Veritas Storage Foundation™ and High Availability Solutions Release Notes

Linux

5.1 Service Pack 1 Rolling Patch 1

Symantec™

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

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

Product version: 5.1 SP1 RP1

Document version: 5.1SP1RP1.2

Legal Notice

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

docs@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Contents

Chapter 1

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

Introduction

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

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

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

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

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

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

Changes introduced in 5.1 SP1 RP1

This section lists the changes in 5.1 SP1 RP1.

Oracle VM Server 2.0 for SPARC

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

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

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

To determine a product's version, use the `-version` option with the `installrp` script. After you install 5.1 SP1 RP1, only the installrp script can detect 5.1 SP1 RP1 versions.

VxVM and ASM co-existence enablement (2194492)

The VxVM and ASM co-existence was disabled by default. This co-existence is now enabled. So VxVM will identify ASM disks and display the same accordingly.

IMF support for DB2 agent

Added Intelligent Monitoring Framework (IMF) capability and support for intelligent resource monitoring for DB2 agent. With IMF, VCS supports intelligent resource monitoring in addition to poll-based monitoring. Poll-based monitoring polls the resources periodically, whereas intelligent monitoring performs asynchronous monitoring. You can enable or disable the intelligent resource monitoring functionality of the agent for DB2.

See the *5.1 SP1 Veritas Cluster Server Administrator’s Guide* for more information about:

- IMF notification module functions
- Administering the AMF kernel driver

Before enabling the agent for IMF

Before you enable the DB2 agent for IMF, ensure that the AMF kernel module is loaded and AMF is configured. For details see the **Administering the AMF kernel driver** section of the *5.1 SP1 Veritas Cluster Server Administration Guide*.

How the DB2 agent supports intelligent resource monitoring

When an IMF-enabled agent starts up, the agent initializes the asynchronous monitoring framework (AMF) kernel driver. After the resource is in a steady state, the agent registers with the AMF kernel driver, the details of the resource that are required to monitor the resource. For example, the agent for DB2 registers the PIDs of the DB2 processes with the IMF notification module. The agent’s ‘imf_getnotification’ function waits for any resource state changes.When the AMF kernel driver module notifies the `imf_getnotification` function about a resource state change, the agent framework runs the **monitor** agent function to ascertain the state of that resource. The agent notifies the state change to VCS, which then takes appropriate action. See the *5.1 SP1 Veritas Cluster Server Administrator’s Guide* for more information.

Agent functions for the IMF functionality

If the DB2 agent is enabled for IMF, the agent supports the following additional functions.

imf_init

This function initializes the DB2 agent to interface with the AMF kernel driver, which is the IMF notification module for the agent for DB2. This function runs when the agent starts up.

imf_getnotification

This function gets notifications about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.

imf_register

This function registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into a steady online state.

Attributes that enable IMF

If the agent for DB2 is enabled for IMF, the agent uses the following type-level attributes in addition to the attributes described in DB2 agent installation and configuration guide.

IMF

This resource type-level attribute determines whether the DB2 agent must perform intelligent resource monitoring. You can also override the value of this attribute at the resource level. This attribute includes the following keys:

Mode

Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows:

- 0—Does not perform intelligent resource monitoring
- 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources
- 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources
- 3—Performs intelligent resource monitoring for both online and for offline resources

Note: The agent for DB2 supports intelligent resource monitoring for online resources only. Hence, Mode should be set to either 0 or 2.

Type and dimension: integer-association

Default: 0

MonitorFreq

This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.

Default: 1

You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:

- After every (*MonitorFreq x MonitorInterval*) number of seconds for online resources

RegisterRetryLimit

If you enable intelligent resource monitoring, the agent invokes the imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes.

Default: 3

IMFRegList

An ordered list of attributes whose values are registered with the IMF notification module.

Type and dimension: string-vector

Default: **static str IMFRegList[]** = { *DB2InstOwner, DB2InstHome* }

Note: In case of an upgrade to VCS 5.1SP1 RP1, please ensure that the new `Db2udbTypes.cf` file is used which contains the definition of **IMFRegList** as above.

System requirements

This section describes the system requirements for this release

Supported Linux operating systems

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

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

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

The Veritas Storage Foundation (SF) 5.1 SP1 RP1 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel), SP3 (2.6.16.60-0.54.5) or SP4 (2.6.16.60-0.75.1), on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) with SP1 (2.6.32.12-0.7 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)

  Note: SFCFSRAC is not supported on SLES 11 SP1.

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

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

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

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

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

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

VMware Environment

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

Note: This TechNote includes information specific to all 5.1 releases. Please check this technote for the latest information.

Supported VCS agents

Table 1-1 lists the agents for enterprise applications and the software that the agents support.

Table 1-1 Supported software for the VCS agents for enterprise applications

Agent	Application	Application version	version
DB2	DB2 Enterprise Server Edition	8.1, 8,2, 9.1, 9.5, 9.7	RHEL 5, SLES 10, SLES 11, OEL5
Oracle	Oracle	11gR1, 10gR2, 11gR2	RHEL 5, SLES10, SLES 11, OEL5
Sybase	Sybase Adaptive Server Enterprise	12.5.x, 15.x	RHEL 5, SLES 10, SLES 11, OEL5

See the *Veritas Cluster Server Installation Guide* for the agent for more details.

For a list of the VCS application agents and the software that the agents support, see the Veritas Cluster Server Agents Support Matrix at Symantec website.

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: Veritas Storage Foundation (SF) and Veritas Storage Foundation Cluster File System (SFCFS) support running Oracle, DB2, and Sybase on VxFS and VxVM.

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

Recommended memory and swap space

Symantec recommends the following memory and swap space sizes:

- On the system where you run the installation, use the following guidelines for memory minimums when you install on:
  - One to eight nodes, use 1 GB of memory
  - More than eight nodes, use 2 GB of memory or more
- On the system where you run the installation, use the following guidelines for swap space when you install on:
  - One to eight nodes, use (*number of nodes* + 1) x 128 MB of free swap space
  - More than eight nodes, 1 GB of free swap space

List of products

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

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

Fixed issues

This section describes the issues fixed in this release.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP1

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

Table 1-2 Veritas Volume Manager 5.1 SP1 RP1 fixed issues

Fixed issues	Description
1426480	The `volcvm_clear_pr`() ioctl now propogates the error returned by DMP to the caller.
1829285	`vxconfigd` no longer dumps core while assigning a unique native name to a disk.
1869002	Introduced a Circular buffer at the vold level for master-slave communication.
1940052	`vxconfigd` no longer hangs on the master after removing the HBA alias from the zone and node leave followed by join
1959513	The `-o noreonline` option of a diskgroup import is now propogated to slave nodes.
1970560	vxconfigd no longer dumps core on the master node when vxconfigd on a passive slave dies and command shipping is in progress.
2015467	Improved performance for NetBackup 6.5.5 on Veritas Storage Foundatoin 5.1 VxVM mapping provider.
2038928	Added support for creating and using pre-5.1 SP1 release diskgroups on CDS-initialized disks.
2062190	vxrootadm : split/join operation fails when there is a rvg present in the rootdg/backupdg
2080730	The vxvm exclude file and vxdmp exclude file contents are now consistent after updating the files using the `vxdiskadm` command and `vxdmpadm` command.
2082450	vxdisk resize should output more meaningful error message
2088007	possibility of reviving only secondary paths in dmp_revive_paths()
2105547	tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations
2125306	Fixed a few issues related to loading the HBA API library and the `vxinstall` script.
2129477	`vxdisk reclaim` no longer fails after resize.
2129989	EVA ASL should report an error message if pref_bit is not set for a LUN

Table 1-2 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2133503	Renaming enclosure results in dmpevents.log reporting 'mode for Enclosure has changed from Private to Private'
2148682	while shipping a command node hangs in master selection on slave nodes and master update on master node
2152830	In a multi-level clone disks environment, a regular diskgroup import is now handled properly, and in the case of a diskgroup import failure, the correct error message is now displayed.
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core.
2160199	An upcoming master can now import a shared diskgroup, which allows the master takeover to succeed.
2166682	checks needed to make sure that a plex is active before reading from it during fsvm mirror read interface
2172488	FMR2 restore doesn't sync the existing snapshot mirrors
2179479	The flags on a disk are no longer incorrectly set as "error" even after running the `vxdisk scandisks` command after creating a PV and volume group.
2181631	striped-mirror volume cannot be grown across sites with -oallowspansites w/ DRL
2183984	system panic in dmp_update_stats() routine
2188590	an ilock acquired by a SLAVE node for a read on a DCL object can lead to IO hang when the node becomes MASTER before reading is done
2191693	`vxdmpadm native list` command now displays output and error messages.
2194492	VxVM-ASM co-existence enablement
2199496	Fixed a data corruption issue with the "site mirror" Campus Cluster feature.
2199836	The system with the root volume group and DMP native support enabled now successfully boots and mounts.
2200670	The `vxattachd` command can now recover disks if even if the disk group is not imported.

Table 1-2 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault
2215216	`vxkprint` now reports TP-related values.
2220926	The `vxprivutil -D set` *`attr`* command no longer causes the `vxprivutil` command to hang.
2226813	Rlinks no longer remain disconnected with the UDP protocol if data ports are specified.
2227923	Renaming an enclosure is now persistent.
2234844	An `asm2vxfs` conversion with Linux partitions no longer fails.

Veritas File System: Issues fixed in 5.1 SP1 RP1

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

Table 1-3 Veritas File System fixed issues

Incident	Description
1296491	Fixed issues seen during a force unmount of a parent cluster file system while a child was being mounted or umounted.
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.
2030119	fsppadm core dumps when analysing a badly formatted XML file, is resolved
2032525	Fixed the cause of an NFS stale file handle.
2061554	Sequential extents are now collated.
2111921	Improved the performance of VxFS file systems with concurrent I/O or direct I/O enabled.
2149659	Fixed the cause of an error that resulted during the truncate operation of a file with a shared extent in the presence of a Storage Checkpoint containing FileSnaps.

