, and `/etc/vcsmmtab` 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, VCSMM, LMX, GAB and LLT in that order, on each of the nodes of the sub-cluster to be upgraded.

- Solaris 9:

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

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

- Solaris 10:

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

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, LMX, VCSMM, GAB, LLT, VXFEN will be upgraded together.

- Upgrade (use patchadd on Solaris) all the packages with new product version, except VCS and agent related packages on the sub-cluster being upgraded.
  - Sol_Sprac 2.9
    142629-05 VRTSvxvm
    141270-02 VRTSsfmh
    142633-03 VRTSvxfs
    143260-03 VRTSllt
    143262-03 VRTSgab
    143706-03 VRTSvxfen
    143279-03 VRTScps
    142631-03 VRTSdbed
    143270-03 VRTSodm
  - Sol_Sparc 2.10
    142629-05 VRTSvxvm
    141270-02 VRTSsfmh
    142634-03 VRTSvxfs
    143261-03 VRTSllt
    143263-03 VRTSgab
    143707-03 VRTSvxfen
    143279-03 VRTScps
    142631-03 VRTSdbed
    143271-03 VRTSodm
  - x-86
    142630-05 VRTSvxvm

141752-02 VRTSsfmh
142635-03 VRTSvxfs
143266-03 VRTSllt
143267-03 VRTSgab
143708-03 VRTSvxfen
143280-03 VRTScps
142632-03 VRTSdbed
143272-03 VRTSodm

- Re-link oracle in case of SFCFS RAC.
- In reverse order, start the components that you previously stopped. For example, on Solaris 10:

```
# svcadm enable llt
# svcs -a|grep llt
online          0:22:44 svc:/system/llt:default
# svcadm enable gab
# svcs -a|grep gab
online          0:25:07 svc:/system/gab:default
# svcadm enable vxodm
# svcs -a|grep vxodm
online          0:25:20 svc:/system/vxodm:default
# svcadm enable vxfen
# svcs -a|grep vxfen
online          0:25:37 svc:/system/vxfen:default
```

For Solaris 9, the commands are listed under /etc/init.d to start modules. For example, to start LLT, use /etc/init.d/llt start, etc.
LLT and LMX start communicating with the new version automatically when they are started.
Start HA.

```
# hastart
```

- Once all the services are started, start use the `hastart` command to start HA. All ports should come up successfully and cluster start communication the other nodes.

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 and LMX 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 `vcsmmconfig -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`, `/etc/gabtab`, and `/etc/vcsmmtab` 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
- Modify the VCS configuration to reflect version specific requirements if any.
- Upgrade VRTSvcs and agent-related patches:
  - Sol_Sparc 2.10
    VRTSvcs 143264-07
    VRTSvcsag 143265-07
    VRTSvcsea 143276-04
    VRTScavf 143274-02
  - Sol_Sparc 2.9
    VRTSvcs 143264-07
    VRTSvcsag 143265-07
    VRTSvcsea 143276-04
    VRTScavf 143273-02
  - x-86
    VRTSvcs 143268-07
    VRTSvcsag 143269-07
    VRTSvcsea 143277-03
    VRTScavf 143275-02
- 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.

- Modify the VCS configuration to reflect version specific requirements if any.
- Upgrade VRTSvcs and agent-related patches:
  - Sol_Sparc 2.10
    VRTSvcs 143264-07
    VRTSvcsag 143265-07
    VRTSvcsea 143276-04
    VRTScavf 143274-02
  - Sol_Sparc 2.9
    VRTSvcs 143264-07
    VRTSvcsag 143265-07
    VRTSvcsea 143276-04
    VRTScavf 143273-02
  - x-86
    VRTSvcs 143268-07
    VRTSvcsag 143269-07
    VRTSvcsea 143277-03
    VRTScavf 143275-02
- 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
  ```

Chapter 5

Removing and rolling back

This chapter includes the following topics:

About removing and rolling back 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 for Veritas Storage Foundation (SF), Veritas Storage Foundation Cluster File System (SFCFS), and Veritas Storage Foundation for Oracle RAC (SFRAC). 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.

Note: Symantec recommends using the following steps to roll back. There is no `uninstallrp` utility to roll back the patches.

Removing 5.1 RP2 from Veritas Cluster Server

Use the following procedure to remove VCS 5.1 RP2 from your cluster manually.

To remove 5.1 RP2 from VCS manually

1. List the service groups in your cluster and their status. On any node, type:

```
# hagrp -state
```

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

```
# hagrp -offline -force ClusterService -sys system
```

3. Make the VCS configuration writable. On any node, type:

```
# haconf -makerw
```

4. Freeze all service groups. On any node, type:

```
# hagrp -freeze service_group -persistent
```

   where *service_group* is the name of the service group. Note that the ClusterService group cannot be frozen.

5. Save the configuration (`main.cf`) file with the groups frozen. On any node, type:

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

6. Make a backup copy of the current `main.cf` and all `types.cf` configuration files. For example, on one node in the cluster, type:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/types.cf.save
```

7. Shut down VCS. On any node, type:

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

8. Shut down CmdServer. On each node, type:

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

9 Verify that VCS has shut down. On any node, type:

```
# /sbin/gabconfig -a
```

The output resembles: GAB Port Memberships Port a gen 23dc0001 membership 01 The output shows no membership for port h.

10 For Solaris 10, on nodes that run non-global zones, check if the non-global zones are in the running state. Boot the non-global zones that are not in the running state.

- Check the zone's state. On each node, type:

```
# zoneadm list -icv
```

- Boot the zone if it is not in the running state. On each node, type:

```
# zoneadm -z zone boot
```

where *zone* is the name of the non-global zone.

