Veritas Storage Foundation and High Availability Solutions Release Notes

AIX

5.1 Rolling Patch 1

symantec™

Storage Foundation and High Availability Solutions Release Notes 5.1 Rolling Patch 1

Product version: 5.1 RP1

Document version: 5.1RP1.1

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

Release Notes

This document includes the following topics:

- Introduction
- System Requirements
- List of products
- Fixed issues
- Known issues
- Software limitations
- Changes in Storage Foundation High Availability
- Downloading the rolling patch archive
- List of patches
- Installing the Veritas software for the first time
- Installing 5.1 RP1 using the web-based installer
- Prerequisites for upgrading to 5.1 RP1
- Supported upgrade paths
- Upgrading 5.1 to 5.1 RP1
- Verifying software versions
- Removing and rolling back
- Documentation addendum

Introduction

This document provides information about the Storage Foundation and High Availability Solutions 5.1 Rolling Patch 1.

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

http://entsupport.symantec.com/docs/334998

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.

System Requirements

This section describes the system requirements for this release

Supported AIX operating systems

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

Product installation scripts verify the required update levels. The installation process terminates if the target systems do not meet the maintenance level requirements.

For any Veritas cluster product, all nodes in the cluster must have the same operating system version and update level.

The minimum system requirements for this release are as follows:

AIX 5.3 at one of the following levels:

- TL7 with SP2
- or any higher TLs

AIX 6.1 at one of the following levels:

- TL0 with SP4
- or any higher TLs

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

http://entsupport.symantec.com/docs/334998

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.

Additional Oracle support for SF Oracle RAC

Table 1-1 Oracle RAC versions that SF Oracle RAC supports

Oracle version	AIX 5.3	AIX 6.1
10gR2 10.2(64-bit)	Yes	Yes
11gR1 11.1(64-bit)	Yes	Yes
11gR2	Yes	Yes

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

The following sections describe the Veritas Storage Foundation High Availability issues that were fixed in this release.

- Veritas Storage Foundation fixed issues in 5.1 RP1
- Veritas Volume Manager fixed issues in 5.1 RP1 release
- Veritas File System fixed issues in 5.1 RP1 release
- Veritas Storage Foundation Cluster File System fixed issues in 5.1 RP1 release
- Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 RP1
- Veritas Cluster Server fixed issues in 5.1 RP1
- Storage Foundation Manager fixed issues in 5.1 RP1
- VEA fixed issues in 5.1 RP1

Veritas Volume Manager fixed issues in 5.1 RP1 release

Table 1-2 Veritas Volume Manager 5.1 RP1 fixed issues

Fixed issues	Description
1948412	64 bit VxMS plugins of VxVM are missing in the path /opt/VRTSvxms/lib/map/aix64
1948167	vxdisk resize cannot handle over 1TB gpt labeled disk as expected
1940166	Problem while creating more than 30 shared DGs
1938484	EFI: Prevent multipathing don't work for EFI disk
1915356	I/O stuck in vxvm caused cluster node panic
1935332	NASGW:vxdisk updateudid marks the disk as clone_disk
1935230	Panic in voldsio_timeout() function.
1932091	Need for dmp_revive_paths() in dmp reconfiguration/restore_demon code path.
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

Table 1-2 Veritas Volume Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1892387	VM commands getting hung on master node with 32-node cluster.
1892372	vxdisk reclaim should skip cache object update
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.
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
1852212	When vxesd is enabled, dmp/dr procedure(CLAB CCT test case #1805-phase 2) with PowerPath panics the system
1846165	Data corruption seen on cdsdisks on Solaris-x86 in several customer cases
1840673	After adding new luns one of the nodes in 3 node CFS cluster hangs
1835139	CERT : pnate test hang I/O greater than 200 seconds during the filer giveback
1826088	After pulling out FC cables of local site array, plex became DETACHED/ACTIVE
1792795	supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
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.
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.

