Veritas Storage Foundation™ and High Availability Solutions Release Notes

Linux

5.1 Service Pack 1 Rolling Patch 2

November 2011

Symantec™

Veritas Storage Foundation and High Availability Solutions Release Notes 5.1 Service Pack 1 Rolling Patch 2

Product version: 5.1 SP1 RP2

Document version: 5.1SP1RP2.1

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

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

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec Web site.

https://sort.symantec.com/documents

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:

doc_feedback@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

Technical Support 4

Chapter 1 About Veritas Storage Foundation and High Availability Solutions 11

Introduction 11
About the installrp script 12
The installrp script options 13
Changes introduced in 5.1 SP1 RP2 16
Changes related to Storage Foundation and High Availability 16
Changes related to installing, upgrading and rolling back 16
Changes related to Veritas Volume Manager 17
Changes related to Veritas Cluster Server 17
System requirements 23
Supported Linux operating systems 23
Database requirements 24
Recommended memory and swap space 24
List of products 25
Fixed issues 25
Veritas Volume Manager fixed issues 26
Veritas File System fixed issues 33
Veritas Storage Foundation Cluster File System fixed issues 39
Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues 40
Veritas Storage Foundation for Oracle RAC fixed issues 41
Veritas Cluster Server fixed issues 41
Storage Foundation Manager fixed issues 47
Veritas Storage Foundation for Databases (SFDB) tools fixed issues 47
Known issues 48
Issues related to installation 48
Veritas Storage Foundation known issues 50
Veritas Volume Manager known issues 54
Veritas File System known issues 59

Veritas Cluster Server known issues 63
Veritas Volume Replicator known issues 68
Veritas Storage Foundation for Databases (SFDB) tools known issues 75
Veritas Storage Foundation for Oracle RAC known issues 78
Software limitations 81
Veritas Storage Foundation software limitations 81
Veritas Volume Manager software limitations 81
Veritas File System software limitations 82
Veritas Volume Replicator software limitations 83
Veritas Storage Foundation for Databases tools software limitations 85
Documentation errata 85
Veritas Cluster Server Administrator's Guide (2444653) 86
Veritas Installation Guides (2521411) 86
List of RPMs 86
Downloading the 5.1 SP1 RP2 archive 90

Chapter 2 Installing the products for the first time 91

Installing the Veritas software using the script-based installer 91
Installing Veritas software using the Web-based installer 92
Starting the Veritas Web-based installer 93
Obtaining a security exception on Mozilla Firefox 93
Installing 5.1 SP1 RP2 with the Veritas Web-based installer 93
Upgrading Veritas product with the Veritas Web-based installer 95

Chapter 3 Upgrading to 5.1 SP1 RP2 97

Prerequisites for upgrading to 5.1 SP1 RP2 97
Downloading required software to upgrade to 5.1 SP1 RP2 97
Supported upgrade paths 98
Upgrading to 5.1 SP1 RP2 98
Performing a full upgrade to 5.1 SP1 RP2 on a cluster 99
Upgrading to 5.1 SP1 RP2 on a standalone system 107
Upgrading to 5.1 SP1 RP2 on a system that has encapsulated boot disk 110
Performing a rolling upgrade using the installer 110
Upgrading the operating system 111
Verifying software versions 112

Chapter 4 Removing Veritas Storage Foundation and High Availability Solutions 113

About removing Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 113
Uninstalling Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 114
Preparing to uninstall Veritas Storage Foundation Cluster File System 114
Uninstalling Veritas Storage Foundation Cluster File System 116
Uninstalling Veritas Storage Foundation for Oracle RAC 117
Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP2 119

Chapter 1

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- Introduction
- About the installrp script
- Changes introduced in 5.1 SP1 RP2
- System requirements
- List of products
- Fixed issues
- Known issues
- Software limitations
- Documentation errata
- List of RPMs
- Downloading the 5.1 SP1 RP2 archive

Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Rolling Patch 2 release.

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

http://www.symantec.com/docs/TECH75506

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

http://www.symantec.com/docs/TECH74012

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.