Table 1-3 Veritas File System fixed issues *(continued)*

Incident	Description
2162822	During online migration from ufs to vxfs, df command returns a non-zero return value.
2163084	The `listxattr`() call now uses rwlock.
2169273	During online migration, nfs export of the migrating file system leads to system panic
2177253	A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for vxupgrade purposes
2178147	Linking a IFSOC file now properly calls `vx_dotdot_op`(), which fixes the cause of a corrupted inode.
2181833	The `vxfilesnap` command no longer gives an incorrect error message on a Storage Checkpoint file system.
2184528	`fsck` no longer fails to repair corrupt directory blocks that have duplicate directory entries.
2194618	Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system
2198553	A forced unmount now properly clears the `bd_super` structure member.
2221623	Fixed a performance loss due to a delxwri_ilist spin lock with the default values for `vx_idelxwri_timelag`.
2226257	Fixed the cause of a system panic in the `in _list_add`() call, which led to corruption in the `vx_ftenter`() codepath when using named data streams.

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

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

Table 1-4 Veritas Storage Foundation Cluster File System fixed issues

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

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

There are no Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues in this release.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP1

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

Table 1-5 Veritas Cluster Server 5.1 SP1 RP1 fixed issues

Fixed Issues	Description
1949294	`fdsetup` can now correctly parse disk names containing characters such as "-".
1949303	`fdsetup` no longer allows volume that are not part of the RVG, which fixes a possible cause of the RVGSnapshot agent failing.
2011536	Added IMF support for the db2udb agent.
2159991	Fixed an issue with messages in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2172181	Fixed an issue with AMF-related messages for the CAVF agent in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2179652	The monitor script of the db2udb agent can now handle empty attribute values.
2184205	Fixed an issue with HAD in which the parent service group did not fail over if the parent service group had an online local firm dependency with a child service group.
2194473	HAD no longer dumps core while overriding the static attribute to the resource level.
2205556	Fixed an issue with the offline EP of the DNS agent, which did not remove all A/AAAA records if OffDelRR=1 for multi-home records.
2205563	A clean EP now properly removes resource records when OffDelRR=1.
2205567	Fixed an issue in which having an attribute set to `master.vfd` caused the DNS agent to fail to query the DNS server.
2208675	There is now a return value check for broadcast ping in NIC/MultiNICA monitor, which fixes one possible cause of the MultiNic resource is going into the FAULTED state in IPv6 with the Link option configuration.
2208901	Fixed an issue with the RVGSnapshot agent.
2209337	Fixed an issue with VCSAPI where the RemoteGroup agent crashed if the VCSAPI log level was set to a non-zero value.
2214539	Fixed an issue in which rebooting a node sometimes set the intentonline of a group to 2, even if the group was online somewhere else. This caused the group to use the autostartlist and not perform a failover.

Table 1-5 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Fixed Issues	Description
2217446	Fixed an issue that caused the installation of `VRTSvcsag` to fail.
2218556	Fixed an issue in the `cpsadm` command in which it sometimes failed if LLT was not installed or configured on a single node cluster.
2218561	Fixed an issue in which MonitorTimeStats incorrectly showed 303 seconds intermittently.
2219955	Fixed an issue in which a split-brain condition occurred even when using VCS Steward.
2220749	Fixed an issue in which the Cluster Manager (Java Console) was not encrypting the `DBAPword` attribute of the Oracle Agent.
2241397	Fixed an issue in which halogin did not work in a secure environment where the root broker was not a VCS node.

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

This section describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in this release.

Table 1-6 Storage Foundation for Databases fixed issues

Incident	Description
2203917	Process table has been changed to use per-hash-bucket locks, and the number of buckets has been increased from 32 to 256.
2237709	The `dbdst_preset_policy` command no longer aborts when you specify the volume class as MEDIUM.

Known issues

This section covers the known issues in this release.

Issues related to installation

This section describes the known issues during installation and upgrade.

EULA changes (2161557)

The locations for all EULAs have changed.

The English EULAs now appear in /*product_dir*/EULA/en/*product_eula.pdf*

The EULAs for Japanese and Chinese now appear in those language in the following locations:

The Japanese EULAs appear in /*product_dir*/EULA/ja/*product_eula.pdf*

The Chinese EULAs appear in /*product_dir*/EULA/zh/*product_eula.pdf*

NBU 6.5 or older version is installed on a VxFS file system (2056282)

NBU 6.5 or older version is installed on a VxFS file system. Before upgrading to SF5.1, user umounts all VxFS file systems including the one which hosts NBU binaries(/usr/openv). While upgrading SF5.1, installer fails to check out NBU is installed on the same machine and uninstalls the shared infrastructure packages: VRTSpbx, VRTSat, VRTSicsco which cause NBU not working.

Workaround: Before you umount the VxFS file system which hosts NBU, copy the two files / usr/openv/netbackup/bin/version and /usr/openv/netbackup/version to /tmp directory. After you umount the NBU file system, manually copy these two version files from /tmp to their original path. If the path doesn’t exist, make the same directory path with the command: `mkdir -p /usr/openv/netbackup/bin` and `mkdir -p /usr/openv/netbackup/bin`. Run the installer to finish the upgrade process. After upgrade process is done, remove the two version files and their directory paths.

How to recover systems already affected by this issue: Manually install VRTSpbx, VRTSat, VRTSicsco packages after the upgrade process is done.

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

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

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

The VRTSacclib is deprecated (2032052)

The VRTSacclib is deprecated. For installation, uninstallation, and upgrades, note the following:

- Fresh installs: Do not install VRTSacclib.
- Upgrade: Ignore VRTSacclib.
- Uninstall: Ignore VRTSacclib.

Error messages in syslog (1630188)

If you install or uninstall a product on a node, you may see the following warnings in syslog: /var/log/message. These warnings are harmless and can be ignored.

```
Jul  6 10:58:50 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:54 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:59 swlx62 setroubleshoot: SELinux is preventing the
restorecon from using potentially mislabeled files
```

Ignore certain errors after an operating system upgrade—after a product upgrade with encapsulated boot disks (2030970)

Ignore certain errors after an operating system upgrade after a product upgrade with encapsulated boot disks.

You can ignore the following errors after you upgrade the operating system after a product upgrade that occurred with an encapsulated boot disk. Examples of the errors follow:

```
The partioning on disk /dev/sda is not readable by
The partioning tool parted, which is used to change the
partition table.
You can use the partitions on disk /dev/sda as they are.
You can format them and assign mount points to them, but you
cannot add, edit, resize, or remove partitions from that
disk with this tool.
```

Or

```
Root device: /dev/vx/dsk/bootdg/rootvol (mounted on / as reiserfs)
Module list: pilix mptspi qla2xxx silmage processor thermal fan
reiserfs aedd (xennet xenblk)

Kernel image; /boot/vmlinuz-2.6.16.60-0.54.5-smp
Initrd image: /boot/initrd-2.6.16.60-0.54.5-smp
```

The operating system upgrade is not failing. The error messages are harmless.

Workaround: Remove the /boot/vmlinuz.b4vxvm and /boot/initrd.b4vxvm files (from an un-encapsulated system) before the operating system upgrade.

Veritas Storage Foundation known issues

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

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.

db2exp may frequently dump core (1854459)

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

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

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

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

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

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

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

DBI1070I  Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

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

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

To communicate in a dual-stack environment

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

  ```
  127.0.0.1 swlx20-v6
  ```

  Or

  ```
  127.0.0.1 swlx20-v6.punipv6.com
  ```

  where *`127.0.0.1`* is the IPv4 loopback address.

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

Process start-up may hang during configuration using the installer (1678116)

After you have installed a Storage Foundation product, some VM process may hang during the configuration phase.

Workaround: Kill the installation program, and rerun the configuration again.

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

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

Tools like `dbca` may hang during database creation.

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

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

After upgrading SF stack from 5.0MP3SP1 RP1 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the SFM GUI.

Workaround:

To resolve this known issue

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

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

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

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

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

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

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

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

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

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

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

To resolve this issue

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

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

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

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

Veritas Volume Manager known issues

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

The vxcdsconvert utility is not supported for EFI disks (2064490)

The vxcdsconvert utility is not supported for EFI disks.

Post encapsulation of the root disk, system comes back up after first reboot unencapsulated (2119038)

In some cases, after encapsulating the root disk and rebooting the system, it may come up without completing the encapsulation. This happens because the `vxvm-reconfig` startup script is unable to complete the encapsulation process.

Workaround

Reboot the system or to run the following command.

```
# service vxvm-reconfig start
```

This will reboot the system and complete the remaining stages of encapsulation.

Required attributes for Veritas Volume Manager (VxVM) devices to avoid boot failures (1411526)

To support iSCSI devices, Veritas Volume Manager (VxVM) does not start non-root devices until runlevel2. The boot process expects all local (non-NFS) mount points in the `/etc/fstab` file to be present at boot time. To avoid boot failures, all VxVM entries in the `/etc/fstab` file must have the _netdev attribute, and must not have the fsck required flag set. These attributes enable VxVM to defer mounting of VxVM devices until after VxVM has started.

vxrestored daemon fails to restore disabled paths (1663167)

The `vxrestored` daemon fails to restore disabled paths on RHEL 5 with direct attached disks.

Workaround:

Enable the `mpt_disable_hotplug_remove` tunable so that path level failover and failback function properly on RHEL 5 machines with direct attached disks.

To enable the mpt_disable_hotplug_remove tunable

1 Edit the /etc/modprobe.conffile and add the following line to the end of the file:

```
options mptsas mpt_disable_hotplug_remove=0
```

2 Rebuild the initrd image:

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

3 Reboot the system.

System hangs or panics after disabling 3 of 4 arrayside ports (1724260)

The system hangs or panics after you disable 3 of 4 arrayside ports.