Veritas File System fixed issues in 5.1 RP1 release

Table 1-3 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number)

Fixed issues	Description
1897458, 1805046	Fixed issue in alert generation from vxfs when file system usage threshold is set.
1933635, 1914625	Fixed issues in fs pattern assignment policy of the file system.
1933975, 1844833	Fixed VX_EBMAPMAX error during filesystem shrinking using fsadm..
1934085, 1871935	We now update ilist on secondary even if error received from primary for a VX_GETIAS_MSG is EIO.
1934095, 1838468	Fixed a race in qiostat update which was resulting in data page fault.
1934096, 1746491	Fix to avoid core dump while running fsvmap by initializing a local pointer.
1934098, 1860701	Moved drop of active level and reaquire to top of loop to stop resize from being locked out during clone removal.
1934107, 1891400	Fixed incorrect ACL inheritance issue by changing the way it cached permission data.
1947356, 1883938	Added utility mkdstfs to create DST policies.
1934094, 1846461	Fixed an issue with vxfsstat(1M) counters.

Veritas Storage Foundation fixed issues in 5.1 RP1

Table 1-4 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.

Table 1-4 Veritas Storage Foundation fixed issues in 5.1 RP1 *(continued)*

Fixed issues	Description
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 Storage Foundation Cluster File System fixed issues in 5.1 RP1 release

Table 1-5 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.
1961790, 1986445	Fixed issue in the mount(1M) command to correctly set the master node.
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 in 5.1 RP1

Table 1-6 Veritas Storage Foundation for Oracle RAC 5.1 RP1fixed issues

Fixed issues	Description
1980842, 1983222	Fixed issue in cfsadmin command for RVG volumes.

Table 1-6 Veritas Storage Foundation for Oracle RAC 5.1 RP1fixed issues *(continued)*

Fixed issues	Description
1932827	Fixed PrivNIC Agent to support AIX LPAR interfaces with native 64k MTU.
1908924	Fixed an issue in MultiPrivNIC Agent where the IP failed over by this agent does not ping on AIX if ARP refresh timeout is set to a high value.
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.
1891389	Issue: cssd agent support for Oracle 11gR2 Resolution: cssd agent is modified to support 11gR2 framework. 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 in 5.1 RP1

Table 1-7 Veritas Cluster Server 5.1 RP1 fixed issues

Fixed issues	Description
1967207	vxfenconfig -l reports multiple (duplicate) paths after phased upgrade on the first sub-cluster.
1946367	Change in LLT delivery thread's priority limits
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
1916004	ASMagent connecting as sysdba instead of sysasm for 11gR2

Table 1-7 Veritas Cluster Server 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1915909	[VCS][281-889-442] hares allows to create resources which has "." special character
1915016	[VCS][281-795-096] Port h halting system due to internal protocol error on gab_sf_dlv_gaps().
1900450	Race script is killed if it exceeds the script time-out.
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
1855196	System panic due to depleted memory during GAB broadcast stress and reboot node 0.
1839091	SxRT5.1:SFRAC:Resource coordpoint became FAULTED from time to time.
1809827	Largenode: Node32 and node31 not able to join cluster

Storage Foundation Manager fixed issues in 5.1 RP1

Table 1-8 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

Table 1-8 Storage Foundation Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
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
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

Table 1-8 Storage Foundation Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
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
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

VEA fixed issues in 5.1 RP1

Table 1-9 VEA 5.1 RP1 fixed issues

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

Table 1-9 VEA 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
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 5.1 RP1 release.

- Veritas Storage Foundation known issues in 5.1 RP1 release
- Veritas Volume Manager known issues in 5.1 RP1 release
- Veritas File System known issues in 5.1 RP1 release
- Veritas Storage Foundation Cluster File System known issues in 5.1 RP1 release
- Veritas Storage Foundation for Oracle RAC known issues in 5.1 RP1
- Veritas Cluster Server known issues in 5.1 RP1
- Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 RP1

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

Veritas Storage Foundation known issues in 5.1 RP1 release

The following are new additional Storage Foundation known issues in this 5.1 RP1 release.

dbed_clonedb of offline checkpoint fails with ORA-00600 with Oracle 11gR2 when ODM is enabled (1982674)

When performing offline checkpoint database cloning on Oracle 11gR2 and ODM is enabled, the `dbed_clonedb` command fails with error:

```
$ dbed_clonedb -S mofcln1 -m /tmp/mofcln1 -c \
Checkpoint_1267604996
SFORA dbed_clonedb ERROR V-81-4920 Database mofcln1 is still in
recovery mode.
SFORA dbed_clonedb ERROR V-81-4881 Log file is at /tmp/oralog.out.10392.
```

The `/tmp/oralog.out.10392` file indicates an error.