This rolling patch applies to the following releases of Storage Foundation and High Availability products:

- Storage Foundation and High Availability Solutions 5.1 SP1
- Storage Foundation and High Availability Solutions 5.1 SP1 RP1
- Storage Foundation and High Availability Solutions 5.1 SP1 PR2
- VirtualStore 5.1 SP1 PR3

This rolling patch is available as:

- 5.1 SP1 RP2
- 5.1 SP1 PR3 RP2

Given that this rolling patch applies to the previously released 5.1 SP1 platform RP releases, Symantec does not plan on the following releases:

- 5.1 SP1 PR2 RP1
- 5.1 SP1 PR3 RP1

About the installrp script

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 provides an upgrade script.

See "Supported upgrade paths" on page 98.

Symantec recommends that you use the 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 1-1 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.
`[ -postcheck ]`	Use the `-postcheck` option after an installation or upgrade to help you determine installation-related problems and provide troubleshooting information.
`[ -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.
`[ -tmppath tmp_path ]`	Use the `-tmppath` option to select a directory other than `/var/tmp` as the working directory for `installrp`. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
`[ -hostfile hostfile_path ]`	Use the `-hostfile` option to specify the location of a file containing the system names for installer.
`[ -keyfile ssh_key_file ]`	Use the `-keyfile` option to specify a key file for SSH. When you use this option the `-i ssh_key_file` is passed to every SSH invocation.

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

Command Line Option	Function
[-kickstart *dir_path*	Use to produce a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec rpms in the correct order for installing, in a format that can be used for Kickstart installations. The dir_path indicates the path to the directory in which to create the file.
[-patchpath *patch_path*]	Use the -patchpath option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by installrp.

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

Command Line Option	Function
`[ -rsh | -redirect | -listpatches | -makeresponsefile | -pkginfo | -serial | -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 `-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. 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 `-version` option to have the installer check and report the installed products and their versions. Identifies the installed and missing RPM and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPM and patches where applicable.

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

Command Line Option	Function
`[-upgrade_kernelpkgs | -upgrade_nonkernelpkgs]`	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.

Changes introduced in 5.1 SP1 RP2

This section lists the changes in 5.1 SP1 RP2.

Changes related to Storage Foundation and High Availability

Storage Foundation and High Availability includes the following changes in 5.1 SP1 RP2:

New README files provide detailed information on the included patches

This release includes README_SYMC files that provide detailed information on the patches included in this release. These README_SYMC files provide the following information:

- Patch IDs, incidents fixed through the patch
- Symptom, description, and resolution for the addressed issues

The README_SYMC files are available in the following directory:

- `<archictecture>/rpms`

The Rolling Patch Release Notes continue to provide a summary of fixed issues.

Changes related to installing, upgrading and rolling back

The following changes are related to installing, upgrading and rolling back of the product in 5.1 SP1 RP2 release.

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

To determine a product's version, use the `-version` option with the `installrp` script. After you install 5.1 SP1 RP2, only the `installrp` script can detect 5.1 SP1 RP2 versions.

Changes related to Veritas Volume Manager

Veritas Volume Manager includes the following changes in 5.1 SP1 RP2.

ASM co-existence now enabled by default

The VxVM and ASM co-existence is now enabled by default. VxVM now identifies ASM disks automatically.

support vxcdsconvert utility for EFI disks.

The vxcdsconvert utility is now supported for EFI disks.

Changes related to Veritas Cluster Server

Veritas Cluster Server includes the following changes in 5.1 SP1 RP2:

Allow Fencing to start when a majority of the coordination points are available

When a node starts up, the following conditions must occur before vxfen starts on that node:

- vxfen must be able to get the Universal Unique Identifier (UUID) or serial number of each coordination point specified in the `/etc/vxfenmode` file.
- vxfen must be able to register the node with a majority of coordination points (CPs) specified in the `/etc/vxfenmode` file.

However, due to accessibility issues, if vxfen fails to get the UUID or serial number of one of the specified CPs, vxfen treats it as a fatal failure. The node cannot then join a cluster or start a cluster. As a result, every coordination point becomes a potential single point of failure, and compromises high availability (HA).

Symantec has modified the fencing module to fix the issue. Each VCS node now stores the UUIDs or serial numbers of all the CPs that the node registers with. As a result, if a node is later unable to access a specified coordination point, vxfen can use the stored UUIDs/serial numbers.

