Veritas Storage Foundation™ and High Availability Solutions Release Notes

Solaris

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

Product version: 5.1 SP1 RP1

Document version: 5.1SP1RP1.2

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

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/334829

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

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

Before installing or upgrading 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.

About rolling back the RP installation

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

See "About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1" on page 109.

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

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

New attribute for Zone agent

When VCS takes a zone resource offline the `.vcspwd` file is removed. This file is used by VCS agents to manage application services inside the local zone.

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

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

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 Solaris operating systems

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

- SPARC
  - Solaris 10 Update 6, 7, 8, 9
  - Solaris 9 Update 7, 8 and 9 + latest recommend cluster patch
    Solaris 9 Update 9 is based on Update 8 but added hardware support for the V445, V245 and Ultra 45 SPARC platforms
- x64
  - Solaris 10 Update 6, 7, 8, 9

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

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

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

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

Required Solaris patches

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

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

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

Table 1-1 Solaris SPARC patches

Operating system	Oracle patch number
Solaris 9	114477-04 122300-29 (required for Live Upgrade) Patches required for FS, SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC: ■ For Solaris 9 Update 7 (to bring OS kernel to FS patch 122300 level): 117171-17 113073-14 ■ For Solaris 9 Update 7, Update 8 and Update 9: 112908-33 (required for 122300-10) 118558-39 (required for 122300-10) 122300-10
Solaris 10	119254-06 119042-02 125731-02 128306-05 127111-01

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

Table 1-2 Solaris x64 patches

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

Supported VCS agents

Table 1-3 lists the agents for enterprise applications and the software that the agents support.

Table 1-3 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	SPARC: Solaris 9, 10; x64: Solaris 10
Oracle	Oracle	11gR1, 10gR2, 11gR2	SPARC: Solaris 9, 10; x64: Solaris 10
Sybase	Sybase Adaptive Server Enterprise	12.5.x, 15.x	SPARC: Solaris 9, 10; x64: Solaris 10

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.

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

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

Recommended memory and swap space

Symantec recommends the following memory and swap space sizes:

- On the system where you run the installation, 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 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 1-4 Veritas Volume Manager 5.1 SP1 RP1 fixed issues

Fixed issues	Description
2160199	Master takeover fails as the upcoming Master could not import shared DG. SUN Bug ID is 6988420.
2148682	while shipping a command node hangs in master selection on slave nodes and master update on master node
2181631	Striped-mirror volume cannot be grown across sites with -oallowspansites with DRL
2133503	Renaming enclosure results in dmpevents.log reporting Mode for Enclosure has changed from Private to Private
2200670	vxattachd does not recover disks if disk group is not imported
2191693	'vxdmpadm native list' command is not displaying any output nor error
2080730	vxvm/vxdmp exclude file contents after updation should be consistent via vxdiskadm and vxdmpadm
2129989	EVA ASL should report an error message if pref_bit is not set for a LUN
2129477	vxdisk reclaim command fails after resize operation.
2194492	VxVM-ASM co-existence enablement
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core
2166682	checks needed to make sure that a plex is active before reading from it during fsvm mirror read interface
2088007	Possibility of reviving only secondary paths in DMP

Table 1-4 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2188590	An ilock acquired by a SLAVE node for a read on a DCL object can lead to IO hang when the node becomes MASTER before reading is done
2015467	Performance improvement work for NetBackup 6.5.5 on SF 5.1 VxVM mapping provider
1829285	vxconfigd coredumps while assigning unique native name to a disk
2226813	VVR: rlinks remain disconnected with UDP protocol if data ports are specified
2227923	renaming of enclosure name is not persistent
2172488	FMR: with dco version 0 restore operation doesn't sync the existing snapshot mirrors
2183984	System panics due to race condition while updating DMP I/O statistics
2038928	creation of pre 5.1SP1 (older) version diskgroup fails
1970560	When vxconfigd is idle (which is not shipping the command) slave dies and command shipping is in progress, vxconfigd core dumped on Master
2082450	In case of failure, vxdisk resize should display more meaningful error message
1869002	Introduction of Circular buffer at vold level for master-slave communication.
1426480	VOLCVM_CLEAR_PR ioctl does not propogate the error returned by DMP to the caller
2105547	tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault
1959513	Propagate -o noreonline option of disk group import to slave nodes
1940052	vxconfigd hung on Master after removing the hba alias from zone and node leave followed by join
2199496	Data Corruption seen with "site mirror" Campus Cluster feature
2234844	asm2vxfs conversion fails
2215216	vxkprint does not report TP related values

Table 1-4 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

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

Veritas File System: Issues fixed in 5.1 SP1 RP1

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

Table 1-5 Veritas File System fixed issues

Incident	Description
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.
2030119	fsppadm core dumps when analysing a badly formatted XML file, is resolved
2162822	During online migration from ufs to vxfs, df command returns a non-zero return value.
2169273	During online migration, nfs export of the migrating file system leads to system panic
2177253	A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for vxupgrade purposes

Table 1-5 Veritas File System fixed issues *(continued)*

Incident	Description
2178147	Linking a IFSOC file now properly calls `vx_dotdot_op`(), which fixes the cause of a corrupted inode.
2184528	`fsck` no longer fails to repair corrupt directory blocks that have duplicate directory entries.
2194618	Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system
2221623	Fixed a performance loss due to a delxwri_ilist spin lock with the default values for `vx_idelxwri_timelag`.
2255786	Failure to full fsck cleanly due to incorrect string comparison of file name.
2228311	Support for mostly online migration from ZFS to VxFS

Veritas Storage Foundation Cluster File System: 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 1-6 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

Table 1-6 Veritas Storage Foundation Cluster File System fixed issues *(continued)*

Incident	Description
2184114	In a large filesystem, stat operation on cluster mount leaves the file system frozen for too long leading to CVMVoldg and CFSMount agent timeouts.
2203917	ODM I/O performance bottleneck due to threads waiting in odm_rwsleep_lock() to get lock on odm_iop_table is resolved
2232554	System panic in vx_iupdat_clustblks() due to an unprotected inode getting corrupted.
2241123	glmdump failure with the error "/opt/VRTSglm/sbin/glmdump: syntax error at line 340: `count=$' unexpected" is resolved.