Sample output of the `/tmp/oralog.out.10392` file:

```
ALTER DATABASE OPEN RESETLOGS
*
ERROR at line 1:
ORA-00600: internal error code, arguments: [ksfdgmsn4],
[ODM ERROR V-41-4-2-207-1 Operation not permitted],
[], [], [], [], [], [], [], [], [], []
ORA-00318: log 1 of thread 1, expected file size 512 doesn't match 512
ORA-00312: online log 1 thread 1:
'/tmp/mofcln1/snap_data11r2/FLAS11r2/redo01.log'
```

Note: This issue may occur in a VVR environment.

Workaround:

Perform the offline checkpoint cloning for 11gR2 on another `ORACLE_HOME` where ODM is disabled.

Dbed_ckptrollback fails for -F datafile option for Oracle database version 11gr2 (1959400)

On Oracle 11gr2 database, dbed_ckptrollback fails with following error "SFORA rb.file ERROR V-81-3038 Error occurred while querying Oracle Database." The root cause of this problem is an Oracle 11GR2 defect (8367917).

Workaround:

To manually recover the datafile

1. Take the corrupt data file offline.
2. Mount the checkpoint using dbed utilities.
3. Restore the corrupt file manually.
4. Recover the datafile.
5. Bring the datafile online.

Veritas Volume Manager known issues in 5.1 RP1 release

The following are new additional Veritas Volume Manager known issues in this 5.1 RP1 release.

vxesd dump core when it starts (1897007)

This issue happens during the case when the system is connected to a switch with more than 64 ports.

Workaround: To fix the issue, change the switch to lesser port number.

Cannot restore root file system that was backed up by mksysb (1989057)

The issue is seen when the boot disk is on SAN with multiple paths and if you back up by using `mksysb` to restore it.

The RCA for this issue is completed, this should be available with next HotFix or Rolling Patch.

Unable to initialize and use ramdisk for VxVM use (1825516)

Cannot initialize a ramdisk with “vxdisk –f init ramdisk type=nopriv volatile".

Veritas File System known issues in 5.1 RP1 release

No additional known issues exist for Veritas File System in the 5.1 RP1 release.

Veritas Storage Foundation Cluster File System known issues in 5.1 RP1 release

The following are new additional Veritas Storage Foundation Cluster File System known issues in this 5.1 RP1 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.

Veritas Storage Foundation for Oracle RAC known issues in 5.1 RP1

The following are new additional Veritas Storage Foundation for Oracle RAC known issues in this 5.1 RP1 release.

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

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

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

Or

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

Workaround:

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

Veritas Cluster Server known issues in 5.1 RP1

The following are new additional Veritas Cluster Server known issues in this 5.1 RP1 release.

Oracle agent

The Oracle agent does not support health check monitoring for 11g R2.

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.

Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 RP1

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in this 5.1 RP1 release.

Health Check monitoring does not work with 11gR1 and 11gR2 (1985055)

Health Check monitoring does not work with 11gR1 and 11gR2.

Workaround: Set MonitorOption attribute for Oracle resource to 0.

An issue with the ohasd process (1985093)

There is an issue the `ohasd` process.

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

Intentional Offline

Intentional Offline does not work.

ASMInstAgent does not support having pfile or spfile

The ASMInstAgent does not support having `pfile` or `spfile` for the ASM Instance on the ASM disk groups.

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

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 RP1 release
- Veritas Volume Manager software limitations in 5.1 RP1 release
- Veritas Storage Foundation for Oracle RAC software limitations in 5.1 RP1

Veritas Storage Foundation software limitations in 5.1 RP1 release

The following are additional Veritas Storage Foundation software limitations in this 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 RP1 release

The following are additional Veritas Volume Manager software limitations in this 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 RP1

The following are additional Veritas Storage Foundation for Oracle RAC software limitations in this release.

CRSResource agent

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

Changes in Storage Foundation High Availability

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

About the new installrp script

Veritas Storage Foundation and High Availability Solutions 5.1 RP1 provides a new upgrade script. To upgrade from Veritas Storage Foundation and High Availability Solutions version 5.1 or later, the recommended upgrade method is to use the new upgrade script. The `installrp` script allows you to upgrade all the patches associated with the packages installed and then starts all the processes.

installrp script options

Table 1-10 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.