By design, the fix works only when a majority of CPs are accessible to the node when the node starts. At the time of a fencing race, the racer needs to have its keys registered on a majority of CPs in order to be able to win the race. In order to enable this, fencing is designed not to start if a majority of CPs are not available at the time of startup.

This fix applies only to clusters that use customized fencing.

As part of the fix, Symantec has introduced two optional attributes to the `/etc/vxfenmode` file.

- `db_ignore_list` Specifies the type(s) of CPs for which a node must not store the UUID/serial number. To specify multiple values, use a comma-separated list. vxfen supports the values `none`, `disk`, and `server`.

  Note: By default, this feature is available only for coordination point servers. To turn it on for disks, you must set the value of the db_ignore_list to `none`.

- `db_entries_limit` Specifies the maximum number of UUIDs/serial numbers that a node can store. The default value for this attribute is 1000. If the default value is used, the node approximately requires 1MB of disk space to store the UUIDs/serial numbers.

Veritas Cluster Server share agent hostname comparison is now case insensitive

The share agent client comparison is now case insensitive.

Added XFS file system support to the mount agent

The mount agent now supports XFS file systems.

Added ext4 file system support to the mount agent

The mount agent now supports ext4 file systems.

Modified owner.vfd to compare only the uid and gid of the user

Modified the `owner.vfd` file to compare only the `uid` and `gid` of the user, instead of the complete output string of `id` command.

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 VCS5.1 SP1 RP2, please ensure that the new `Db2udbTypes.cf` file is used which contains the definition of **IMFRegList** as above.

The full value displays when you use hares -display

For `hares -display`, the maximum number of characters was 20. Any attribute value that was more than 20 characters would be truncated.

Symantec has removed the 20 characters limit. Now the full value displays when you use `hares -display`.

Restrict the max value of FencingWeight to 9999

When Preferred Fencing is enabled, the max weight value that you can assign to the FencingWeight attribute is 9999. `had` adds 1 to every weight that you assign for each node. If the value of FencingWeight is set to 10000, VCS fails to set the node weight to 10001 because the maximum node weight accepted by vxfen driver is 10000.

Below message displays if the value of FencingWeight that you assign is equal to or more than 10000.

```
VCS WARNING V-16-1-50003 FencingWeight should be an \
integer between [0..9999] inclusive
```

Changes to Application Agent

The following are changes about Application Agent for shared disk support

Shared disk support

Symantec enhanced the Application agent to support use of shared disks for StartProgram, StopProgram, CleanProgram and MonitorProgram attributes.

Support for UNIX style return values in MonitorProgram

The Application agent handles the standard UNIX style return values, that is, "0" for success and "1" for failure, and now reports the resource state based on following set of values:

100 or 1 --> OFFLINE

101 to 110 or 0 --> ONLINE

Any other value --> UNKNOWN.

Validating the user that is specified in the User attribute

The Application agent is modified to validate the user name configured in the User attribute. If the user does not exist on the system, the agent reports the state of a configured resource on that system as UNKNOWN.

User home directory validation for the user that is specified in the User attribute

The Application agent validates the configured user for a home directory. If a home directory is not set for a configured user, then the Application agent will report the resource state as UNKNOWN.

System requirements

This section describes the system requirements for this release

Supported Linux operating systems

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

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

http://www.symantec.com/docs/TECH75506

The Veritas Storage Foundation (SF) 5.1 SP1 RP2 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
  Red Hat Enterprise Linux 6 (RHEL 6) on AMD Opteron or Intel Xeon EM64T (x86_64)
  Red Hat Enterprise Linux 6 (RHEL 6) with Update 1 on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel), SP3 (2.6.16.60-0.54.5) or SP4 (2.6.16.60-0.85.1), on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) with SP1 (2.6.32.12-0.7 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)

  Note: SFCFSRAC is not supported on SLES 11 SP1.

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

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

Symantec supports only Oracle Enterprise Linux, Red Hat Enterprise Linux (RHEL Compatible Mode only), and SUSE Linux Enterprise Server distributed kernel binaries.