Veritas 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 1-7 Veritas Cluster Server 5.1 SP1 RP1 fixed issues

Fixed Issues	Description
1949294	`fdsetup` can now correctly parse disk names containing characters such as "-".
1949303	`fdsetup` no longer allows volume that are not part of the RVG, which fixes a possible cause of the RVGSnapshot agent failing.
2011536	Added IMF support for the db2udb agent.
2159991	Fixed an issue with messages in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2172181	Fixed an issue with AMF-related messages for the CAVF agent in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2175599	Fixed an issue with fencing configuration on a 64-node cluster.
2179652	The monitor script of the db2udb agent can now handle empty attribute values.

Table 1-7 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Fixed Issues	Description
2184205	Fixed an issue with HAD in which the parent service group did not fail over if the parent service group had an online local firm dependency with a child service group.
2194473	HAD no longer dumps core while overriding the static attribute to the resource level.
2203430	Fixed an issue with haping.
2205556	Fixed an issue with the offline EP of the DNS agent, which did not remove all A/AAAA records if OffDelRR=1 for multi-home records.
2205563	A clean EP now properly removes resource records when OffDelRR=1.
2205567	Fixed an issue in which having an attribute set to `master.vfd` caused the DNS agent to fail to query the DNS server.
2208416	The zone agent now properly starts the network services when the `BootState` attribute is set to multi-user-server.
2208901	Fixed an issue with the RVGSnapshot agent.
2209337	Fixed an issue with VCSAPI where the RemoteGroup agent crashed if the VCSAPI log level was set to a non-zero value.
2210281	Fixed an issue with the agent framework in which MonitorTimeStats incorrectly showed 303 seconds intermittently.
2213725	Fixed an issue with the zone agent, in which some of the system services did not start if BootState was configured.
2214539	Fixed an issue in which rebooting a node sometimes set the intentonline of a group to 2, even if the group was online somewhere else. This caused the group to use the autostartlist and not perform a failover.
2218556	Fixed an issue in the `cpsadm` command in which it sometimes failed if LLT was not installed or configured on a single node cluster.
2230869	Fixed an issue with the `IPMultiNICB.xml` file in which the VCS GUI showed the `UseMpathd` attribute for IPMutliNICB resources.
2232633	Fixed an issue with the zone agent in which offline zone resource removed all VCS zone users.
2241397	Fixed an issue in which halogin did not work in a secure environment where the root broker was not a VCS node.

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.

PATCH_REQUIRES filed is not set for VCS patches in patchinfo (2249560)

The `PATCH_REQUIRES` field is not set for VCS patches. Due to this 5.1 SP1 RP1 patches may get install on any 5.1 release which is not correct. This problem occurs only when you install 5.1 SP1 RP1 patches individually with `patchadd` command.

Workaround: Use product installer `installrp` script to install the 5.1 SP1 RP1 patches which installs 5.1 SP1 RP1 patches only if 5.1 SP1 is already installed.

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*

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

When you upgrade or uninstall Veritas High Availability product, some modules may fail to unload with error messages similar to the following messages:

```
fdd failed to stop on node_name
vxfs failed to stop on node_name
```

The issue may be observed on any one or all the nodes in the sub-cluster.

Workaround: After the upgrade or uninstallation completes, follow the instructions provided by the installer to resolve the issue.

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

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

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

The VRTSacclib package is deprecated package (2032052)

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

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

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

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

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

http://wikis.sun.com

Veritas Storage Foundation known issues

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

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

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

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

db2exp may frequently dump core (1854459)

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

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

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

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

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

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

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

DBI1070I  Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

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

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

To communicate in a dual-stack environment

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

  127.0.0.1 swlx20-v6

  Or

  127.0.0.1 swlx20-v6.punipv6.com

  where *127.0.0.1* is the IPv4 loopback address.

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

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

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

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

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

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

For example:

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

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

```
    changed /kernel/drv/vxspec.conf

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

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

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

Tools like `dbca` may hang during database creation.

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

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

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

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

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

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

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

After upgrading SF stack from 5.0MP3SP1 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.

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

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

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

Work-around:

To resolve the issue, perform the following steps.

To recover from the error

1 Stop the I/O.

2 Reboot the system with the following command:

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

3 Retry the operation.

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

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

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

Expanding a LUN to a size greater than 1 TB fails to show correct expanded size (2123677)

This issue occurs when you perform a Dynamic LUN Expansion for a LUN that is smaller than 1 TB and increase the size to greater than 1 Tb. After the expansion, Veritas Volume Manager (VxVM) fails ongoing I/O, and the public region size is reset to original size. After you run the `vxdisk scandisks` command, VxVM does not show the correct expanded size of the LUN. The issue is due to underlying Solaris issues. Refer to Sun Bug Id 6929449 and Sun Bug Id 6912703.

Workaround: There is no workaround for this issue.

Co-existence check might fail for CDS disks

In Veritas Volume Manager (VxVM) 5.1 SP1 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.

Removing a volume from a thin LUN in an alternate boot disk group triggers disk reclamation (2080609)

If you remove a volume from an alternate boot disk group on a thin LUN, this operation triggers thin reclamation, which may remove information required for the disk to be bootable. This issue does not affect the current boot disk, since VxVM avoids performing a reclaim on disks under the bootdg.

Workaround: If you remove a volume or plex from an alternate boot disk group with the `vxedit` command, specify the `-n` option to avoid triggering thin reclamation. For example:

```
# vxedit -g diskgroup -rfn rm volumename
```

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

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

Suppressing the primary path of an encapsulated SAN boot disk from Veritas Volume Manager causes the system reboot to fail (1933631)

If you suppress the primary path of an array from VxVM control and then reboot the system, the system boot fails.

If you have an encapsulated SAN boot device with multiple primary paths, the issue occurs when you suppress the first primary path. When you configure a SAN boot device, the primary path is set as a boot device. In general, the first path of the SAN boot device corresponds to the first configured path during SAN boot. Even if another primary path is configured as a boot device, suppressing the first device from VxVM causes the boot to fail.

Workaround:

When the boot device is suppressed from VxVM, change the OS boot device sequencing accordingly.

For Solaris SPARC system, use the `eeprom boot-device` command to set the boot device sequencing.

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

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 1-8 shows the Hitachi arrays that have new array names.

Table 1-8 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

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 -F vxfs /dev/vx/dsk/dg/snapshot_volume
```

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

For example:

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

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

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

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

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

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

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

When VVR is used for replicating shared disk groups in an SFCFS or SFRAC environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster

nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg`commands:

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

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

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

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

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

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

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

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

Workaround:

To resolve this issue

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

   ```
   export IPM_HEARTBEAT_TIMEOUT
   IPM_HEARTBEAT_TIMEOUT=120
   ```

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

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

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

Workaround: There is no workaround for this issue.

vxassist relayout removes the DCM (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

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

3. Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

Veritas File System known issues

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

Storage Checkpoints can exceed the quota limit (2102201)

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

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

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

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

Workaround: There is no workaround for this issue.

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.

Installing the VRTSvxfs 5.1 SP1 RP1 package on non-global zones can fail (2086894)

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

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

Workaround: The following procedure installs the VRTSvxfs 5.1 SP1 RP1 package on a non-global zone.

To install VRTSvxfs 5.1 SP1 RP1 on a non-global zone

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

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

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

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

Workaround: You must reboot after the upgrade completes.

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).

imf_register reports error while doing registeration (2011536)

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

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.

Manual installation or uninstall of kernel component packages and patches require boot archive updates

After manually installing and uninstalling the packages or the patches corresponding to kernel components such as LLT, GAB, and VXFEN, make sure to update the boot archive. [2159242]

- In case of primary boot environment, run the following command:

```
# bootadm update-archive
```

- In case of an alternate root environment, run the following command:

```
# bootadm update-archive -R [altrootpath]
```

See the bootadm manual page for Solaris.

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.

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]
```

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

When you upgrade or uninstall Veritas High Availability product, some modules may fail to unload with error messages similar to the following messages:

```
fdd failed to stop on node_name
vxfs failed to stop on node_name
```

The issue may be observed on any one or all the nodes in the sub-cluster.

Workaround: After the upgrade or uninstallation completes, follow the instructions provided by the installer to resolve the issue.

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 Veritas High Availability product 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 Veritas High Availability product, 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.

Installed 5.0 MP3 without configuration, then upgrade to 5.1 SP1, installer can not continue (2016346)

If you install 5.0MP3 without configuration, you cannot upgrade to 5.1SP1. This upgrade path is not supported.

Workaround: Uninstall 5.0 MP3, and then install 5.1 SP1.

Issues related to the VCS engine

Missing host names in engine_A.log file

The GUI does not read the engine_A.log file. It reads the engine_A.ldf file, gets the message id from it, and then queries for the message from the bmc file of the

appropriate locale (Japanese or English). The bmc file does not have system names present and so they are read as missing. [1736295]

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.

Parent group faulting in zone 1 results in the child group being automatically failed over to zone 2

Parent group faulting in zone 1 results in the child group being automatically failed over to zone 2. [1859387]

Error messages displayed while uninstalling VCS 5.1SP1 patch (2128536)

Description: The following messages may be displayed while uninstalling the VCS 5.1SP1 patch.

```
Pkgadd failed. See /var/tmp/<*.log> for details
Patchrm is terminating.
WARNING: patchrm returned <7> for global zone
Done!
```

Resolution: This is an ignorable error/warning message and the patch is uninstalled successfully.

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

English text while using 'hares -action' command (1786742)

Description: The output of hares –action is displayed in English text and not in your configured locale.

Resolution: No resolution.

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 Live Upgrade

The VRTSvcsea package removal fails while removing it from the alternate disk during live upgrade (2096925)

Description: This is because the pre-remove script of VRTSvcsea checks if any of the following agent is running:

- ASMDG
- ASMInst
- Db2udb
- Netlsnr
- Oracle
- Sybase
- SybaseBk

Therefore, even if agents are running from the first disk, we cannot remove the package from second disk.

Resolution: Stop any of the running agents before removing the package.

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

Live Upgrade may fail on a Solaris 9 host if a VxFS file system is in `/etc/vfstab`.

Workaround: On the Solaris 9 host, install the Live Upgrade packages `SUNWlucfg`, `SUNWluu`, and `SUNWlur` from a Solaris 10 image. After you install the packages, install the latest Live Upgrade patch.

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

http://wikis.sun.com

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 Veritas High Availability product 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 Veritas High Availability product 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.

Delay in rebooting Solaris 10 nodes due to vxfen service timeout issues

When you reboot the nodes using the `shutdown -i6 -g0 -y` command, the following error messages may appear:

```
svc:/system/vxfen:default:Method or service exit
timed out. Killing contract 142
svc:/system/vxfen:default:Method "/lib/svc/method/vxfen stop"
failed due to signal Kill.
```

This error occurs because the vxfen client is still active when VCS attempts to stop I/O fencing. As a result, the vxfen stop service times out and delays the system reboot. [1897449]

Workaround: Perform the following steps to avoid this vxfen stop service timeout error.

To avoid the vxfen stop service timeout error

1 Stop VCS. On any node in the cluster, run the following command:

```
# hastop -all
```

2 Reboot the systems:

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

VXFEN service goes to maintenance mode when restarted if the VCS engine is running

On Solaris 10, if you restart the vxfen service when the VCS engine is running, then the vxfen service goes into maintenance mode. You must stop the VCS engine before you restart or disable the vxfen service, and then enable the vxfen service. [2116219]

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 High Availability product may report AT error during system reboot

When you reboot the systems using the `shutdown -i6 -g0 -y` command, Veritas High Availability product reports the following AT error in the /var/adm/messages file and the /var/VRTSvcs/log/engine_A.log file, and VCS faults the VxSS service group.[1765594]

```
VCS ERROR V-16-1-13067 (host_name) Agent is calling clean
for resource (vxatd) because the resource became OFFLINE
unexpectedly, on its own.
VxSS State s245sf2 |OFFLINE|FAULTED|
VxSS State s245sf3 |ONLINE|
```

Workaround: This error occurs due to a timing issue and this message may be safely ignored.

Veritas Cluster Server agents for Veritas Volume Replicator known issues in

No known issues exist for Veritas Storage Foundation Cluster File System in the 5.1SP1 release.

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.

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.

Converting a multi-pathed disk (2695660)

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

DMP settings for NetApp storage attached environment

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

Table 1-9

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
```

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

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

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

Work-around:

To resolve the issue, perform the following steps.

To recover from the error

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

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

3. Retry the operation.

Veritas File System software limitations

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

Veritas Volume Replicator software limitations

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

Replication in a shared environment

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

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.
- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.
- A replication configuration in which an IPv6 address is specified for the `local_host` attribute of a primary RLINK and an IPv4 address is specified for the `remote_host` attribute of the same RLINK.
- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.
- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 5.1SP1 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 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.

Mount resources can cause core dumps

Due to a known Solaris issue, certain system calls create memory leaks that can lead to a core dump. This happens in situations where the Mount resource's FSType

attribute has a value of nfs, and is exacerbated when the resource is for a non-global zone and the value of the SecondLevelMonitor attribute is 1. [1827036]

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.

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.

Engine hangs when you perform a global cluster upgrade from 5.0MP3 in mixed-stack environments

If you try to upgrade a mixed stack VCS environment (where IPv4 and IPv6 are in use), from 5.0MP3 to 5.1SP1, HAD may hang.

Workaround: When you perform an upgrade from 5.0MP3, make sure no IPv6 addresses are plumbed on the system.[1820327].

Use VCS installer to install or upgrade VCS when the zone root is on VxFS shared storage

You must use the VCS installer program to install or upgrade VCS when the zone root is on Veritas File System (VxFS). [1215671]

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.

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".

The operating system does not distinguish between IPv4 and IPv6 packet counts

In a dual-stack configuration, when you use packet counts and the IPv6 network is disabled, the NIC agent might not detect a faulted NIC. It might not detect a fault because while the IPv6 network is down its packet count still increases. The packet count increases because the operating system does not distinguish between the packet counts for IPv4 and IPv6 networks. The agent then concludes that the NIC is up. If you are using the same NIC device for IPv4 as well as IPv6 resources, set PingOptimize to 0 and specify a value for the NetworkHosts attribute for either the IPv6 or the IPv4 NIC resource. [1061253]

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 patches

This section lists the patches and packages for 5.1 SP1 RP1.

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

Table 1-10 Patches and packages for Solaris SPARC

Patch ID	5.1 Package Name	Products Affected	Patch Size
142633-07	VRTSvxfs(5.9)	FS, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	39M
142634-07	VRTSvxfs(5.10)	FS, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	50M
142629-09	VRTSvxvm	DMP, VM, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	199M
144159-01	VRTSsfmh	FS, DMP, VM, SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	32M
143287-03	VRTSvcs	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	68M

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

Patch ID	5.1 Package Name	Products Affected	Patch Size
143288-04	VRTSvcsag	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	409K
145454-02	VRTSvxfen (5.9)	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	331K
145455-02	VRTSvxfen (5.10)	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	335K
145471-01	VRTSamf (5.9)	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	51K
145473-01	VRTSamf (5.10)	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	51K
143289-02	VRTScps	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	2.3M
143290-03	VRTSvcsea	VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	149K
143270-06	VRTSodm(5.9)	SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.3M