Table 1-10 shows command line options for the product upgrade script *(continued)*

Command Line Option	Function
[-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.
[-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.
[-tmppath *tmp_path*]	The -tmppath option is used to select a directory other than /var/tmp as the working directory for installrp. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
[-hostfile *hostfile_path*]	The -hostfile option specifies the location of a file containing the system names for installer.
[-keyfile *ssh_key_file*]	The -keyfile option specifies a key file for SSH. When this option is used, -i <*ssh_key_file*> is passed to every SSH invocation.
[-patchpath *patch_path*]	The -patchpath option is used to define the complete path of a directory available to all install systems (usually NFS mounted) that contains all patches to be installed by installrp.

Table 1-10 shows command line options for the product upgrade script *(continued)*

Command Line Option	Function
`[ -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. This option is available with or without one of following options: -allpkgs, -minpkgs, and -recpkgs. 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.

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 64.

Downloading the rolling patch archive

The patches included in the 5.1 RP1 release are available for download from the Symantec website. After downloading the 5.1 RP1 file, use the gunzip and tar to uncompress and extract.

For the 5.1 RP1 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

http://entsupport.symantec.com/docs/334998

List of patches

This section lists the patches and filesets.

Table 1-11 Filesets and products affected

5.1 fileset names	AIX OS release level	Fileset size	Fileset version	Products affected
VRTScavf.bff	5.3/6.1	250 KB	05.01.0001.0000	SFCFS, SF Oracle RAC
VRTScps.bff	5.3/6.1	2.9 MB	05.01.0001.0000	VCS, SFHA, SF Oracle RAC
VRTSdbac.bff	5.3/6.1	9 MB	05.01.0001.0000	SF Oracle RAC
VRTSdbed.bff	5.3/6.1	37 MB	05.01.0001.0000	SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC
VRTSgab.bff	5.3/6.1	4.5 MB	05.01.0001.0000	VCS, SFHA, SFCFS, SF Oracle RAC
VRTSllt.bff	5.3/6.1	2.7 MB	05.01.0001.0000	VCS, SFHA, SFCFS, SF Oracle RAC

Table 1-11 Filesets and products affected *(continued)*

5.1 fileset names	AIX OS release level	Fileset size	Fileset version	Products affected
VRTSob.bff	5.3/6.1	60 MB	03.04.0235.0027	SF, SFHA, SFCFS, SF Oracle RAC
VRTSodm.bff	5.3/6.1	250 KB	05.01.0001.0000	SF, SFHA, SFCFS, SF Oracle RAC
VRTSsfmh.bff	5.3/6.1	29 MB	02.01.0198.0031	SF, SFHA, SFCFS, SF Oracle RAC
VRTSvcs.bff	5.3/6.1	63 MB	05.01.0001.0000	VCS, SFHA, SFCFS, SF Oracle RAC
VRTSvcsag.bff	5.3/6.1	3 MB	05.01.0001.0000	VCS, SFHA, SFCFS, SF Oracle RAC
VRTSvcsea.bff	5.3/6.1	200 KB	05.01.0001.0000	VCS, SFHA, SFCFS, SF Oracle RAC
VRTSvxfen.bff	5.3/6.1	3 MB	05.01.0001.0000	VCS, SFHA, SFCFS, SF Oracle RAC
VRTSvxfs.bff	5.3/6.1	18 MB	05.01.0001.0000	SF, SFHA, SFCFS, SF Oracle RAC
VRTSvxvm.bff	5.3/6.1	62 MB	05.01.0001.0000	SF, SFHA, SFCFS, SF Oracle RAC

Installing the Veritas software for the first time

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 RP1. 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 GA binaries. Choose one of the following to start the installation:

   - For Storage Foundation:

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

   - For Storage Foundation HA:

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

   - For Storage Foundation Cluster File System:

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

   - For Storage Foundation Cluster File System HA:

     ```
     # ./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 RP1.

   See "Prerequisites for upgrading to 5.1 RP1" on page 35.

3. Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program.

   - If the 5.1 product is installed and configured, then run the `installrp` script to install 5.1 RP1.

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

     See "About the new installrp script" on page 25.

   - If the 5.1 product is installed and not configured, run the `installrp` script to install 5.1 RP1 and configure the product.