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

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

http://www.symantec.com/docs/TECH75506

VMware Environment

For information about the use of this product in a VMware Environment, refer to http://www.symantec.com/docs/TECH51941

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

Database requirements

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

http://www.symantec.com/docs/TECH74389

Note: Veritas Storage Foundation (SF) and Veritas Storage Foundation Cluster File System (SFCFS) do not support running SFDB tools with DB2 and Sybase, but they support running Oracle, DB2, and Sybase on VxFS and VxVM.

SF Oracle RAC will announce support for RHEL6 once Oracle supports it. For latest information on support for Oracle database versions with SF Oracle RAC, see the Technical Support TechNote:

http://www.symantec.com/docs/TECH44807

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
  - For a minimum of 256 MB for 1 node and a maximum of 1 GB of swap space for 8 or more nodes

List of products

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

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

Fixed issues

This section describes the issues fixed in this release.

See the `README_SYMC.xxxxx-xx` files in the `<archictecture>/rpms` directory on the installation media for the symptom, description, and resolution of the fixed issue.

Veritas Volume Manager fixed issues

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP2

Table 1-2 describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP2.

Table 1-2 Veritas Volume Manager 5.1 SP1 RP2 fixed issues

Fixed issues	Description
2484685	Race between two vol_subdisk sios while doing done processing which causes one thread to free sio_fsvm_priv before other thread accesses it
2480600	I/O permanent hung on master node when IO size larger than 512K, and 32+ threads write in parallel
2440349	DCO volume may grow into any 'site' even when 'alloc=site:xxxx' is specified by a list of 'site' to be limited
2431470	vxpfto uses DM name when calling vxdisk, but vxdisk will match DA name first and thus cause corruption
2431423	CVR: Panic in vol_mv_commit_check after I/O error on DCM
2428875	I/O on both nodes (wait for the DCM flush started), and crash the slave node, lead to the master reconfiguration hang
2428631	Allow same fence key to be used for all Disk groups
2425722	vxsd move operation failed for disk size greater than or equal to 2TB
2425551	IO hung for 6 mintues when reboot the slave node, if there is I/O on both master and slave
2424833	while autosync_deport#2 primary logowner hits ted assert nmcom_send_msg_tcp
2422058	The VxVM diskgroup can NOT import with I/O fencing enabled of both dmp and raw mode
2421067	Vxconfigd hung in both nodes of primary
2419348	DMP panic: race between dmp reconfig and dmp pass through ioctl
2413904	Multiple issues are seen while performing Dynamic LUN reconfiguration

Table 1-2 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2411698	VVR:iohang: On I/O to both master and slave
2410845	Lots of 'reservation conflict' messages seen on 51SP1RP1P1 clusters with XIV arrays
2408771	vxconfigd does not scan and discover all the storage device; some storage devices are skipped
2407192	Application I/O hangs because of race between CVM reconfiguration and Log-owner change protocol
2406292	Panic in vol_subdisksio_delete()
2400654	Stale arrayinfo file can cause vxdmpadm commands to hang
2400076	vxconfigd produced kernel panic when you run "vxinstall" command
2396293	I/Os loaded, sanboot failed with vxconfigd core dump
2388725	Panic in dmp_get_dmpsymbols when attempting to load an APM
2387993	While testing including/excluding libvxppso vxconfigd goes into disabled mode
2386120	Enhancement request to add diagnostic logging to help triage a CVM master takeover failure situation
2385694	IO hung if the slave node rebooted
2385680	vol_rv_async_childdone+1147
2383158	VVR: vxio panic in vol_rv_mdship_srv_done+680
2379029	Changing of enclosure name is not working for all devices in enclosure
2369786	VVR:A deadloop about NM_ERR_HEADR_IO
2369177	DDL: do_diskio function should be able to handle offset greater than 2 TB
2365951	Growto failing with error V-5-1-10128 Unexpected kernel error in configuration update
2364253	VVR: Kernel memory is leaked on VVR secondary while using SO snapshots
2359814	vxconfigbackup doesn't handle errors well