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

Patch ID	5.1 Package Name	Products Affected	Patch Size
143271-06	VRTSodm(5.10)	SF, SFHA, SFCFS, SFCFSHA, SFRAC, SVS	1.2M
143680-02	VRTSglm(5.9)	SFCFS, SFCFSHA, SFRAC, SVS	906K
143681-02	VRTSglm(5.10)	SFCFS, SFCFSHA, SFRAC, SVS	576K
143273-06	VRTScavf(5.9)	SFCFS, SFCFSHA, SFRAC, SVS	816K
143274-06	VRTScavf(5.10)	SFCFS, SFCFSHA, SFRAC, SVS	818K

Table 1-11 Patches and packages for Solaris x64

Patch ID	5.1 Package Name	Products Affected	Patch Size
143294-02	VRTSvcs	VCS, SFHA, SFCFSHA, SFRAC	71M
143295-04	VRTSvcsag	VCS, SFHA, SFCFSHA, SFRAC	418K
145456-02	VRTSvxfen	VCS, SFHA, SFCFSHA, SFRAC	356K
145472-01	VRTSamf	VCS, SFHA, SFCFSHA, SFRAC	52K
145458-01	VRTSsfmh	VCS, SFHA, SFCFSHA, SFRAC	37M
143296-02	VRTScps	VCS, SFHA, SFCFSHA, SFRAC	2.4M
143297-03	VRTSvcsea	VCS, SFHA, SFCFSHA, SFRAC	145K
142630-09	VRTSvxvm	DMP,VM,SF,SFHA,SFCFS,SFCFSHA,SFRAC,SVS	181M
142635-07	VRTSvxfs	FS,SF,SFHA,SFCFS,SFCFSHA,SFRAC,SVS	32M

Table 1-11 Patches and packages for Solaris x64 *(continued)*

Patch ID	5.1 Package Name	Products Affected	Patch Size
143682-02	VRTSglm	SFCFS,SFCFSHA,SFRAC,SVS	611K
143275-06	VRTScavf	SFCFS,SFCFSHA,SFRAC,SVS	875K
143272-06	VRTSodm	SF,SFHA,SFCFS,SFCFSHA,SFRAC,SVS	1.1M

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/334829

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 1-12 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.

Table 1-12 The command line options for the product upgrade script *(continued)*

Command Line Option	Function
[-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.
[-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.
[-jumpstart *jumpstart_path*]	Use the -jumpstart option to generate finish scripts. You can use the finish scripts with the Solaris JumpStart Server to automate installation of all packages and patches for every product. You need to specify an available location to store the finish scripts as a complete path. The -jumpstart option is supported on Solaris only.

Table 1-12 The command line options for the product upgrade script *(continued)*

Command Line Option	Function
`[ -keyfile ssh_key_file ]`	Use the `-keyfile` option to specify a key file for SSH. When you use this option the `-i ssh_key_file` is passed to every SSH invocation.
`[ -patchpath patch_path ]`	Use the `-patchpath` option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by `installrp`.
`[ -rootpath root_path ]`	Use the `-rootpath` option to re-root the installation of all packages to the given path. On Solaris, -rootpath passes -R <root_path> to pkgadd.

Table 1-12 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 packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable.

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 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 91.

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**.

Chapter 3

Installing the products using JumpStart

This chapter includes the following topics:

Installing the products using JumpStart

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

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

Generating the finish scripts

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

To generate the finish script

1. Mount the 5.1SP1 RP1 media and run the `installrp` program to generate the scripts.

   ```
   # installrp -jumpstart directory_to_generate_scripts
   ```

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

   For example:

   ```
   # ./installrp -jumpstart /js_scripts
   ```

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

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

4. Specify private region length.

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

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

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

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

```
The finish scripts for AT is generated at
/js_scripts/jumpstart_at.fin
The finish scripts for DMP is generated at
/js_scripts/jumpstart_dmp.fin
The finish scripts for FS is generated at
/js_scripts/jumpstart_fs.fin
The finish scripts for SF is generated at
/js_scripts/jumpstart_sf.fin
The finish scripts for SFCFS is generated at
/js_scripts/jumpstart_sfcfs.fin
The finish scripts for SFCFSHA is generated at
/js_scripts/jumpstart_sfcfsha.fin
The finish scripts for SFHA is generated at
/js_scripts/jumpstart_sfha.fin
The finish scripts for SFRAC is generated at
/js_scripts/jumpstart_sfrac.fin
The finish scripts for SVS is generated at
/js_scripts/jumpstart_svs.fin
The finish scripts for VCS is generated at
/js_scripts/jumpstart_vcs.fin
The finish scripts for VM is generated at
/js_scripts/jumpstart_vm.fin
The encapsulation boot disk script for VM is generated at
/js_scripts/encap_bootdisk_vm51101000.fin
```

List the js_scripts directory.

```
# ls /js_scripts
```

You could select scripts according to the products you want to install and copy them to the `BUILDSRC NFS` shared location. For example, `/export/config` where you mounted the `BUILDSRC`.

For SF:

```
encap_bootdisk_vm51101000.fin jumpstart_sf.fin
```

For SFHA:

```
encap_bootdisk_vm51101000.fin jumpstart_sfha.fin
```

For SFCFS:

```
encap_bootdisk_vm51101000.fin jumpstart_sfcfs.fin
```

For SF Oracle RAC:

```
encap_bootdisk_vm51101000.fin jumpstart_sfrac.fin
```

For VCS:

```
encap_bootdisk_vm51101000.fin jumpstart_vcs51.fin
```

7 Copy the install and uninstall scripts of Veritas product that you are installing on to `BUILDSRC` shared location `/export/config` from 5.1 SP1 media.