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

See "About the new installrp script" on page 25.
The `installrp` script will give you an option to configure the product. If you choose not to configure the product at the time of the 5.1 RP1 installation, then 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 HA:

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

- For Storage Foundation Cluster File System:

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

- For Storage Foundation Cluster File System HA:

```
# ./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 for your product.

Installing 5.1 RP1 using the web-based installer

This section describes how to install 5.1 RP1 using the web-based installer.

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

About the Web-based installer

The `webinstaller` script is used to start and stop the Veritas XPortal Server `xprtlwid` process. The `webinstaller` script can also be used to check the status of the XPortal Server.

When the `webinstaller` script starts the `xprtlwid` process, the script displays a URL. Use this URL to access the Web-based installer from a Web browser such as Internet Explorer or FireFox.

The Web installer creates log files whenever the Web installer is operating. While the installation processes are operating, the log files are located in a session-based directory under the `/var/tmp` directory. After the install process completes, the log files are located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

The location of the Veritas XPortal Server configuration file is `/var/opt/webinstaller/xprtlwid.conf`.

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 1-12 Web-based installer requirements

System	Function	Requirements
Target system	The system(s) where the Veritas products will be installed.	Must be a supported platform for Veritas product 5.1 RP1
Installation server	The server from which to initiate the installation. The installation media is mounted and accessible from the installation server.	Must be the same OS as the system(s) on which to install.
Administrative system	The system on which you run the web browser to perform the installation.	Web browser

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. On the administrative server, start the Web browser.
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.

Performing a pre-installation check with the Veritas Web-based installer

This section describes performing a pre-installation check with the Veritas Web-based installer.

To perform a pre-installation check

1. Start the Web-based installer.

   See "Starting the Veritas Web-based installer" on page 32.
2. On the Select a task and a product page, select **Perform a Pre-installation check** from the **Task** drop-down list.

3 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Validate**.

4 The installer performs the precheck and displays the results.

5 If the validation completes successfully, click **Next**. The installer prompts you to begin the installation. Click **Yes** to install the Veritas product on the selected system. Click **No** to install later.

6 Click **Finish**. The installer prompts you for another task.

Installing products with the Veritas Web-based installer

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

To install Veritas product

1 Perform preliminary steps.

2 Start the Web-based installer.

See "Starting the Veritas Web-based installer" on page 32.

3 On the Select a task and product page, select **Install RP1** from the **Task** drop-down list.

4 Select Veritas product or Veritas product High Availability from the Product drop-down list, and click Next.

5 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.

6 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.

7 After the validation completes successfully, click **Next** to install Veritas product on the selected system.

8 For Storage Foundation, click Next to complete the configuration and start the product processes.

For Storage Foundation High Availability, the installer prompts you to configure the cluster.

Note that you are prompted to configure only if the product is not yet configured.

If you select n, you can exit the installer. You must configure the product before you can use Veritas product.

See the Veritas product's Installation Guide to configure the product.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

9 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**. The installer prompts you for another task.

Prerequisites for upgrading to 5.1 RP1

The following list describes prerequisites for upgrading to the 5.1 RP1 release:

- For any product in the Storage Foundation stack, regardless of your operating system, you must have the 5.1 release installed before you can upgrade that product to the 5.1 RP1 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 RP1
- 5.1 P1 to 5.1 RP1
- 5.1 to 5.1 P1 to 5.1 RP1

Upgrading 5.1 to 5.1 RP1

This section describes how to upgrade from 5.1 to 5.1 RP1 on a cluster or a standalone system.

- Performing a full upgrade to 5.1 RP1 on a cluster
  Use the procedures to perform a full upgrade to 5.1 RP1 on a cluster that has VCS, SFHA, SFCFS, or SF Oracle RAC installed and configured.
- Upgrading Veritas product with the Veritas Web-based installer
  Use the procedure to upgrade your Veritas product with the Web-based installer.
- Performing a rolling upgrade using the installer
  Use the procedure to upgrade your Veritas product with a rolling upgrade.
- Performing a rolling upgrade manually
  Use the procedure to upgrade your Veritas product manually with the rolling upgrade.
- Upgrading to 5.1 RP1 on a standalone system
  Use the procedure to upgrade to 5.1 RP1 on a system that has SF and VCS installed.

Performing a full upgrade to 5.1 RP1 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 RP1:

- Performing a full upgrade to 5.1 RP1 for VCS
- Performing a full upgrade to 5.1 RP1 on a SFHA cluster
- Performing a full upgrade to 5.1 RP1 on a SFCFS cluster
- Performing a full upgrade to 5.1 RP1 on a SF Oracle RAC cluster

Performing a full upgrade to 5.1 RP1 for VCS

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

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

To upgrade VCS

1 Review the installation prerequisites.

See "Prerequisites for upgrading to 5.1 RP1" on page 35.

2 Check the readiness of the nodes where you plan to upgrade. Start the pre-upgrade check:

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

See "About the new installrp script" on page 25.

3 Resolve any issues that the precheck finds.

4 Start the upgrade:

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

5 After the upgrade, review the log files.

6 Verify the upgrade.

See "Verifying software versions" on page 56.

Performing a full upgrade to 5.1 RP1 on a SFHA cluster

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

To perform a full upgrade to 5.1 RP1 on a 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 On each node, enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

7 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
# 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 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 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
```