Workaround:

This issue is fixed with a Novell patch for SLES 11 as indicated in Bugzilla ID 524347:

https://bugzilla.novell.com/show_bug.cgi?id=524347

Machine fails to boot after root disk encapsulation on servers with UEFI firmware (1842096)

Certain new servers in the market such as IBM x3650 M2, Dell PowerEdge T610, come with support for the UEFI firmware. UEFI supports booting from legacy MBR type disks with certain restrictions on the disk partitions. One of the restrictions is that each partition must not overlap with other partitions. During root disk encapsulation, it creates an overlapping partition that spans the public region of the root disk. If the check for overlapping partitions is not disabled from the UEFI firmware, then the machine fails to come up following the reboot initiated after running the commands to encapsulate the root disk.

Workaround:

The following workarounds have been tested and are recommended in a single-node environment.

For the IBM x3650 series servers, the UEFI firmware settings should be set to boot with the "Legacy Only" option.

For the Dell PowerEdge T610 system, set "Boot Mode" to "BIOS" from the "Boot Settings" menu.

Root disk encapsulation issue (1603309)

Encapsulation of root disk will fail if it has been assigned a customized name with vxdmpadm(1M) command. If you wish to encapsulate the root disk, make sure that you have not assigned a customized name to its corresponding DMP node. See vxdmpadm(1M) and the section "Setting customized names for DMP nodes" on page 173 for details.

VxVM starts before OS device scan is done (1635274)

While working with some arrays, VxVM may start before all devices are scanned by the OS. This slow OS device discovery may result in malfunctioning of VM, fencing and VCS due to partial disks seen by VxVM.

Workaround:

After the fabric discovery is finished, issue the `vxdisk scandisks` command to bring newly discovered devices into the VxVM configuration.

vxdisk -f init can overwrite some of the public region contents (1190117)

If a disk was initialized by a previous VxVM version or defined with a smaller private region than the new default of 32Mb, then the public region data will be overridden. The work around is to specify explicitly the length of privoffset, puboffset, publen, privlen while initializing the disk.

The relayout operation fails when there are too many disks in the disk group. (2015135)

The attempted relayout operation on diskgroup containing approximately more than 300 Luns/disks may fail with error 'Cannot setup space

Converting LVM volumes to VxVM volumes fails when multipathed storage devices are present (1471781, 1931727)

The `vxvmconvert` utility cannot convert LVM volumes to VxVM volumes when multipathed storage devices are present. This issue occurs with SLES 11 and RHEL5, due to changes in the LVM utilities. If multipathed devices are detected, the `vxvmconvert` utility exits with the following error:

```
vxvmconvert cannot convert multipathed devices on SLES11 systems.
... Exiting.
```

Workaround: There is no workaround for this issue.

Co-existence check might fail for CDS disks

In Veritas Volume Manager (VxVM) 5.1 SP1 RP1, VxVM introduces the ability to support Cross-platform Data Sharing (CDS) on disks larger than 1 TB. VxVM uses the SUN VTOC Table to initialize the cdsdisk layout on devices up to 1 TB. VxVM uses the GUID Partition Table (GPT) to initialize the cdsdisk layout on devices larger than 1 TB.

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

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

Workaround: There is no workaround for this issue.

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

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

Node is not able to join the cluster with high I/O load on the array with VCS (2124595)

When the array has a high I/O load, the DMP database exchange between master node and joining node takes a longer time. This situation results in VCS resource online timeout, and then VCS stops the join operation.

Workaround:

Increase the online timeout value for the HA resource to 600 seconds. The default value is 300 seconds.

To set the OnlineTimeout attribute for the HA resource type CVMCluster

1 Make the VCS configuration to be ReadWrite:

```
# haconf -makerw
```

2 Change the OnlineTimeout attribute value of CVMCluster:

```
# hatype -modify CVMCluster OnlineTimeout 600
```

3 Display the current value of OnlineTimeout attribute of CVMCluster:

```
# hatype -display CVMCluster -attribute OnlineTimeout
```

4 Save and close the VCS configuration:

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

Changes in enclosure attributes are not persistent after an upgrade to VxVM 5.1 SP1 RP1 (2082414)

The Veritas Volume Manager (VxVM) 5.1 SP1 RP1 includes several array names that differ from the array names in previous releases. Therefore, if you upgrade from a previous release to VxVM 5.1 SP1 RP1, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 5.1 SP1 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-7 shows the Hitachi arrays that have new array names.

Table 1-7 Hitachi arrays with new array names

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

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. The persistence of the enclosure name is achieved with the `/etc/vx/array.info` file, which stores the mapping between cabinet serial number and array name. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 5.1 SP1 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

DMP disables subpaths and initiates failover when an iSCSI link is failed and recovered within 5 seconds. (2100039)

When using iSCSI S/W initiator with an EMC CLARiiON array, iSCSI connection errors may cause DMP to disable subpaths and initiate failover. This situation occurs when an iSCSI link is failed and recovered within 5 seconds.

Work around:

When using iSCSI S/W initiator with an EMC CLARiiON array, set the node.session.timeo.replacement_timeout iSCSI tunable value to 40 secs or higher.

Device discovery failure during VxVM startup in SUSE (2122771)

When the system boots up, some devices are not displayed in the `vxdisk list` output. This issue occurs if the vold daemon does the device discovery before the operating system (OS) completes its device discovery. Therefore, the vold daemon may miss some devices.

Work-around:

Configure the vxvm-startup script to wait until the operating system discovery is completed before starting vold. In the /etc/vx/vxvm-startup file, set DEV_DISCOVER_DELAY to the expected device discovery time taken by the OS. By default, DEV_DISCOVER_DELAY is set to 0.

You must reboot the system before this configuration applies. To discover the missed devices, run the `vxdisk scandisks` command or the `vxdctl enable` command.

DMP marks the subpaths as DISABLED while these subpaths are accessible from OS level (2037222)

For iSCSI devices on SLES 10 SP3, the DMP tunable parameter dmp_fast_recovery needs to be turned off.

```
# vxdmpadm settune dmp_fast_recovery=off
```

DMP panics if a DDL device discovery is initiated immediately after loss of connectivity to the storage (2040929)

When using EMC Powerpath with VxVM 5.1SP1 RP1 on SLES11, set the fast_io_fail_tmo on the HBA port to any non-zero value that is less than the dev_loss_tmo value so as to avoid a panic in case a DDL device discovery is initiated by the `vxdisk scandisks` command or the `vxdctl enable` command immediately after loss of connectivity to the storage.

EFI labelled disks of size greater than 2TB size; will not be supported. (2270880)

On Solaris 10, if the size of EFI labelled disk is greater than 2TB, the disk capacity will be truncated to 2TB when it is initialized under Veritas Volume Manager.

Workaround: There is no workaround for this issue.

Veritas Volume Replicator known issues

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

vradmin syncvol command compatibility with IPv6 addresses (2075307)

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

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

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

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

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname could not be
imported on bunker host hostname. Operation failed with error 256
and message VxVM VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote
server unreachable... Timestamp VCS ERROR V-16-2-13066 (hostname)
```

```
Agent is calling clean for resource(RVGPrimary) because the resource
is not up even after online completed.
```

Workaround:

To resolve this issue

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

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

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

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

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

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

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

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

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround: The following workarounds resolve these issues.

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

For example:

```
# fsck -t vxfs /dev/vx/dsk/dg/snapshot_volume
```

For issue 2, run the `fsck` command on the replicated data volumes on the new Primary site, to restore the consistency of the file system residing on the data volume.

For example:

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

Running SUSE Linux and using Novell's YaST tool to configure an IPv6 address may result in an error

When Novell's YaST tool is invoked to configure an IPv6 address on a different network interface and if:

- the host name, the DNS server name and domain name are specified to the YaST tool.

- IPv6 address is assigned by the Dynamic Host Configuration Protocol (DHCP).
- the "Write Hostname to /etc/hosts" option is selected (this is selected by default).

This results in the `vradmin` command returning the following error:

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

This happens because the YaST tool can replace the `/etc/hosts` entry containing `127.0.0.2` from the IPv4 host name to the specified new IPv6 host name. For example:

```
127.0.0.2 v6hostname.space.ipv6.com v6hostname
```

Workaround: The following procedure resolves this issue.

To resolve this issue

1. Edit the `/etc/hosts` file to specify the correct IPv6 address.
2. Restart the `vradmind` daemon on all VVR hosts:

```
# /etc/init.d/vras-vradmind.sh restart
```

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

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

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

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

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

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

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

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

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

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

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

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

```
# /etc/init.d/vras-vradmind.sh restart
```

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

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

Workaround:

To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

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

Workaround: There is no workaround for this issue.

vxassist relayout removes the DCM (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

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

3. Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

Veritas File System known issues

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

Storage Checkpoints can exceed the quota limit (2102201)

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

- The Storage Checkpoint quota has been enabled.
- The Storage Checkpoint quota is not exceeded.

- A file content modification operation, including removing a file, needs to push some or all blocks of the file to the Storage Checkpoint.
- Number of blocks that need to be pushed to the Storage Checkpoint is enough to exceed Storage Checkpoint quota hard limit.

Workaround: There is no workaround for this issue.

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

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

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

A null pointer dereference in the `vx_unlockmap`() call can cause a panic. A fix for this issue will be released in the a future patch.

Workaround: There is no workaround for this issue.

VxFS module loading fails when freevxfs module is loaded (1736305)

The following module loading error can occur during RPM installation if the `freevxfs` module is loaded:

```
Error in loading module "vxfs". See documentation.