For example:

```
# cp -p /dvd_mount/product_name/installprod  /export/config
# cp -p /dvd_mount/product_name/uninstallprod  /export/config
```

Here *prod* is one of at/dmp/fs/sf/sfcfs/sfha/sfrac/svs/vcs/vm.

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

```
BUILDSRC="hostname_or_ip:/path_to_pkgs_patches_scripts"
```

Example: **BUILDSRC=10.209.100.100:/export/config**

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

9
- If you want to install other products packages in the Storage Foundation and High Availability suite, then use the product-specific install command with one of the following options to get a list of packages in the order to be installed:
  - `minpkgs`
  - `recpkgs`
  - `allpkgs`

    Use the list of packages that is generated to replace the package list in the finish scripts.
- If you want to install other products patches in the Storage Foundation and High Availability suite, then use the product-specific install command with one of the following options to get a list of patches in the order to be installed:

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

See “About the installrp script” on page 75.
See “The installrp script options” on page 75.

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

Overview of JumpStart installation tasks

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

Summary of tasks

1 Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.

2 Read the JumpStart installation instructions.

3 Generate the finish scripts.

See “Generating the finish scripts” on page 83.

4 Modify the rules file for JumpStart.

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

5 Prepare installation resources.

See “Preparing installation resources” on page 88.

6 On each client node, run the following command to install the Veritas product packages and patches:

For Solaris SPARC:

```
Ok> boot net -  install
```

For Solaris x64:

Press **F12** and select the network boot mode.

Note: The system is restarted after the packages are installed. If you choose to encapsulate the root disk on your systems, the systems start with an encapsulated root disk.

---

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

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

where *`installprod`* is the product's installation command.

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1 Copy the contents of 5.1 SP1 disk and 5.1 SP1 RP1 disk both to buildsrc.

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

Note: After you copied the patches, you must uncompress them using the gunzip and tar commands.

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

See "Generating the finish scripts" on page 83.

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

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

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

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

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

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

See "Generating the finish scripts" on page 83.

5 Modify the rules file as required.

For example:

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

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

Adding language pack information to the finish file

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

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

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

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

do
...
done
```

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

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

do
.
.
.
done
```

Chapter 4

Upgrading to 5.1 SP1 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)

- 5.1 SP1 P2 to 5.1 SP1 RP1

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), or Veritas Storage Foundation for Oracle RAC (SFRAC) 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.
- Upgrading Veritas products using Live Upgrade
  Use the procedure to upgrade your Veritas product with a Live Upgrade.
- Performing a rolling upgrade using the installer
  Use the procedure to upgrade your Veritas product with a rolling upgrade.

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 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 16.
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 75.
4. Resolve any issues that the precheck finds.
5. Start the upgrade:

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

   See "About the installrp script" on page 75.
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. Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and uncompressed 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
   ```

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

   See "About the installrp script" on page 75.
3. Resolve any issue that the precheck finds.
4. Start the upgrade:

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

   See "About the installrp script" on page 75.

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. Verify that `/opt/VRTS/bin` and `/opt/VRTSvcs/bin` is in your `PATH` so that you can execute all product commands.
3. From any node in the cluster, make the VCS configuration writable:

   ```
   # haconf -makerw
   ```
4. Enter the following command to freeze HA service group operations on any node:

   ```
   # hagrp -freeze groupname -persistent
   ```
5. Make the configuration read-only:

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

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

```
# df -F vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# df -F vxfs
```

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

```
# umount /filesystem
```

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

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

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

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

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

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

9 Stop activity to all VxVM volumes.

For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

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

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

Verify that no volumes remain open:

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

11 Stop VCS:

```
# hastop -all
```

12 On each node, stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

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

- Solaris 9:

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

- Solaris 10:

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

14 If required, apply the OS kernel patches.

See "System requirements" on page 16.

See Oracle's documentation for the procedures.

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

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

If the VEA service is running, stop it:

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

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

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

Review the output.

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

18 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

19 Enter the following command on any node to unfreeze HA service group operations:

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

20 Make the configuration read-only:

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

21 Bring the CVM service group online on each node:

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

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

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

23 If you stopped any RVGs in step 10, restart each RVG:

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

24 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

25 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

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 16.

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

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

See Oracle's documentation for the procedures.

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