Table 1-2 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2358321	Remove usage of __invalidate_device() from VxVM Symbol is no longer in kABI whitelist
2357798	CVR:Memory leak due to unfreed vol_ru_update structure
2357507	In presence of large number of NR (Not-Ready) devices, server panics due to NMI triggered and when DMP continuously generates large no of path disable/enable events
2356744	VxVM script daemons should not allow its duplication instance in itself
2349352	During LUN provisioning in single path IO mode environment a data corruption is observed
2346470	Excluding and including a LUN in a loop triggers a huge memory leak
2337694	TP "vxdisk -o thin list" showing size 0 for over 2TB LUNs on RHEL5
2337353	vxdmpadm include vxvm dmpnodename=*emcpower#* includes all excluded dmpnodes along with the requested one
2334534	In CVM environment, vxconfigd level join is hung when Master returns error "VE_NO_JOINERS" to a joining node and cluster nidmap is changed in new reconfiguration
2323925	If rootdisk is encapsulated and if install-db is present, clear warning should be displayed on system boot
2322752	Duplicate DA records seen for NR devices upon restart of vxconfigd
2320917	vxconfigd core dump and lost dg config after removing volume and disk on thin reclaim LUN
2317703	Vxesd/Vxconfigd leaks file descriptors
2316297	After applying 51SP1RP1 error message "Device is in use" appears during boot time
2299670	Disk Groups created on EFI LUNs do not auto import at boot time using VxVM version 51SP1 and later
2286559	kernel heap corruption detected panic after array controller reboot
2263317	CLONE: Diskgroup import with dgid needs to be clearly documented in manual for the case in which original dg was destroyed and cloned disks are present

Table 1-2 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated
2255182	Handling misconfiguration of CLARiiON array reporting one failovermode value through one HBA and different from other HBA
2253970	Support per-disk maxiosize for private region I/Os
2253552	Leak in vxsfdefault_parsey at function vxsf_getdefault (*val)
2249113	vol_ru_recover_primlog_done return the same start address to be read from SRL, if the dummy update is greater than MAX_WRITE
2248730	vxdg import command hangs as vxrecover daemon (spawned by vxdg) doesn't close standard error stream
2242268	panic in voldrl_unlog
2240056	vxdg move' transaction not completing and backups fail
2237089	vxrecover might start the recovery of data volumes before the recovery of the associated cache volume is recovered
2234821	etrack 1946267 - DMP can't detect the re-enabled os device status does not work on RHEL5
2232411	supporting NetApp Metro Cluster
2228531	cvm master vxconfigd process hung in vol_klog_lock()
2218470	Some of the VxVM init scripts need to be compliant to the Linux Standard Base
2205108	SVS 51SP1: vxconfigd clubbing all luns in a single dmpnode
2204752	Multiple VM commands succeed but throw "GPT entries checksum mismatch" error message for hpdisk format
2200670	vxattachd does not recover disks if disk group is not imported
2197254	While creating volumes on thinrclm disks, the option "logtype=none" does not work with vxassist command
2196918	Snapshot creation with cachesize fails, as it doesn't take into account diskgroup alignment
2196480	The disk initialization failed due to wrong number of cylinders reported in devintf_disk_geom_raw() gotten from raw geometry

Table 1-2 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2194685	vxconfigd daemon core dump during array side switch ports disable and re-enable
2193429	IO policy not getting preserved when vold is restarted and migration from one devlist to other is taking place
2190020	SUSE complains dmp_deamon applying 1m continuous memory paging is too large
2179259	DMP SCSI bypass needs to be enhanced to handle I/O greater than 2 TB
2165394	CLONE: dg imported by selecting wrong disks After destroying original dg, when try to import clone devices without useclonedev option with dgname, then it import dg with original disks
2154287	Improve handling of Not-Ready(NR)devices which are triggering "VxVM vxdmp V-5-3-1062 dmp_restore_node: Unstable path" messages
2152830	In multilevel clone disks environment, regular DG import should be handled properly and in case of DG import failure, it should report correct error message
2144775	Failoverpolicy "local" is not getting preserved after upgrade from 51RP1/Sles10Sp2 to 51Sp1/Sles10Sp3
2139179	SSB check invalid when lun copy
2094672	CVR: vxconfigd on master hangs while reconfig is running in cvr stress with 8 users
2033909	In SF-RAC configuration, IO hung after disable secondary path of A/PG array Fujitsu ETERNUS3000
1791397	VVR:RU thread keeps spinning sending START_UPDATE message repeatedly to the secondary
1675599	Memory leaks in DDL and ASLs

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 1-3 Veritas Volume Manager 5.1 SP1 RP1 fixed issues

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

Table 1-3 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

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

Table 1-3 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

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

Veritas File System fixed issues

This section describes Veritas File System fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas File System: Issues fixed in 5.1 SP1 RP2

Table 1-4 describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP2.