ERROR: No appropriate VxFS drivers found that can be loaded.
See VxFS documentation for the list of supported platforms.
```

Workaround: Ensure that the `freevxfs` module is not loaded before installing the `VRTSvxfs` RPM. The following command shows if the `freevxfs` module is loaded:

```
# lsmod | grep freevxfs
```

If the `freevxfs` module is loaded, unload the module:

```
# rmmod freevxfs
```

A mount can become busy after being used for NFS advisory locking (1508386)

If you export a VxFS file system using NFS and you perform file locking from the NFS client, the file system can become unable to be unmounted. In this case, the `umount` command fails with the `EBUSY` error.

Workaround: Force unmount the file system:

```
# vxumount -o force /mount1
```

umount can hang when inotify watches are used (1590324)

If inotify watches are used, then an unmount can hang in the `vx_softcnt_flush`() call. The hang occurs because inotify watches increment the `i_count` variable and cause the `v_os_hold` value to remain elevated until the inotify watcher releases the hold.

Workaround: There is no workaround for this issue.

5.1SP1 sol_sprac Test LM-stress_S5 hits an assert of "f:vx_idelxwri_off:5a vai vx_trunc_tran"(2169326)

When a removable clone mounted in “rw” mode, and if files from the mounted clone are being modified, operations on the mounted clone may fail if the file system is full. Even though the clone will continue to be seen mounted, no operation except an umount will be allowed. Also `fsckptadm list` *`primary mount point`* will not show the clone right away.

Workaround: It is advised to create non-removable clones to avoid encountering this issue.

Veritas Cluster Server known issues

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

Stale NFS file handle on the client across failover of a VCS service group containing LVMLogicalVolume resource

A VCS service group for a LVM volume group will be online automatically after a failover. However, the client applications may fail or be interrupted by stale NFS file handle error.

Workaround: To avoid the stale NFS file handle on the client across service group failover, specify "fsid=" in the Options attribute for Share resources.

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

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

Workaround:

1. Set VCS_REMOTE_BROKER to the remote AB:

   ```
   # export VCS_REMOTE_BROKER=remote_broker
   ```

2. Set VCS_DOMAIN and VCS_DOMAINTYPE:

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

3. Run halogin:

   ```
   # halogin ldap_user
   ```

   Provide password when prompted.

4. Unset VCS_DOMAIN and VCS_DOMAINTYPE:

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

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

NFS cluster I/O fails when storage is disabled

The I/O from the NFS clusters are saved on a shared disk or a shared storage. When the shared disks or shared storage connected to the NFS clusters are disabled, the I/O from the NFS Client fails and an I/O error occurs.

Workaround: If the application exits (fails/stops), restart the application.

Issues related to installation

This section describes the known issues during installation and upgrade.

While configuring authentication passwords through the Veritas product installer, the double quote character is not accepted (1245237)

The Veritas product installer prompts you to configure authentication passwords when you configure Veritas Cluster Server (VCS) as a secure cluster, or when you configure Symantec Product Authentication Service (AT) in authentication broker (AB) mode. If you use the Veritas product installer to configure authentication passwords, the double quote character (\") is not accepted. Even though this special character is accepted by authentication, the installer does not correctly pass the characters through to the nodes.

Workaround: There is no workaround for this issue. When entering authentication passwords, do not use the double quote character (\").

Errors observed during partial upgrade of SFHA

While upgrading the VCS packages during an SFHA upgrade from 5.0 MP3 RP2 to 5.1SP1, CPI failed to uninstall the I/O fencing packages (VRTSvxfen, VRTSllt, and VRTSgab). [1779129]

Workaround

Before upgrading SFHA from 5.0 MP3 RP2 to 5.1SP1, you must apply the I/O fencing hotfix 5.0MP3RP2HF2.

While upgrading the VCS stack from a version prior to VCS 5.1, reconfiguration of MultiNICA IPv4RouteOptions attribute is required

The 5.1SP1 MultiNICA agent now uses `ip` command by default. Due to behavioral differences in `ip` and `ifconfig` commands in regards to route configuration, MultiNICA flushes routes and sets them back for the new active device. If the MultiNICA resource configuration is not intended to make use of `ifconfig` command (see table below), you must configure IPv4RouteOptions attribute in MultiNICA resource definition.

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

Table 1-8 Whether attributes are configured and required actions that you need to perform during upgrade

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

Issues with keyless licensing reminders after upgrading VRTSvlic

After upgrading from 5.1 to 5.1SP1, some keyless licenses may be left in the system. As a result, you may see periodic reminders being logged if the VOM server is not configured.

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

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.
2. Set the product level to *NONE* with the command:

   ```
   # vxkeyless set NONE
   ```

3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic`:
   - Verify if the key has `VXKEYLESS` feature Enabled using the following command:

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

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

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

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

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

SELinux error during installation of VRTSvcsag package

Description: During the installation of VRTSvcsag package on RHEL 5 SELinux enabled machine, you may observe following SELinux error:

```
/usr/sbin/semodule:  Failed on /opt/VRTSvcs/bin/selinux/vcsag.pp!
```

This error occurs due to improper installation of the SELinux package. As a result, SELinux commands may not function properly.

Workaround: Reinstall the SELinux package, if you observe this issue.

Installer is unable to split a cluster that is registered with one or more CP servers

Splitting a cluster that uses server-based fencing is currently not supported. [2110148]

You can split a cluster into two and reconfigure on the two clusters using the installer. For example, you can split a cluster *clus1* into *clus1A* and *clus1B*.

However, if you use the installer to reconfigure the , the installer retains the same cluster UUID of *clus1* in both *clus1A* and *clus1B*. If both *clus1A* and *clus1B* use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: None.

Issues related to the VCS engine

Systems with multiple CPUs and copious memory shut-down time may exceed the ShutdownTimeout attribute

The time taken by the system to go down may exceed the default value of the ShutdownTimeout attribute for systems that have a large numbers of CPUs and memory. [1472734]

Workaround: Increase the value of the ShutdownTimeout attribute based on your configuration.

VCS Engine logs messages when it eventually connects to a remote cluster

Description: In a global cluster, if a local cluster fails to connect with a remote cluster in the first attempt but succeeds eventually, then you may see the following warning messages in the engine logs. [2110108]

```
VCS WARNING V-16-1-10510 IpmHandle: pen Bind Failed.
unable to bind to source address 10.209.125.125. errno = 67
```

Workaround: There is currently no workaround for this issue. This issue has no impact on any functionality.

Agent framework can reject `hares -action` command

When a probed resource is disabled and later enabled then, the agent framework can reject `hares -action` command till the agent successfully monitors the resource.

New nodes get added to SystemList and AutoStartList attributes of ClusterService even if AutoAddSystemToCSG is disabled

The AutoAddSystemToCSG attribute determines whether the newly joined or added systems in a cluster become part of the SystemList of the ClusterService service group if the service group is configured. The value 1 (default) indicates that the new systems are added to SystemList of ClusterService. AutoAddSystemToCSG has an impact only when you execute the hasys –add command or when a new node joins the cluster. [2159139]

However, when you use the installer to add a new node to the cluster, the installer modifies the SystemList and AutoStartList attributes irrespective of whether AutoAddSystemToCSG is enabled or disabled. The installer adds the new system to the SystemList and AutoStartList. To add nodes, the installer uses the following commands that are not affected by the value of AutoAddSystemToCSG:

```
# hagrp –modify ClusterService SystemList –add newnode n
# hagrp –modify ClusterService AutoStartList –add newnode
```

Workaround

The installer will be modified in future to prevent automatic addition of nodes to SystemList and AutoStartList.

As a workaround, use the following commands to remove the nodes from the SystemList and AutoStartList:

```
# hagrp –modify ClusterService SystemList –delete newnode
# hagrp –modify ClusterService AutoStartList –delete newnode
```

Issues related to the agent framework

Agent may fail to heartbeat under heavy load (2073018)

Description: An agent may fail to heart beat with the VCS engine under heavy load.

This may happen when agent does not get enough CPU to perform its tasks and when the agent heartbeat exceeds the time set in the AgentReplyTimeout attribute. The VCS engine therefore stops the agent and restarts it. The VCS engine generates the following log when it stops and restarts the agent:

Resolution: If you are aware that the system load is likely to be high, then:

- The value of AgentReplyTimeout attribute can be set to a high value
- The scheduling class and scheduling priority of agent can be increased to avoid CPU starvation for the agent, using the AgentClass and AgentPriority attributes.

Agent framework cannot handle leading and trailing spaces for the dependent attribute

Agent framework does not allow spaces in the target resource attribute name of the dependent resource.

Workaround

Do not provide leading and trailing spaces in the target resource attribute name of the dependent resource.

The agent framework does not detect if service threads hang inside an entry point

In rare cases, the agent framework does not detect if all service threads hang inside a C entry point. In this case it may not cancel them successfully. [1511211]

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

The ArgListValues attribute values for dependent resources may not populate correctly when a target resource is deleted and re-added

For resource attributes, deleting a resource prevents a dependent attribute’s value from refreshing in the dependent resource’s value.

For example, you have resource (*rD*), which depends on a resource’s attribute value (*rT:Attr_rt*). When you delete the target resource (*rT*), and re-add it (*rT*), the dependent resource (*rD*) does not get the correct value for the attribute (*Attr_rt*). [1539927]

Workaround: Set the value of the reference attribute (*target_res_name*) to an empty string.

```
# hares -modify rD target_res_name ""
```

Where *rD* is the name of the dependent resource, and *target_res_name* is the name of the reference attribute that contains the name of the target resource.

Set the value of the reference attribute (*target_res_name*) to the name of the target resource (*rT*).