12 Navigate to the folder that contains the `installrp` program and check the readiness of the systems where you plan to upgrade. The command to start the pre-upgrade check is:

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

The program proceeds in a noninteractive mode to examine the systems for licenses, filesets, disk space, system-to-system communications, etc.

13 Review the output as the program displays the results of the check and saves the results of the check in a log file.

14 Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved.

See "Prerequisites for upgrading to 5.1 RP1" on page 35.

15 Navigate to the folder that contains the installrp program and start the `installrp` program:

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

Review the output.

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 Restart all the volumes. Enter the following command for each disk group:

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

21 If you stopped any RVGs in step 8, restart each RVG:

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

22 Remount all VxFS file systems on all nodes in the selected group:

```
# mount /filesystem
```

23 Remount all Storage Checkpoints on all nodes in the selected group:

```
# mount /checkpoint_name
```

Performing a full upgrade to 5.1 RP1 on a SFCFS cluster

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

To perform a full upgrade to 5.1 RP1 on an SFCFS cluster

1 Log in as superuser.

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

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

```
# haconf -makerw
```

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

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

5 Make the configuration read-only:

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

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

```
# mount | grep vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# mount | grep vxfs
```

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

```
# umount /filesystem
```

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 -local
```

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

```
# /etc/rc.d/rc2.d/S99odm stop
```

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

```
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/rc.d/rc2.d/S70llt stop
```

15 If required, apply the OS kernel patches.

See "System Requirements" on page 8.

See *IBM's* documentation for the procedures.

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

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

If the VEA service is running, stop it:

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

17 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. Enter 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.

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

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

20 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

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

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

22 Make the configuration read-only:

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

23 Bring the CVM service group online on each node:

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

24 Restart all the volumes by entering the following command 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 RP1 on a SF Oracle RAC cluster

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

To upgrade to 5.1 RP1 on a SF Oracle RAC cluster

1 Log in as superuser.

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

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

```
# haconf -makerw
# hasys -freeze -persistent nodename
# haconf -dump -makero
```

4 Stop Oracle database on the cluster:

- If the Oracle RAC instance is managed by VCS:

```
# hagrp -offline oracle_group -sys galaxy
# hagrp -offline oracle_group -sys nebula
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and shut down the instances:

```
$ srvctl stop instance -d database_name -i instance_name
```

5 Stop all applications on the cluster that are not configured under VCS. Use native application commands to stop the application.

6 Unmount the VxFS and CFS file systems that are not managed by VCS.

- Ensure that no processes are running that make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point, enter the following commands:

```
# mount | grep vxfs
# fuser -cu /mount_point
# umount /mount_point
```

Unmount the VxFS or CFS file system:

```
# umount /mount_point
```

7 Stop all VxVM and CVM volumes for each diskgroup that are not managed by VCS on the cluster:

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

Verify that no volumes remain open:

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

8 Stop VCS on the cluster:

```
# hastop -all
```

9 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. Enter the `installrp` script:

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

10 Relink the SF Oracle RAC libraries with Oracle.

Refer to *Veritas Storage Foundation for Oracle RAC 5.1 Installation and Configuration Guide* for more information.