Note: Do not configure one or more Solaris zones to boot from the shared storage.

---

11 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:

```
# /sbin/vxfenconfig -U
```

12 Unload vxfen. On each node, perform the following steps:

- Identify the vxfen kernel module, for example:

```
# modinfo|grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.1RP2)
```

- Unload vxfen using the module number.

```
# modunload -i 210
```

13 Unconfigure GAB. On each node, type:

```
# /sbin/gabconfig -U
```

14 Unload GAB. On each node, perform the following steps:

- Identify the GAB kernel module. For example:

```
# modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.1RP2)
```

- Unload GAB using the module number:

```
# modunload -i 149
```

15 Unconfigure LLT. On each node, perform the following steps:

- Type:

```
# /sbin/lltconfig -U
```

- Type **y** on each node in response to the message.

16 Unload LLT. On each node, perform the following steps:

- Identify the LLT kernel module. For example:

```
# modinfo | grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.1RP2)
```

- Unload LLT using the module number:

```
# modunload -i 147
```

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

- Get the list of 5.1 RP2 patches, type:

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

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

```
# patchrm 143260-02
```

18 Verify that the patches have been removed. On each node, type:

```
# showrev -p | grep VRTS
```

19 If the LLT, GAB, or VXFEN modules cannot be stopped or unloaded following the patch removal, reboot all nodes in the cluster.

20 If you do not perform step 19, start the VCS components manually. On each node, type:

```
# /sbin/lltconfig -c
# /sbin/gabconfig -cx
# /sbin/vxfenconfig -c
# /opt/VRTSvcs/bin/hastart
```

You do not have to start vxfen unless you use the fencing option.

21 After VCS has started, perform the following steps:

- Verify all resources have been probed. On any node, type:

```
# hastatus -summary
```

- Unfreeze all service groups. On any node, type:

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

where *service_group* is the name of the service group.

22 Bring online the ClusterService service group, if necessary. On any node type:

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

where *system* is the node name.

Removing 5.1 RP2 on Veritas Storage Foundation or Veritas Storage Foundation Cluster File System

You can use the following procedure to uninstall 5.1 RP2 on SF or SFCFS.

To uninstall 5.1 RP2 on Veritas Storage Foundation or Veritas Storage Foundation Cluster File System

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

4 Check if the root disk is under VxVM control by running this command:

```
# df -v /
```

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

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

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

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

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the rootdg disk group in addition to the root disk for vxunroot to succeed.

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

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

5 Enter the following command to check if any VxFS file systems are mounted:

```
# df -F vxfs
```

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

```
# umount /filesystem
```

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

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

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

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

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

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

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

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

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

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

9. Stop VCS along with all its resources. Then, stop the remaining resources manually:

   - For Solaris 9:

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

   - For Solaris 10:

     ```
     # svcadm disable vcs
     ```

10. If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

    ```
    # rm /etc/vxfenmode
    ```

11. Unmount `/dev/odm`:

    ```
    # umount /dev/odm
    ```

12. Unload the ODM module:

    ```
    # modinfo | grep odm
    # modunload -i odm_mod_id
    ```

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

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

```
# svcadm disable gab
# svcadm disable llt
```

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 To shut down and remove the installed Veritas packages, use the appropriate command in the /opt/VRTS/install directory. For example, to uninstall the Storage Foundation, enter the following commands:

```
# cd /opt/VRTS/install
# ./uninstallsf [-rsh]
```

You can use this command to remove the packages from one or more systems. For other products, substitute the appropriate script for `uninstallsf` such as `uninstallsfcfs` for the Storage Foundation Cluster File System software. The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to uninstall the software simultaneously on several systems.

Note: Provided that the remote shell (RSH) or secure shell (SSH) has been configured correctly, this command can be run on a single node of the cluster to install the software on all the nodes of the sub-cluster.

---

17 After uninstalling the Veritas software, refer to the appropriate product's 5.1 Installation Guide document to reinstall the 5.1 software.

Removing 5.1 RP2 on Veritas Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 5.1 RP2 on Storage Foundation for Oracle RAC systems.

To uninstall the 5.1 RP2 on Veritas Storage Foundation for Oracle RAC

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

```
# hagrp -offline oracle_group -sys node_name
```

If the database is not managed by VCS, stop the Oracle database as follows:

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

2 If CRS is not under VCS Control, then enter the following command on each node of the cluster to stop CRS.

- For 10gR2 or 11gR1:

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

- For 11gR2:

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

3 Stop the applications that use CVM or CFS that are not under VCS control

- Using native application commands, stop the applications that use CVM or CFS on all nodes.
- Verify that no processes use the CFS mount point:

```
# fuser -c mount_point
```

4 Unmount CFS file systems that are not under VCS control

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

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

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

```
# umount mount_point
```

5 Stop VCS to take the service groups on all nodes offline

On any one node execute following command to stop VCS:

```
# hastop -all
```

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

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

```
# fuser -c mount_point
```

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

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

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

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

```
# umount mount_point
```

8 Remove SF for Oracle RAC.

- On any one node, navigate to the directory that contains the uninstallsfrac program:

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

- Start the uninstallsfrac program:

```
# ./uninstallsfrac
```

9 After uninstalling the SF for Oracle RAC 5.1 RP2, refer to the *Veritas Storage Foundation for Oracle RAC 5.1 Installation and Configuration Guide* document to reinstall the 5.1 software.

Chapter 6

Verifying software versions

This chapter includes the following topics:

Verifying software versions

To verify if the Veritas patches are installed on your system, enter the following command:

```
# showrev -p|grep patch_id
```