```
# hares -modify rD target_res_name rT
```

Agent performance and heartbeat issues

Depending on the system capacity and the number of resources configured under VCS, the agent may not get enough CPU cycles to function properly. This can prevent the agent from producing a heartbeat synchronously with the engine. If you notice poor agent performance and an agent's inability to heartbeat to the engine, check for the following symptoms.

Navigate to /var/VRTSvcs/diag/agents/ and look for files that resemble:

```
FFDC_AGFWMain_729_agent_type.log  FFDC_AGFWTimer_729_agent_type.log core
FFDC_AGFWSvc_729_agent_type.log   agent_typeAgent_stack_729.txt
```

Where *agent_type* is the type of agent, for example Application or FileOnOff. If you find these files, perform the next step.

Navigate to /var/VRTSvcs/log/ and check the engine_*.log file for messages that resemble:

```
2009/10/06 15:31:58 VCS WARNING V-16-1-10023 Agent agent_type
not sending alive messages since Tue Oct 06 15:29:27 2009
2009/10/06 15:31:58 VCS NOTICE V-16-1-53026 Agent agent_type
ipm connection still valid
2009/10/06 15:31:58 VCS NOTICE V-16-1-53030 Termination request sent to
agent_type agent process with pid 729
```

Workaround: If you see that both of the above criteria are true, increase the value of the AgentReplyTimeout attribute value. (Up to 300 seconds or as necessary.) [1853285]

Issues related to global clusters

The engine log file receives too many log messages on the secure site in global cluster environments

When the WAC process runs in secure mode on one site, and the other site does not use secure mode, the engine log file on the secure site gets logs every five seconds. [1539646]

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

Application group attempts to come online on primary site before fire drill service group goes offline on the secondary site

The application service group comes online on the primary site while the fire drill service group attempts to go offline at the same time, causing the application group to fault. (2107386)

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

Issues related to LLT

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

LLT port stats sometimes shows recvcnt larger than recvbytes

With each received packet, LLT increments the following variables:

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

Both these variables are integers. With constant traffic, recvbytes hits and rolls over MAX_INT quickly. This can cause the value of recvbytes to be less than the value of recvcnt. [1788315]

This does not impact the LLT functionality.

LLT may incorrectly declare port-level connection for nodes in large cluster configurations

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This occurs in some corner cases even though a port is not even registered on the peer node. [1809827]

Issues related to I/O fencing

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

All nodes in a sub-cluster panic if the node that races for I/O fencing panics

At the time of a network partition the lowest node in each sub-cluster races for the coordination points on behalf of that sub-cluster. If the lowest node is unable to contact a majority of the coordination points or the lowest node itself unexpectedly panics during the race, then all the nodes in that sub-cluster will panic. [1965954]

Coordination Point agent does not provide detailed log message for inaccessible CP servers

The Coordination Point agent does not log detailed information of the CP servers that are inaccessible. When CP server is not accessible, the agent does not mention the UUID or the virtual IP of the CP server in the engine log. [1907648]

Preferred fencing does not work as expected for large clusters in certain cases

If you have configured system-based or group-based preferred fencing policy, preferred fencing does not work if all the following cases are true:

- The fencing setup uses customized mode with one or more CP servers.
- The application cluster has more than eight nodes.
- The node weight for a single node (say galaxy with node id 0) is more than the sum total of node weights for the rest of the nodes.
- A network fault occurs and the cluster partitions into two with the single node (galaxy) on one part and the rest of the nodes on the other part.

Under such circumstances, for group-based preferred fencing, the single node panics even though more high priority services are online on that node. For system-based preferred fencing, the single node panics even though more weight is assigned to the node. [2161816]

See the *Veritas product Administrator's Guide* for more information on preferred fencing.

Server-based I/O fencing fails to start after configuration on nodes with different locale settings

On each (application cluster) node, the vxfen module retrieves and stores the list of the UUIDs of coordination points. When different nodes have different locale settings, the list of UUIDs on one (application) node does not match with that of the other (application) nodes. Hence, I/O fencing does not start after configuration. [2112742]

Workaround: Start I/O fencing after fixing the locale settings to use the same values on all the (application) cluster nodes.

Reconfiguring with I/O fencing fails if you use the same CP servers

When you reconfigure an application cluster that uses server-based I/O fencing (customized fencing mode), the installer does not remove the application cluster information from the CP servers before the reconfiguration. As a result, if you reconfigure the application cluster and choose to configure I/O fencing in customized mode using the same CP servers, then reconfiguration of server-based fencing for the application cluster fails. [2076240]

Workaround: Manually remove the application cluster information from the CP servers after you reconfigure but before you reconfigure server-based I/O fencing for the application cluster.

See the *Veritas Cluster Server Administrator's Guide* for instructions to remove the application cluster information from the CP servers.

CP server cannot bind to multiple IPs (2085941)

Coordination point server (CP server) binds only to a single virtual IP and listens on the same. Application clusters cannot access the CP server if it fails to establish connection to this virtual IP. Therefore, if the connection fails because of the subnet in which the virtual IP of the CP server exists, you cannot access the CP server even if there is another subnet through which the client can connect to the CP server over a different IP.

Resolution: No known resolution for this issue.

Issues related to Symantec Product Authentication Service with VCS

This section covers the known issues related to Symantec Product Authentication Service (AT) in this release.

The vcsat and cpsat commands may appear to be hung

The following commands may appear to be hung when you invoke them from the command shell:

- /opt/VRTScps/bin/cpsat
- /opt/VRTSvcs/bin/vcsat

This issue occurs when the command requires some user interaction. [1841185]

Workaround:

- To fix the issue for vcsat, run the commands as follows:

  ```
  # export EAT_HOME_DIR=/opt/VRTSvcs
  # /opt/VRTSvcs/bin/vssatvcs command_line_argument
  # unset EAT_HOME_DIR
  ```

- To fix the issue for cpsat, run the commands as follows:

  ```
  # export EAT_HOME_DIR=/opt/VRTScps
  # /opt/VRTScps/bin/vssatcps command_line_argument
  # unset EAT_HOME_DIR
  ```

Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 SP1 RP1

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 SP1 RP1 release.

RVGPrimary online script does not function correctly (1949293)

The RVGPrimary online script does not function correctly.

Issues with bunker replay

When ClusterFailoverPolicy is set to Auto and the AppGroup is configured only on some nodes of the primary cluster, global cluster immediately detects any system fault at the primary site and quickly fails over the AppGroup to the remote site. VVR might take longer to detect the fault at the primary site and to complete its configuration changes to reflect the fault.

This causes the RVGPrimary online at the failover site to fail and the following message is displayed:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname
could not be imported on bunker host hostname. Operation
failed with error 256 and message VxVM
VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote server
unreachable...

Timestamp VCS ERROR V-16-2-13066 (hostname) Agent is calling
 clean for resource(RVGPrimary) because the resource
 is not up even after online completed.
```

Resolution: To ensure that global clustering successfully initiates a bunker replay, Symantec recommends that you set the value of the OnlineRetryLimit attribute to a non-zero value for RVGPrimary resource when the primary site has a bunker configured.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

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

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

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

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

Workaround:

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

Database fails over during Flashsnap operations (1469310)

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

Workaround:

There is no workaround for this issue.

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

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

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

Workaround

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

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

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

```
# vxdg join snapshop_disk_group_name
        primary_disk_group_name
```

2 Start all the volumes in primary disk group.

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

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

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

Repeat this step for all the volumes.

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

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

Workaround:

There is no workaround for this issue.

Clone command errors in a Data Guard environment using the MEMORY_TARGET feature for Oracle 11g (1824713)

The `dbed_vmclonedb` command displays errors when attempting to take a clone on a STANDBY database in a dataguard environment when you are using the MEMORY_TARGET feature for Oracle 11g.

When you attempt to take a clone of a STANDBY database, the `dbed_vmclonedb` displays the following error messages:

```
dbed_vmclonedb started at 2009-08-26 11:32:16
Editing remote_login_passwordfile in initclone2.ora.
Altering instance_name parameter in initclone2.ora.
Altering instance_number parameter in initclone2.ora.
Altering thread parameter in initclone2.ora.
SFORA dbed_vmclonedb ERROR V-81-4918 Database clone2 has not been
correctly recovered.
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.20569/recover.log.
```

This is Oracle 11g-specific issue known regarding the MEMORY_TARGET feature, and the issue has existed since the Oracle 11gr1 release. The MEMORY_TARGET feature requires the `/dev/shm` file system to be mounted and to have at least 1,660,944,384 bytes of available space. The issue occurs if the `/dev/shm`file system is not mounted or if the file system is mounted but has available space that is less than the required minimum size.

Workaround

To avoid the issue, remount the `/dev/shm` file system with sufficient available space.

1. Shut down the database.
2. Unmount the `/dev/shm` file system:

   ```
   # umount /dev/shm
   ```

3. Mount the `/dev/shm` file system with the following options:

   ```
   # mount -t tmpfs shmfs -o size=4096m /dev/shm
   ```

4. Start the database.

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.

Veritas Storage Foundation for Oracle RAC known issues

The following are new additional Veritas Storage Foundation for Oracle RAC known issues in this 5.1 SP1 RP1 release.

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

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

Workaround:

1. Log into each node in the cluster as the root user.
2. Perform the following operations:

   - If you have not yet installed Oracle Clusterware/Grid Infrastructure, create the directory and set the correct ownership as follows before you invoke the installation program:

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

     where:
     oracle_base is the name of the Oracle base directory.
     user_name is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).
     oraInventory_group_name is the name of the oraInventory group.
     Complete the other preparatory tasks before you invoke the installation program. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

   - If you faced this issue during the installation of Oracle Clusterware/Grid Infrastructure, open another terminal session, and modify the ownership of the directory on all nodes in the cluster as follows:

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

     where:
     oracle_base is the name of the Oracle base directory.
     user_name is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).

oraInventory_group_name is the name of the oraInventory group.

Return to the former session and proceed with the installation.

Software limitations

This section covers the software limitations of this release.

Veritas Storage Foundation software limitations

There are no Veritas Storage Foundation software limitations in the 5.1 SP1 RP1 release.

Veritas Volume Manager software limitations

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

DMP cannot detect the re-enabled OS device status on SLES11 after restoring the daemon (1718573)

Because the remove_on_dev_loss parameter of the scsi_transport_fc modile is removed in SLES 11, the OS device files are removed after a device loss with the dev_loss_tmo parameter. Whenthe device comes back online, the port names may have changed, in which case DMP cannot recognize the device's status with the restored daemon.

Workaround:

Set the dev_loss_tmo parameter to 8000000. This workaround only works with QLogic drivers, as Emulex drivers are not supported on SLES 11.

DMP settings for NetApp storage attached environment

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

Table 1-9

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

The change is persistent across reboots.

To change the tunable parameters

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

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

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

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

DMP behavior on Linux SLES11 when connectivity to a path is lost (2049371)

On SLES 11, when the connectivity to a path is lost, the SLES 11 kernel removes the device path from its database. DMP reacts to the UDEV event that is raised in this process, and marks the device path as DISABLED[M]. DMP will not use the path for further I/Os. Unlike on other flavours of Linux, the path state is DISABLED[M] instead of DISABLED. Subsequently, if the path comes back online, DMP responds to the UDEV event to signal the addition of device path into SLES 11 kernel. DMP enables the path and changes its state to ENABLED.