11 Restart the nodes:

```
# shutdown -r now
```

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

```
# haconf -makerw
# hasys -unfreeze -persistent nodename
# haconf -dump -makero
```

13 Start VCS on each of the nodes:

- For parallel groups:

```
# hagrp -online group_name -sys nodename
```

- For failover groups:

```
# hagrp -online group_name -any
```

14 If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and start the instances:

```
$ srvctl start instance -d database_name -i instance_name
```

15 Manually mount the VxFS and CFS file systems that are not managed by VCS.

16 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.

Upgrading Veritas product with the Veritas Web-based installer

This section describes upgrading Veritas product with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

Note: Upgrading SF Oracle RAC with the Web-based installer is not supported.

To upgrade Veritas product

1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.

2 Start the Web-based installer.

See "Starting the Veritas Web-based installer" on page 32.

3 Select **Install RP1**.

The installer detects the product that is installed on the specified system.

4 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Validate**.

5 Click **Next** to complete the upgrade.

After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

6 Click **Finish**. After the upgrade, if the product is not configured, the web-based installer asks: "Do you want to configure this product?" If the product is already configured, it will not ask any questions.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native operating system accounts.

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

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 RP1 or from 5.1 P1 to 5.1 RP1.

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.

Performing a rolling upgrade on kernel packages: phase 1

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

To perform the rolling upgrade on kernel packages: phase 1

1 On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.

```
./installrp  -upgrade_kernelpkgs nodeA
```

Review the EULA, if you accept its terms, enter **y** to proceed.

2 Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.

3 Note the installation log location. The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.

4 The installer further replaces kernel components. Review the output.

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

6 Repeat step 1 to 5 on the second subcluster.

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

You now upgrade the non-kernel packages..

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

1 Start the installer for the rolling upgrade with the `-upgrade_nonkernelpkgs` option. Specify all the nodes in the cluster:

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

Review the EULA, if you accept its terms, enter **y** to proceed.

2 Note the installation log location. The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.

3 The installer upgrades non-kernel components. Review the output.

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

5 Manually check 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.

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.

10 Stop 'had' on the sub-cluster to be upgraded, and switch any remaining failover Service Groups 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, GMS, GLM, VCSMM, LMX, GAB and LLT in that order, on each of the nodes of the sub-cluster to be upgraded.

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

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

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

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

- LLT 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
    ```

  - Install new version of VRTSvcs and agent related packages.
  - 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. There will be no HA for the sub-cluster being upgraded, from this step onwards.

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

  - Install new version of VRTSvcs and agent related packages.

- Force stop VCS on the remaining sub-cluster. There is no HA for the entire cluster from this point onwards.

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

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

  ```
  # hastart
  ```

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

  ```
  # hastart
  ```

Upgrading to 5.1 RP1 on a standalone system

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

To upgrade to 5.1 RP1 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 8.

   See *IBM's* documentation for the procedures.
4. Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

   ```
   # mount | grep vxfs
   ```
5. Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

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

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

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

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

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

Verify that no volumes remain open:

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

9 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

10 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. Enter the `installrp` script:

```
# ./installrp nodename
```

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

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

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

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

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

14 Remount all VxFS file systems and Storage Checkpoints:

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

15 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

Verifying software versions

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

```
# lslpp -L VRTS\*
```

Removing and rolling back

Roll back of the 5.1 RP1 to the release 5.1 version is not supported for certain products. It is recommended that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the release 5.1 software. You can roll back 5.1 RP1 to the release 5.1 version for Veritas Cluster Server.

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

- Rolling back 5.1 RP1 to 5.1 for Veritas Cluster Server
- Removing 5.1 RP1 on SF or SFCFS

■ Removing 5.1 RP1 on Storage Foundation for Oracle RAC

Rolling back 5.1 RP1 to 5.1 for Veritas Cluster Server

Use the following procedure to roll back VCS 5.1 RP1 to VCS 5.1 on your cluster manually. To uninstall VCS, see the Veritas Cluster Server Installation Guide.

Note: Use this procedure only when rolling back VCS. Do not roll back VCS when it is part of other products that rely on VCS, for example Storage Foundation Clustered File System or Storage Foundation for Oracle RAC.