Table 1-4 Veritas File System fixed issues

Fixed issues	Description
2529356 (2340953)	cfs.stress.enterprise hit an assert f:vx_iget:1a.
2508164 (2481984)	file system will hang if a customer creates 400 shares
2494464 (2247387)	LM stress.S3 test hit an assert "vx_ino_update:2"
2486597 (2486589)	threads blocked behind vx_ireuse_steal

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

Fixed issues	Description
2484815 (2440584)	node panic in vx_sync() during shutdown
2482344 (2424240)	Dedup ioctl sharing extents incorrectly under certain scenarios
2482337 (2431674)	panic in vx_common_msgprint() via vx_inactive()
2480949 (2480935)	V-3-26626: File Change Log IOTEMP and ACCESSTEMP index creation failure for /vx/fsvm with message Argument list too long
2430679 (1892045)	Improve the memory allocation for per-cpu data.
2427281 (2413172)	There is a priority 1 issue reported by AXA Rosenburg for Filestore replication and issue seems related to VxFS
2427269 (2399228)	TRuncate up size updates can be missed
2426039 (2412604)	It does not work when set homedir user softlimit numspace quota after generate data
2425439 (2242630)	Remove limits on inode and buffer cache sizes
2425429 (2422574)	Reboot one node and the node can't mount file system , after turn on the homedir quota on
2420060 (2403126)	cfs recovery didn't finished timely in the primary node after one slave left.
2418819 (2283893)	Add functionality of free space defragmentation through fsadm.
2412181 (2372093)	new fsadm -C hung

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

Fixed issues	Description
2412179 (2387609)	User quota corruption
2412177 (2371710)	user quota information corrupts on 5.1SP1
2412169 (2371903)	newline in vx_osdep.c: snprintf(cmp->cm_name, sizeof(cmp->cm_name), "vxclonefs-%d" breaks native LVM(pvs)
2412029 (2384831)	vxfs panic in iput() from vx_softcnt_flush() ,after filesystem full fsck,and run reboot
2402643 (2399178)	fsck : pass2c needs performance enhancements
2398798 (2374887)	Accessing FS hung. FS marked full fsck after reboot of node.
2386483 (2374887)	Accessing FS hung. FS marked full fsck after reboot of node.
2373565 (2283315)	cfs-stress_S5 hits assert of "f:vx_reorg_emap:10 via vx_extmap_reorg"
2368738 (2368737)	RCQ processing code should set FULLFSCK flag if it finds a corrupt indirect block.
2360821 (1956458)	fsckpt_fbmap for changed blocks failed with ENXIO due to inode mapped to hole in ILIST of down stream checkpoint
2360819 (2337470)	In the process of shrink fs, the fs out of inodes, fs version is 5.0MP4HF*
2360817 (2332460)	vxedquota slow on some systems
2341007 (2300682)	Question about IOTemp on fsppadm query

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

Fixed issues	Description
2340839 (2316793)	After removing files df command takes 10 seconds to complete
2340834 (2302426)	Unaligned Reference Fault in vx_copy_getemap_structs
2340831 (2272072)	Threads stuck in vx_rwsleep_rec_lock_em
2340825 (2290800)	investigation on ilist HOLE
2340817 (2192895)	Panic while set/get acls - possible race condition
2340799 (2059611)	Panic in vx_unlockmap() due to NULL ml_tranp
2340741 (2282201)	vxdump core dumped whilst backing up layout 7 local VxFS file system
2338010 (2337737)	killing IOs to a CFS and ls command to the same CFS is hanging.
2329893 (2316094)	There was discrepancy between vxi_bcache_maxkbyte and vx_bc_bufhwm.
2329887 (2253938)	EAU delegation timeouts
2320049 (2419991)	ncheck: no way to limit output to specific filesets, as with limiting output to specific inodes
2320044 (2419989)	ncheck -i does not limit output to the specified inodes when using -o device/block/sector
2311490 (2074806)	dm_punch_hole request does not invalidate pages