Veritas File System software limitations

The following are software limitations in the 5.1 SP1 RP1 release of Veritas Storage Foundation.

There are no Veritas Storage Foundation software limitations in the 5.1 SP1 RP1 release.

Linux I/O Scheduler for Database Workloads (1446361)

Symantec recommends using the Linux deadline I/O scheduler for database workloads on both Red Hat and SUSE distributions.

To configure a system to use this scheduler, include the `elevator=deadline` parameter in the boot arguments of the GRUB or LILO configuration file.

The location of the appropriate configuration file depends on the system's architecture and Linux distribution:

Configuration File	Architecture and Distribution
/boot/grub/menu.lst	RHEL5 x86_64, SLES10 x86_64, and SLES11 x86_64

For the GRUB configuration files, add the `elevator=deadline` parameter to the kernel command. For example, change:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2
    initrd /boot/initrd-2.6.18-128.el5.img
```

To:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2 \
    elevator=deadline
    initrd /boot/initrd-2.6.18-128.el5.img
```

A setting for the elevator parameter is always included by SUSE in its LILO and GRUB configuration files. In this case, change the parameter from `elevator=cfq` to `elevator=deadline`.

Reboot the system once the appropriate file has been modified.

See the Linux operating system documentation for more information on I/O schedulers.

Veritas Volume Replicator software limitations

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

Replication in a shared environment

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

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.
- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.

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

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 5.1SP1 RP1 and the prior major releases of Storage Foundation (5.0 MP3 and 5.1). Replication between versions is supported for disk group versions 140, 150, and 160 only. Both the Primary and Secondary hosts must be using a supported disk group version.

Cannot relayout data volumes in an RVG from concat to striped-mirror (2162537)

This issue occurs when you try a relayout operation on a data volume which is associated to an RVG, and the target layout is a striped-mirror.

Workaround:

To relayout a data volume in an RVG from concat to striped-mirror

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

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

3. Stop the affected RVG. Enter the following:

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

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

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

5. Relayout the volumes to striped-mirror. Enter the following:

   ```
   # vxassist -g diskgroup relayout vol layout=stripe-mirror
   ```

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

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

7. Start the RVG. Enter the following:

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

8. Resume or start the applications.

Veritas Cluster Server software limitations

The following are software limitations in this release of Veritas Cluster Server.

Limitations related to installing and upgrading VCS

Limitation when you use the installer from a remote system

If you use the installer from a remote system, then the remote system must have the same operating system and architecture as that of the target systems where you want to install VCS. [589334]

Limitations related to the VCS engine

VCS deletes user-defined VCS objects that use the HostMonitor object names

If you had defined the following objects in the main.cf file using the reserved words for the HostMonitor daemon, then VCS deletes these objects when the VCS engine starts. [1293092]

- Any group that you defined as VCShmg along with all its resources.
- Any resource type that you defined as HostMonitor along with all the resources of such resource type.
- Any resource that you defined as VCShm.

Limitations related to the bundled agents

Programs using networked services may stop responding if the host is disconnected

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can stop responding if the host is disconnected from the network. If such a program is used as an agent entry point, a network disconnect can cause the entry point to stop responding and possibly time out.

For example, if the host is configured to use NIS maps as a client, basic commands such as `ps -ef` can hang if there is network disconnect.

Symantec recommends creating users locally. To reflect local users, configure:

/etc/nsswitch.conf

Limitations of the DiskGroup agent

Volumes in disk group are started automatically if the Veritas Volume Manager default value of AutoStartVolumes at system level will be set to ON irrespective of the value of the StartVolumes attribute defined inside the VCS. Set AutoStartVolumes to OFF at system level if you do not want to start the volumes as part of import disk group.

Volume agent clean may forcibly stop volume resources

When the attribute FaultOnMonitorTimeouts calls the Volume agent clean entry point after a monitor time-out, the `vxvol -f stop` command is also issued. This command forcibly stops all volumes, even if they are still mounted.

False concurrency violation when using PidFiles to monitor application resources

The PID files created by an application contain the PIDs for the processes that are monitored by Application agent. These files may continue to exist even after a node running the application crashes. On restarting the node, the operating system may assign the PIDs listed in the PID files to other processes running on the node.

Thus, if the Application agent monitors the resource using the PidFiles attribute only, the agent may discover the processes running and report a false concurrency violation. This could result in some processes being stopped that are not under VCS control.

Mount agent

The Mount agent mounts a block device at only one mount point on a system. After a block device is mounted, the agent cannot mount another device at the same mount point.

Share agent

To ensure proper monitoring by the Share agent, verify that the /var/lib/nfs/etab file is clear upon system reboot. Clients in the Share agent must be specified as fully qualified host names to ensure seamless failover.

Driver requirements for DiskReservation agent

The DiskReservation agent has a reserver module in the kernel mode that reserves disks persistently. Any driver that works correctly with the scsiutil utility shipped with the VRTSvcsdr package is supported. Refer to the manual page for scsiutil functionality.

Volumes in a disk group start automatically irrespective of the value of the StartVolumes attribute in VCS

Volumes in a disk group are started automatically if the value of the system level attribute autostartvolumes in Veritas Volume Manager is set to On, irrespective of the value of the StartVolumes attribute in VCS.

Workaround

If you do not want the volumes in a disk group to start automatically after the import of a disk group, set the autostartvolumes attribute to Off at the system level.

Limitations related to IMF

- IMF registration on Linux for “bind” file system type is not supported.
- In case of SLES11 SP1:
  - IMF should not be enabled for the resources where the BlockDevice can get mounted on multiple MountPoints.
  - If FSType attribute value is nfs, then IMF registration for “nfs” file system type is not supported.

Limitations related to the VCS database agents

DB2 RestartLimit value

When multiple DB2 resources all start at the same time with no dependencies, they tend to interfere or race with each other. This is a known DB2 issue.

The default value for the DB2 agent RestartLimit is 3. This higher value spreads out the re-start of the DB2 resources (after a resource online failure), which lowers the chances of DB2 resources all starting simultaneously. [1231311]

Limitation with intentional offline functionality of VCS agent for Oracle

The Oracle resource never faults after an intentional offline.

Intentional offline functionality of VCS agent for Oracle requires you to enable health check monitoring. The agent uses Oracle's Health Check API to find the state of the database. If the API returns a graceful shutdown for the database, then the agent marks the resource state as INTENTIONAL OFFLINE. Later if the Oracle agent's online function does not succeed, the agent does not mark the resource as FAULTED. The state remains as INTENTIONAL OFFLINE because the agent receives the database state from the API as graceful shutdown during each monitor cycle. [1805719]

Limitations related to global clusters

- Cluster address for global cluster requires resolved virtual IP.

  The virtual IP address must have a DNS entry if virtual IP is used for heartbeat agents.
- Total number of clusters in a global cluster configuration can not exceed four.
- Cluster may not be declared as faulted when Symm heartbeat agent is configured even when all hosts are down.

  The Symm agent is used to monitor the link between two Symmetrix arrays. When all the hosts are down in a cluster but the Symm agent is able to see the replication link between the local and remote storage, it would report the heartbeat as ALIVE. Due to this, DR site does not declare the primary site as faulted.

Security-Enhanced Linux is not supported on SLES distributions

VCS does not support Security-Enhanced Linux (SELinux) on SLES10 and SLES11. [1056433]

Systems in a cluster must have same system locale setting

VCS does not support clustering of systems with different system locales. All systems in a cluster must be set to the same locale.

VxVM site for the disk group remains detached after node reboot in campus clusters with fire drill

When you bring the DiskGroupSnap resource online, the DiskGroupSnap agent detaches the site from the target disk group defined. The DiskGroupSnap agent invokes VCS action entry points to run VxVM commands to detach the site. These commands must be run on the node where the disk group is imported, which is at the primary site.

If you attempt to shut down the node where the fire drill service group or the disk group is online, the node goes to a LEAVING state. The VCS engine attempts to take all the service groups offline on that node and rejects all action entry point requests. Therefore, the DiskGroupSnap agent cannot invoke the action to reattach the fire drill site to the target disk group. The agent logs a message that the node is in a leaving state and then removes the lock file. The agent's monitor function declares that the resource is offline. After the node restarts, the disk group site still remains detached. [1272012]

Workaround:

You must take the fire drill service group offline using the `hagrp -offline` command before you shut down the node or before you stop VCS locally.

If the node has restarted, you must manually reattach the fire drill site to the disk group that is imported at the primary site.

If the secondary node has crashed or restarted, you must manually reattach the fire drill site to the target disk group that is imported at the primary site using the following command: `/opt/VRTSvcs/bin/hares -action $targetres joindg -actionargs $fdsitename $is_fenced -sys $targetsys`.

Limitations with DiskGroupSnap agent

The DiskGroupSnap agent has the following limitations:

- The DiskGroupSnap agent does not support layered volumes. [1368385]
- If you use the Bronze configuration for the DiskGroupSnap resource, you could end up with inconsistent data at the secondary site in the following cases: [1391445]
  - After the fire drill service group is brought online, a disaster occurs at the primary site during the fire drill.
  - After the fire drill service group is taken offline, a disaster occurs at the primary while the disks at the secondary are resynchronizing.

  Symantec recommends that you use the Gold configuration for the DiskGroupSnap resource.

Using the KDE desktop

Some menus and dialog boxes on Cluster Manager (Java Console) may appear misaligned or incorrectly sized on a KDE desktop. To ensure the proper appearance and functionality of the console on a KDE desktop, use the Sawfish window manager. You must explicitly select the Sawfish window manager even if it is supposed to appear as the default window manager on a KDE desktop.

System reboot after panic

If the VCS kernel module issues a system panic, a system reboot is required [293447]. The supported Linux kernels do not automatically halt (CPU) processing.

Set the Linux “panic” kernel parameter to a value other than zero to forcibly reboot the system. Append the following two lines at the end of the /etc/sysctl.conf file:

```
# force a reboot after 60 seconds
  kernel.panic = 60
```

Cluster Manager (Java console) limitations

This section covers the software limitations for Cluster Manager (Java Console).