To roll back 5.1 RP1

1 Verify that all of the VCS 5.1 RP1 patches are in the APPLIED state. Create a text file called `filesets.to.reject` that contains the name and version of each fileset, one per line, exactly as shown below.

```
VRTScps     5.1.1.0
VRTSgab     5.1.1.0
VRTSllt     5.1.1.0
VRTSvcs     5.1.1.0
VRTSvcsag   5.1.1.0
VRTSvcsea   5.1.1.0
VRTSvxfen   5.1.1.0
```

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

```
# hagrp -state
```

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

```
# hagrp -offline -force ClusterService -any
```

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

```
# haconf -makerw
```

5 Freeze all service groups except the ClusterService service group. On any node, type:

```
# hagrp -freeze $grp -persistent
# hagrp -list | sort -u +0b -1 | \
 while read grp sys ; do
 hagrp -freeze $grp -persistent
 done
```

You can safely ignore the warning about the failure to freeze the ClusterService group.

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

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

7 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
```

8 Shut down VCS. On any node, type:

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

9 Shut down CmdServer. On each node, type:

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

10 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
  ```

  Output for membership for port h does not appear.

- On each node, run the command

```
# ps -ef | egrep "had|hashadow|CmdServer"
```

Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.

- Stop fencing, GAB, and LLT.

```
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/rc.d/rc2.d/S70llt stop
```

11 Preview the patch removal selection and validity tests. On each node, type:

```
# installp -pr -gXv -f filesets.to.reject
```

Confirm that the patches to be removed are exactly the same as those listed in the `filesets.to.reject` file that you created in step 1.

12 Perform the patch removal. On each node, type:

```
# installp -r -gXv -f filesets.to.reject
```

Review the summaries at the end of each run and confirm that all of the intended patches removed successfully.

13 Reboot all nodes in the cluster.

14 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 -list | sort -u +0b -1 | \
 while read grp sys ; do
  hagrp -unfreeze $grp -persistent
     done
# haconf -dump -makero
```

You can safely ignore the warning about the failure to unfreeze the ClusterService group.

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

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

where *system* is the node name.

Removing 5.1 RP1 on SF or SFCFS

You can use the following procedure to uninstall 5.1 RP1 on SF or SFCFS.

To uninstall 5.1 RP1 on SF or SFCFS

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

```
# mount | grep vxfs
```

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

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

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

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

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

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

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

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

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

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

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

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

```
# /etc/rc.d/rc2.d/vcs stop
```

9 Uninstall VCS:

```
# cd /opt/VRTS/install
# ./uninstallvcs [-usersh]
```

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:

```
# genkex | grep odm
# vxkextadm vxodm unload
```

13 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
```

14 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 or Veritas Storage Foundation Cluster File System, 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.

---

15 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 RP1 on Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 5.1 RP1 on Storage Foundation for Oracle RAC systems.

Note: This procedure will remove the complete SF for Oracle RAC stack from all nodes.

To uninstall the 5.1 RP1 on SF Oracle RAC

1 Stop Oracle and CRS on each node of the cluster.

- If CRS is controlled by VCS, log in as superuser on each system in the cluster and enter the following command:

```
# hastop -all
```

- If CRS is not controlled by VCS, enter the following command on each node of the cluster to stop CRS:

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

  Unmount all VxFS file system used by a database or application and enter the following command to each node of the cluster:

  ```
  # hastop -local
  ```

2. Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped. In the `gabconfig -a` command output, the VCS engine or high availability daemon (HAD) port h is not displayed. This indicates that VCS has been stopped.

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

   Sample output:

   ```
   GAB Port Memberships
     ===============================
     Port a gen 5c3d0b membership 01
     Port b gen 5c3d10 membership 01
     Port d gen 5c3d0c membership 01
     Port o gen 5c3d0f membership 01
   ```

3. Uninstall Storage Foundation for Oracle RAC.

   ```
   # cd /opt/VRTS/install
   # ./uninstallsfrac MyNode1 MyNode2
   ```

   See the *Veritas Storage Foundation for Oracle RAC 5.1 Installation and Configuration Guide* for more information.

4. After uninstalling the packages, refer to the Storage Foundation for Oracle RAC 5.1 Installation and Configuration Guide to reinstall the 5.1 software.

Documentation addendum

The following sections contain additions to current documents.

Using the preonline_vvr trigger for RVGLogowner resources

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

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

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

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

  Where the service group is the service group that contains the RVGLogowner resource (*RVGLogowner_resource_sg*). 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.)