```
# df -F vxfs
```

5 Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

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

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

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

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

   Verify that no volumes remain open:

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

9. Check if the VEA service is running:

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

   If the VEA service is running, stop it:

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

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

    ```
    # ./installrp nodename
    ```

11. If necessary, reinstate any missing mount points in the `/etc/vfstab` file.
12. Restart all the volumes by entering the following command for each disk group:

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

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

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

14 Remount all VxFS file systems and Storage Checkpoints:

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

15 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

Upgrading Veritas products using Live Upgrade

This section describes how to upgrade 5.1 SP1 to 5.1 SP1 RP1 using Live Upgrade.

Supported live upgrade paths:

- Upgrading Veritas Products without Solaris OS upgrade:
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 9 Update x 5.1 SP1 RP1
  - Upgrading Solaris 10 Update x 5.1 SP1 to Solaris 10 Update x 5.1 SP1 RP1
- Upgrading Veritas Products with Solaris OS upgrade
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 9 Update y 5.1 SP1 RP1
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 10 Update y 5.1 SP1 RP1
  - Upgrading Solaris 10 Update x 5.1 SP1 to Solaris 10 Update y 5.1 SP1 RP1

Prerequisites to upgrade to 5.1 SP1 RP1 using Live Upgrade:

- The node should have an alternate boot disk that is identical to the primary boot disk.
- Installation disc for 5.1 SP1 and 5.1 SP1 RP1 to be installed on the ABE.
- Installation disc for target OS to be installed on ABE.
- The latest list of required patches is available in the Oracle Solaris Live Upgrade Software:
  Patch Requirements (Doc ID 1004881.1) document in My Oracle Support (https://support.oracle.com/).
- If OS upgrade is involved, then remove the currently installed SUNWluu, SUNWlur and SUNWlucfg packages and install SUNWluu, SUNWlur, SUNWlucfg packages from target OS. Also replace SUNWluzone if zones are involved.

- The `vxlustart` script takes around 2-3 hours to complete uninterrupted. Symantec recommends to have a network connection that does not time out in the interim.

Upgrading Veritas products using Live Upgrade from 5.1 SP1 to 5.1 SP1 RP1 without OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, or SF for Oracle RAC from 5.1 SP1 to 5.1 SP1 RP1 using Live Upgrade where OS upgrade is not involved..

To upgrade your Veritas product using Live Upgrade

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

   See your Veritas product 5.1 SP1 Installation Guide for more information.

2. Run the `vxlustart -V` command to ensure there are no problems before beginning the Live Upgrade process.

   If the `vxlustart -V` command reports success, proceed with running the `vxlustart` command.

   If the `vxlustart -V` command reports errors, correct the problem, and run the `vxlustart -V` command again.

   Note: This `vxlustart -V` command does not catch failures that are reported by Solaris Live Upgrade commands.

3. Run the `vxlustart` command to start the Live Upgrade for your Veritas product:

   ```
   # ./vxlustart -v -u target_os_version -U -d disk_name
   ```

4. Run the `installrp` command to upgrade your Veritas product:

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

5. Run the `vxlufinish` command to complete the Live Upgrade:

   - If the primary root disk is not encapsulated, run the following command:

     ```
     # ./vxlufinish -u target_os_version
     ```

   - If the primary root disk is encapsulated by VxVM, run the following command:

```
# ./vxlufinish -u target_os_version -g diskgroup
```

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

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

7 Verify that the alternate boot environment is active.

```
# lustatus
```

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

```
# gabconfig -a
```

Upgrading Veritas products using Live Upgrade from 5.1 SP1 to 5.1 SP1 RP1 with OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, or SF for Oracle RAC from 5.1 SP1 to 5.1 SP1 RP1 using Live Upgrade where OS upgrade is involved..

To upgrade your Veritas product using Live Upgrade

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

See your Veritas product 5.1 SP1 Installation Guide for more information.

2 Run the `vxlustart -V` command to ensure there are no problems before beginning the Live Upgrade process.

If the `vxlustart -V` command reports success, proceed with running the `vxlustart` command.

If the `vxlustart -V` command reports errors, correct the problem, and run the `vxlustart -V` command again.

Note: This `vxlustart -V` command does not catch failures that are reported by Solaris Live Upgrade commands.

---

3 Run the `vxlustart` command to start the Live Upgrade for your Veritas product:

```
# ./vxlustart -v -u target_os_version -s osimage_path -d disk_name
```

4 If you are upgrading from Solaris 9 Update x to Solaris 10 Update y, uninstall Veritas products on ABE, else go to Step 6.

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

5 Install Veritas products again on ABE.

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

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

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

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

8 Run the `vxlufinish` command to complete the Live Upgrade:

- If the primary root disk is not encapsulated, run the following command:

  ```
  # ./vxlufinish -u target_os_version
  ```

- If the primary root disk is encapsulated by VxVM, run the following command:

  ```
  # ./vxlufinish -u target_os_version -g diskgroup
  ```

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

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

10 Verify that the alternate boot environment is active.

```
# lustatus
```

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

```
# gabconfig -a
```

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

Performing a rolling upgrade using the installer

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

About rolling upgrades

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

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

You can perform a rolling upgrade for the following products:

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

You can perform a rolling upgrade from 5.1 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.

Limitation: During VCS and agents upgrade, you must bring down the application High Availability (HA) for several minutes. This does not affect the application

running on the cluster. You can restore the application's high availability after VCS and the agent packages are upgraded.

Performing a rolling upgrade on kernel packages: phase 1

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

To perform the rolling upgrade on kernel packages: phase 1

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

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

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

4. Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
5. The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
6. The installer loads new kernel modules.
7. The installer starts all the relevant processes and brings all the service groups online.
8. If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step 9.
9. Before you proceed to phase 2, complete step 2 to step 7 on the second subcluster.

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

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

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

1. 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 packages.
4. The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.
5. Verify the cluster's status:

   ```
   # hastatus -sum
   ```

Verifying software versions

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

```
# pkginfo -l  pkgname
```

Chapter 5

Rolling back

This chapter includes the following topics:

About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1

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

Roll back of version 5.1 SP1 RP1 to the 5.1 SP1 release is supported for the following products:

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation Cluster File System (SFCFS)
- Veritas Storage Foundation for Oracle RAC (SF for Oracle RAC)
- Veritas Cluster Server (VCS)
- Symantec VirtualStore (SVS)
- Dynamic Multi-Pathing (DMP)

Rolling back using the uninstallrp script

Use the following procedure to roll back from any Veritas product to 5.1 SP1 using the `uninstallrp` script.

Note: If any of the systems that you plan to roll back have encapsulated boot disks, you must reboot them.

To roll back

1 Run the `uninstallrp` command. On each node in the cluster, type:

```
# ./uninstallrp
```

2 If you performed a roll back on a system that has an encapsualted boot disk, you must reboot the system. After reboot, you may need to run `hagrp -list Frozen=1` to get the frozen SG list . Then run `hagrp -unfreeze <group> -persistent` to unfreeze all the frozen SGs manually.

Rolling back manually

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

- Rolling back Storage Foundation or Storage Foundation and High Availability manually
- Rolling back Storage Foundation Cluster File System manually
- Rolling back Storage Foundation for Oracle RAC manually
- Rolling back Veritas Cluster Server manually
- Rolling back Symantec VirtualStore manually
- Rolling back Dynamic Multi-Pathing manually

Note: You must reboot systems that you roll back manually at the end of the roll back procedure.

Rolling back Storage Foundation or Storage Foundation and High Availability manually

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

To roll back SF or SFHA

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# df -v /
```

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

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

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

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

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

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

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

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

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

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

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

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

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

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

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

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

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

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

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

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

- For Solaris 9:

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

- For Solaris 10:

```
# svcadm disable -t vcs
```

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

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

11 Unmount `/dev/odm`:

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

12 Unload the ODM module:

- For Solaris 9:

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

- For Solaris 10:

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

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

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

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

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

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SF 5.1 SP1 RP1 patches.

- Get the list of 5.1 SP1 RP1 patches, type:

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

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

```
# patchrm 143287-02
```

Rolling back Storage Foundation Cluster File System manually

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

To roll back SFCFS or SFCFS HA manually

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# df -v /
```

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

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

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

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

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

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

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

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

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

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

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

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

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

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

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

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

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

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

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

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

- For Solaris 9:

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

- For Solaris 10:

```
# svcadm disable -t vcs
```

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

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

11 Unmount `/dev/odm`:

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

12 Unload the ODM module:

- For Solaris 9:

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

- For Solaris 10:

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

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

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

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

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

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SFCFS 5.1 SP1 RP1 patches.

- Get the list of 5.1 SP1 RP1 patches, type:

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

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

```
# patchrm 143287-02
```

Rolling back Storage Foundation for Oracle RAC manually

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

To roll back SF for Oracle RAC manually

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

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

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

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

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

   - For 10gR2 or 11gR1:

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

   - For 11gR2:

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

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

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

   ```
   # fuser -c mount_point
   ```

4. Unmount CFS file systems that are not under VCS control.

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

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

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

     ```
     # umount mount_point
     ```

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

On any node execute following command to stop VCS:

```
# hastop -all
```

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

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

```
# fuser -c mount_point
```

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

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

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

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

```
# umount mount_point
```

8 To stop the process, type:

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

9 Remove the SF for Oracle RAC 5.1 SP1 RP1 patches.

- Get the list of 5.1 SP1 RP1 patches, type:

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

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

```
# patchrm 143287-03
```

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

11 Reboot the nodes point.

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

Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 5.1 SP1 RP1 to VCS 5.1 SP1 on your cluster manually. To uninstall VCS, see the *Veritas Cluster Server Installation Guide*.

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

To roll back VCS manually

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

   ```
   # hagrp -state
   ```

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

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

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

   ```
   # haconf -makerw
   ```

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

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

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

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

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

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

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

7 Shut down VCS. On any node, type:

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

8 Shut down CmdServer. On each node, type:

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

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

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

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

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

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

  ```
  # zoneadm list -icv
  ```

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

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

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

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

---

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

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

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

- Identify the vxfen kernel module, for example:

  ```
  # modinfo|grep vxfen
  210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.1SP1 RP1)
  ```

- Unload vxfen using the module number.

```
# modunload -i 210
```

13 Unconfigure GAB. On each node, type:

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

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

- Identify the GAB kernel module. For example:

```
# modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.1SP1 RP1)
```

- Unload GAB using the module number:

```
# modunload -i 149
```

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

- Type:

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

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

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

- Identify the LLT kernel module. For example:

```
# modinfo | grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.1SP1 RP1)
```

- Unload LLT using the module number:

```
# modunload -i 147
```

17 Remove the VCS 5.1 SP1 RP1 patches. On each node, perform the following steps:

- Get the list of 5.1 SP1 RP1 patches, type:

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

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

```
# patchrm 143287-02
```

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

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

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

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

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

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

21 After VCS has started, perform the following steps:

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

```
# hastatus -summary
```

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

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

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

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

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

where *system* is the node name.

Rolling back Symantec VirtualStore manually

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

To roll back SVS manually

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# df -v /
```

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

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

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

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

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

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

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

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

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

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

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

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

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

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

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

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

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

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

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

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

- For Solaris 9:

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

- For Solaris 10:

```
# svcadm disable -t vcs
```

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

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

11 Unmount `/dev/odm`:

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

12 Unload the ODM module:

- For Solaris 9:

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

- For Solaris 10:

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

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

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

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

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SVS 5.1 SP1 RP1 patches.

- Get the list of 5.1 SP1 RP1 patches, type:

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

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

```
# patchrm 143287-02
```

Rolling back Dynamic Multi-Pathing manually

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

To roll back DMP manually

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# df -v /
```

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

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

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

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

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

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

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

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

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

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

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

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

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

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

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

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

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

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

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

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

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

- For Solaris 9:

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

- For Solaris 10:

```
# svcadm disable -t vcs
```

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

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

11 Unmount `/dev/odm`:

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

12 Unload the ODM module:

- For Solaris 9:

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

- For Solaris 10:

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

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

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

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

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

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the DMP 5.1 SP1 RP1 patches.

- Get the list of 5.1 SP1 RP1 patches, type:

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

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

```
# patchrm 143287-02
```