Use the VCS 5.1 Java Console to manage clusters

Cluster Manager (Java Console) from versions lower than VCS 5.1 cannot be used to manage VCS 5.1SP1 clusters. Symantec recommends using the latest version of Cluster Manager.

See the *Veritas Cluster Server Installation Guide* for instructions on upgrading Cluster Manager.

Run Java Console on a non-cluster system

Symantec recommends not running Cluster Manager (Java Console) for an extended period on a node in the cluster. The Solaris version of the Java Virtual Machine has a memory leak that can gradually consume the host system’s swap space. This leak does not occur on Windows systems.

Cluster Manager does not work if the hosts file contains IPv6 entries

VCS Cluster Manager fails to connect to the VCS engine if the /etc/hosts file contains IPv6 entries.

Workaround: Remove IPv6 entries from the /etc/hosts file.

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the UseFence attribute is set to the default, “None”.

Limitations related to I/O fencing

This section covers I/O fencing-related software limitations.

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *Veritas Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

I/O fencing uses SCSI-3 PR keys to implement data protection. Keys are placed on I/O fencing coordinator points and on data disks. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator points and data disks to prevent possible difficulties with subsequent cluster startup.

Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordinator points and data disks. Depending on the order of reboot and subsequent startup events, the cluster may warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown -r command on one node at a time and wait for each node to complete shutdown.

Veritas Storage Foundation for Databases tools software limitations

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

Oracle Data Guard in an Oracle RAC environment

Database snapshots and Database Checkpoints are not supported in a Data Guard and Oracle RAC environment.

Upgrading if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 5.1SP1 RP1: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to SP1.

List of RPMs

This section lists the RPMs for 5.1 SP1 RP1.

Note: You can also view the following list using the `installrp` command, type:

```
./installrp -listpatches
```

Table 1-10 RPMs for Red Hat Enterprise Linux 5

Name	Version	Arch	Size in bytes
VRTSamf	5.1.101.000	x86_64	766250
VRTScavf	5.1.101.000	i386	176085

Table 1-10 RPMs for Red Hat Enterprise Linux 5 *(continued)*

Name	Version	Arch	Size in bytes
VRTScps	5.1.101.000	i686	15381359
VRTSdbed	5.1.101.000	i686	6159251
VRTSfssdk	5.1.101.000	x86_64	350130
VRTSglm	5.1.101.000	x86_64	106310
VRTSodm	5.1.101.000	x86_64	236750
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcs	5.1.101.000	i686	46645997
VRTSvcsag	5.1.101.000	i686	676267
VRTSvcsea	5.1.101.000	i686	3810859
VRTSvxfen	5.1.101.000	x86_64	445171
VRTSvxfs	5.1.101.000	x86_64	9589271
VRTSvxvm	5.1.101.000	x86_64	27106642

Table 1-11 RPMs for SUSE Linux Enterprise Server 10

Name	Version	Arch	Size in bytes
VRTSamf	5.1.101.000	x86_64	4146005
VRTScavf	5.1.101.000	i386	150082
VRTScps	5.1.101.000	i686	14430684
VRTSdbed	5.1.101.000	i586	5517947
VRTSfssdk	5.1.101.000	x86_64	307170
VRTSglm	5.1.101.000	x86_64	177000
VRTSodm	5.1.101.000	x86_64	358926
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcs	5.1.101.000	i586	42164375
VRTSvcsag	5.1.101.000	i586	541250
VRTSvcsea	5.1.101.000	i586	3592924

Table 1-11 RPMs for SUSE Linux Enterprise Server 10 *(continued)*

Name	Version	Arch	Size in bytes
VRTSvxfen	5.1.101.000	x86_64	2190354
VRTSvxfs	5.1.101.000	x86_64	11565706
VRTSvxvm	5.1.101.000	x86_64	25777582

Table 1-12 RPMs for SUSE Linux Enterprise Server 11

Name	Version	Arch	Size in bytes
VRTSamf	5.1.101.000	x86_64	1148912
VRTScavf	5.1.101.000	i386	153921
VRTScps	5.1.101.000	i686	11568539
VRTSdbed	5.1.101.000	i586	5517947
VRTSfssdk	5.1.101.000	x86_64	267687
VRTSglm	5.1.101.000	x86_64	166511
VRTSodm	5.1.101.000	x86_64	275614
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcs	5.1.101.000	i686	30344584
VRTSvcsag	5.1.101.000	i686	488671
VRTSvcsea	5.1.101.000	i686	2627991
VRTSvxfen	5.1.101.000	x86_64	1394205
VRTSvxfs	5.1.101.000	x86_64	6377866
VRTSvxvm	5.1.101.000	x86_64	18203603

Downloading the 5.1 SP1 RP1 archive

The patches that are included in the 5.1 SP1 RP1 release are available for download from the Symantec website. After downloading the 5.1 SP1 RP1 rolling patch, use gunzip and tar commands to uncompress and extract it.

For the 5.1 SP1 RP1 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

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

About the installrp script

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1 provides a new upgrade script. To upgrade from Veritas Storage Foundation and High Availability Solutions version 5.1 SP1 or later, Symantec recommends that you use the new upgrade script. The `installrp` script allows you to upgrade all the patches associated with the packages installed, after which you can reboot to start all of the processes.

The installrp script options

Table 1-13 The command line options for the product upgrade script

Command Line Option	Function
`[ system1 system2... ]`	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.
`[ -precheck ]`	Use the `-precheck` option to confirm that systems meet the products' installation requirements before the installation.
`[ -logpath log_path ]`	Use the `-logpath` option to select a directory other than `/opt/VRTS/install/logs` as the location where the installrp log files, summary file, and response file are saved.
`[ -responsefile response_file ]`	Use the `-responsefile` option to perform automated installations or uninstallations using information stored in a file rather than prompting for information. *response_file* is the full path of the file that contains configuration definitions.
`[ -makeresponsefile ]`	Use the `-makeresponsefile` option to generate a response file without doing an actual installation. The text that displays install, uninstall, start, and stop actions are simulations. These actions are not performed on the system.

Table 1-13 The command line options for the product upgrade script *(continued)*

Command Line Option	Function
`[ -tmppath` *`tmp_path`* `]`	Use the `-tmppath` option to select a directory other than `/var/tmp` as the working directory for `installrp`. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
`[ -hostfile` *`hostfile_path`* `]`	Use the `-hostfile` option to specify the location of a file containing the system names for installer.
`[ -keyfile` *`ssh_key_file`* `]`	Use the `-keyfile` option to specify a key file for SSH. When you use this option the `-i` *`ssh_key_file`* is passed to every SSH invocation.
`[-kickstart` *`dir_path`*	Use to produce a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec rpms in the correct order for installing, in a format that can be used for Kickstart installations. The dir_path indicates the path to the directory in which to create the file.
`[ -patchpath` *`patch_path`* `]`	Use the `-patchpath` option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by `installrp`.

Table 1-13 The command line options for the product upgrade script *(continued)*

Command Line Option	Function
`[ -rsh | -redirect | -listpatches | -pkginfo | -serial | -upgrade_kernelpkgs | -upgrade_nonkernelpkgs | -version ]`	Use the `-rsh` option when rsh and rcp are to be forced for communication though ssh and scp is also setup between the systems. Use the `-redirect` option to display progress details without showing the progress bar. Use the `-listpatches` option to display product patches in the correct installation order. Use the `-pkginfo` option is used to display the correct installation order of packages and patches. This option is available with or without one of following options: -allpkgs, -minpkgs, and -recpkgs. Use the `-serial` option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion. Use the `-upgrade_kernelpkgs` option for the rolling upgrade's upgrade of kernel packages to the latest version Use the `-upgrade_nonkernelpkgs` option for the rolling upgrade's upgrade of non-kernel packages. In this phase, VCS packages and other agent packages are upgraded to the latest versions. Product kernel drivers are upgraded to the latest protocol version. Use the `-version` option to have the installer check and report the installed products and their versions. Identifies the installed and missing and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing and patches where applicable.

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 Veritas Storage Foundation and High Availability Solutions product for the first time on a host and install 5.1 SP1 RP1. Review the 5.1 SP1 Installation Guide and Release Notes for your product.

To install the Veritas software for the first time

1 Mount the 5.1 SP1 product disc and navigate to the folder that contains the installation program to install 5.1 SP1. Choose one of the following to start the installation:

- For Veritas Storage Foundation:

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

- For Veritas Storage Foundation High Availability:

  ```
  # ./installsfha node1 node2 ... nodeN
  ```

- For Veritas Storage Foundation Cluster File System and Veritas Storage Foundation Cluster File System High Availability:

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

- For Veritas Storage Foundation Cluster File System for Oracle RAC:

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

- For Veritas Storage Foundation for Oracle RAC:

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

- For Veritas Cluster Server:

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

- For Symantec VirtualStore:

  ```
  # ./installsvs node1 node2 ... nodeN
  ```

- For Dynamic Multi-Pathing:

  ```
  # ./installdmp node1 node2 ... nodeN
  ```

2 Review the installation prerequisites for upgrading to 5.1 SP1 RP1.

See "Prerequisites for upgrading to 5.1 SP1 RP1" on page 87.

3 Copy the patch archive downloaded from patch central to temporary location, untar the archive and browse to the directory containing the installrp script.

- If the 5.1 SP1 product is installed, run the `installrp` script to install 5.1 SP1 RP1.

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

  If you did not configure the product after the 5.1 SP1 installation, the installer prompts you to configure it. If you do not want to configure the product now, answer **n** to the prompt. To configure the product in the future, run the product installation script from the 5.1 SP1 installation media with the `-configure` option.

Installing Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 5.1

SP1 RP1 using the Web-based installer. For detailed instructions on how to install 5.1 SP1 using the Web-based installer, follow the procedures in the 5.1 SP1 Installation Guide and Release Notes for your products.

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 5.1SP1 version of the Veritas product must be installed before upgrading to 5.1 SP1 RP1.

2 On the **Select a task and product** page, select **Install SP1 RP1** from the **Task** drop-down list, and click **Next**

3 Stop all applications accessing the filesystem. Unmount all mounted file systems before installation.

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

5 After the validation completes successfully, click **Next** to install 5.1 SP1 RP1 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**.

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.

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 83.

3 Stop all applications accessing the file system. Unmount all mounted filesystems before installation.

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.

Chapter 3

Upgrading to 5.1 SP1 RP1

This chapter includes the following topics:

Prerequisites for upgrading to 5.1 SP1 RP1

The following list describes prerequisites for upgrading to the 5.1 SP1 RP1 release:

- For any product in the Veritas Storage Foundation stack, you must have the 5.1 SP1 release installed before you can upgrade that product to the 5.1 SP1 RP1 release.
- Each system must have sufficient free space to accommodate patches.
- Stop all applications accessing the vxfs filesystem. Unmount all mounted vxfs file systems before upgrading.
- The full list of prerequisites can be obtained by running `./installrp -precheck`

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 5.1 SP1 to 5.1 SP1 RP1
- 5.1 SP1 P1 to 5.1 SP1 RP1 (upgrade procedure for this path is the same as the above path)

Upgrading from 5.1 SP1 to 5.1 SP1 RP1

This section describes how to upgrade from 5.1 SP1 to 5.1 SP1 RP1 on a cluster or a standalone system.

- Performing a full upgrade to 5.1 SP1 RP1 on a cluster
  Use the procedures to perform a full upgrade to 5.1 SP1 RP1 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System (SFCFS), Veritas Storage Foundation for Oracle RAC (SFRAC), or Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC) installed and configured.
- Upgrading to 5.1 SP1 RP1 on a standalone system
  Use the procedure to upgrade to 5.1 SP1 RP1 on a system that has SF installed.
- Performing a rolling upgrade using the installer
  Use the procedure to upgrade your Veritas product with a rolling upgrade.

Performing a full upgrade to 5.1 SP1 RP1 on a cluster

Performing a full upgrade on a cluster requires stopping cluster failover functionality during the entire procedure. However, if you use Veritas Storage Foundation Cluster File System (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 Veritas Cluster Server (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 SP1 RP1:

- Performing a full upgrade to 5.1 SP1 RP1 for Veritas Cluster Server
- Performing a full upgrade to 5.1 SP1 RP1 on an SFHA cluster
- Performing a full upgrade to 5.1 SP1 RP1 on an SFCFS cluster
- Performing a full upgrade to 5.1 SP1 RP1 on an SFCFS RAC cluster
- Performing a full upgrade to 5.1 SP1 RP1 on an SF Oracle RAC cluster

Performing a full upgrade to 5.1 SP1 RP1 for Veritas Cluster Server

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

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

To upgrade VCS

1. Log in as superuser.

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

   See "System requirements" on page 15.

3. Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

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

   See "About the installrp script" on page 77.

4. Resolve any issues that the precheck finds.

5. Start the upgrade:

   ```
   # ./installrp node1node2 ... nodeN
   ```

   See "About the installrp script" on page 77.

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

Performing a full upgrade to 5.1 SP1 RP1 on an SFHA cluster

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

To perform a full upgrade to 5.1 SP1 RP1 on an SFHA cluster

1 Log in as superuser.

2 From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

```
# ./installrp node1node2
```

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

Performing a full upgrade to 5.1 SP1 RP1 on an SFCFS cluster

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

To perform a full upgrade to 5.1 SP1 RP1 on an SFCFS cluster

1 Log in as superuser.

2 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
```

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

```
# mount | grep vxfs
```

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

```
# umount /filesystem
```

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

4 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.

5 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

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

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

Verify that no volumes remain open:

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

7 If required, apply the OS kernel patches.

8 From the directory that contains the extracted and untarred 5.1 SP1 RP1 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.

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

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

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

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

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

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

13 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

14 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

Performing a full upgrade to 5.1 SP1 RP1 on an SFCFS RAC cluster

To prepare for a full upgrade to 5.1 SP1 RP1 on an SFCFS RAC cluster

1 Log in as superuser.

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

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

4 Stop Oracle Clusterware.

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

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

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

6 Unmount all file systems:

```
# umount /filesystem
```

7 Check if the VEA service is running:

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

8 If the VEA service is running, stop it:

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

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

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

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

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

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

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

13 Disable the startup scripts before upgrading the operating system.

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

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

To upgrade Storage Foundation Cluster File System for Oracle RAC

1 From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

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

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

4 Shut down and reboot the systems.

5 Upgrade Oracle RAC, if required.

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

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

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

  - Back up libodm10.so file.

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

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

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

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

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

  - Back up libodm11.so file.

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

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

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

To bring the upgraded cluster online and restore components

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

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

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

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

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

5 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

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

Performing a full upgrade to 5.1 SP1 RP1 on an SF Oracle RAC cluster

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

To upgrade to 5.1 SP1 RP1 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:

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

7 Stop VCS.

```
# hastop -all
```

8 If required, apply the OS kernel patches.

See "System requirements" on page 15.

See *Oracle's* documentation for the procedures.

9 From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. If ssh key authentication is configured then enter:

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

If ssh is not configured then enter:

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

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

10 Follow the instructions from the installer. If there is some module load/unload issue, reboot all of the nodes of the cluster.

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

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

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

```
# haconf -makerw
```

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

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

14 Make the configuration read-only:

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

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

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

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

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

17 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

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

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

19 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

Upgrading to 5.1 SP1 RP1 on a standalone system

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

To upgrade to 5.1 SP1 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.

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

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

5 Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

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

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

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

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

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

Verify that no volumes remain open:

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

9 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

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

```
# ./installrp nodename
```

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

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

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

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

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

14 Remount all VxFS file systems and Storage Checkpoints:

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

15 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

Performing a rolling upgrade using the installer

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

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

About rolling upgrades

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

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

You can perform a rolling upgrade for the following products:

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

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

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

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

Limitation: During VCS and agents upgrade, you must bring down the application High Availability (HA) for several minutes. This does not affect the application running on the cluster. You can restore the application's high availability after VCS and the agent packages are upgraded.

Performing a rolling upgrade on kernel packages: phase 1

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

To perform the rolling upgrade on kernel packages: phase 1

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

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

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

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

9 Before you proceed to phase 2, complete step 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 RPMs on all the nodes in cluster and restarts VCS cluster.

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

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

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

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

3 The installer upgrades non-kernel RPMs.

4 The installer reboots nodes that use encapsulated boot disks.

5 The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.

6 Verify the cluster's status:

```
# hastatus -sum
```

Upgrading the operating system and upgrading to 5.1 SP1 RP1

This section describes how to upgrade the operating system on a Storage Foundation node where you plan to upgrade to 5.1 SP1 RP1. This section includes the following topics:

- Upgrading RHEL 5
- Upgrading OEL 5
- Upgrading SLES 10

Upgrading RHEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 SP1 RP1.

To upgrade to a later version of RHEL 5

1. Stop Storage Foundation.
2. Upgrade to the latest update version of RHEL 5.
3. Upgrade to 5.1 SP1 RP1.
4. Start Storage Foundation.

Upgrading OEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 SP1 RP1.

To upgrade to a later version of OEL 5

1. Stop Storage Foundation.
2. Upgrade to the latest update version of OEL 5.
3. Upgrade to 5.1 SP1 RP1.
4. Start Storage Foundation.

Upgrading SLES 10

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 SP1 RP1.

To upgrade to a later version of SLES 10

1. Stop Storage Foundation.
2. Upgrade to the latest update version of SLES 10.
3. Upgrade to 5.1 SP1 RP1.
4. Start Storage Foundation.

Verifying software versions

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

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

Chapter 4

Removing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

About removing Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1

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 previous release.

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

Uninstalling Veritas Storage Foundation Cluster File System 5.1 SP1 RP1

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

Preparing to uninstall Veritas Storage Foundation Cluster File System

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

To prepare to uninstall Veritas Storage Foundation Cluster File System

1. Log in as the root user on any node in the cluster.
2. Verify that the following directories are set in your `PATH` environment variable:

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

3. Back up the following configuration files:

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

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

     ```
     # df -v /
     ```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7. Unmount all Storage Checkpoints and file systems:

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

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

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

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

To verify that no volumes are open:

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

9 Stop VCS:

```
# hastop -all
```

Uninstalling Veritas Storage Foundation Cluster File System

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

To uninstall Veritas Storage Foundation Cluster File System

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

2 Navigate to the directory that contains the uninstallation program:

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

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

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

4 Confirm the uninstallation:

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

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

5 Reboot the nodes:

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

After uninstalling the Veritas Storage Foundation Cluster File System, refer to the *Veritas Storage Foundation Cluster File System 5.1 SP1 Installation Guide* to reinstall the 5.1 SP1 software.

Uninstalling Veritas Storage Foundation for Oracle RAC

The following procedure uninstalls Veritas Storage Foundation for Oracle RAC (SFRAC).

Note: This procedure will remove the complete SFRAC stack from all nodes.

To uninstall 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 SFRAC, refer to the *Veritas Storage Foundation for Oracle RAC 5.1 SP1 Installation and Configuration Guide* document to reinstall the SFRAC 5.1 SP1 software.

Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP1

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

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

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

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

3 Back up the following configuration files:

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

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

For example, to stop Oracle Clusterware:

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

5 Stop VCS:

```
# hastop -all
```

6 Verify that port h is not open:

```
# gabconfig -a
```

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

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

8 Unmount all file systems:

```
# umount /filesystem
```

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

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

To verify that no volumes are open:

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

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

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

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

2 Navigate to the directory that contains the uninstallation program:

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

3 Start the uninstallation program:

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

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

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

The installer checks the RPMs installed on the system.

5 Confirm the uninstallation at the following prompt:

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

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

6 Reboot the nodes:

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

After uninstalling the Veritas Storage Foundation Cluster File System for Oracle RAC, refer to the *Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 Installation and Configuration Guide* to reinstall the 5.1 SP1 software.

Rolling back from 5.1 SP1 RP1 to 5.1 SP1

Rollback is not supported on Linux platform.