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

Fixed issues	Description
2296277 (2296107)	Operation not applicable appear on fsppadm query result
2280552 (2246579)	Panic at getblk() when growing a full filesystem with fsadm
2280386 (2061177)	fsadm -de' command erroring with 'bad file number' on filesystem(s) on 5.0MP3RP1
2275543 (1475345)	write() system call hangs for over 10 seconds on VxFS 3.5 on 11.23
2257904 (2251223)	df -h after removing files takes 10 seconds
2255786 (2253617)	LM stress aborted due to "run_fsck : Failed to full fsck cleanly".
2249658 (2220300)	vx_sched' is hogging CPU resources.
2247299 (2161379)	repeated hangs in vx_event_wait()
2243064 (2111921)	Linux DB2 readv() behavior with CIO/DIO enabled
2243063 (1949445)	hang due to large number of files in a directory
2243061 (1296491)	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2239415 (2239412)	system panics while writing to cfs share exported as NFS to ESX server4.1.
2169326 (2169324)	Test LM-stress_S5 hits an assert of "f:vx_idelxwri_off:5a vai vx_trunc_tran"

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

Fixed issues	Description
1296491	Fixed issues seen during a force unmount of a parent cluster file system while a child was being mounted or umounted.
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.
2030119	fsppadm core dumps when analysing a badly formatted XML file, is resolved
2032525	Fixed the cause of an NFS stale file handle.
2061554	Sequential extents are now collated.
2111921	Improved the performance of VxFS file systems with concurrent I/O or direct I/O enabled.
2149659	Fixed the cause of an error that resulted during the truncate operation of a file with a shared extent in the presence of a Storage Checkpoint containing FileSnaps.
2162822	During online migration from ufs to vxfs, df command returns a non-zero return value.
2163084	The `listxattr`() call now uses rwlock.
2169273	During online migration, nfs export of the migrating file system leads to system panic
2177253	A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for vxupgrade purposes
2178147	Linking a IFSOC file now properly calls `vx_dotdot_op`(), which fixes the cause of a corrupted inode.
2181833	The `vxfilesnap` command no longer gives an incorrect error message on a Storage Checkpoint file system.
2184528	`fsck` no longer fails to repair corrupt directory blocks that have duplicate directory entries.

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

Fixed issues	Description
2194618	Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system
2198553	A forced unmount now properly clears the `bd_super` structure member.
2221623	Fixed a performance loss due to a delxwri_ilist spin lock with the default values for `vx_idelxwri_timelag`.
2226257	Fixed the cause of a system panic in the `in _list_add`() call, which led to corruption in the `vx_ftenter`() codepath when using named data streams.

Veritas Storage Foundation Cluster File System fixed issues

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP2

Table 1-6 describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP2.

Table 1-6 Veritas Storage Foundation Cluster File System fixed issues

Fixed issues	Description
2406572 (2146573)	qdetails performance downgraded

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 1-7 Veritas Storage Foundation Cluster File System fixed issues

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

Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues

This section describes Veritas Storage Foundation Cluster File System for Oracle RAC fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation Cluster File System for Oracle RAC: Issues fixed in 5.1 SP1 RP2

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

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

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

Veritas Storage Foundation for Oracle RAC fixed issues

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP2

Table 1-8 describes the incidents that are fixed in Veritas Storage Foundation for Oracle RAC in 5.1 SP1 RP2.

Table 1-8 Veritas Storage Foundation for Oracle RAC fixed issues

Fixed issues	Description
2429449	The cssd agent explicitly uses hard-coded string "cssd" as resource name.
2390892	Starting the VCSMM driver on two or more nodes in the cluster causes a memory leak in the vcsmm_set_cluster_proto function during memory allocation
2374987	Failed to remove original IP address by PrivNIC and MultiPrivNIC agents during failover/failback operation
2374970	CRSResource agent support for 11gR2

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

There are no fixed issues in this release.

Veritas Cluster Server fixed issues

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP2

Table 1-9 describes the incidents that are fixed in Veritas Cluster Server in 5.1 SP1 RP2.

Table 1-9 Veritas Cluster Server 5.1 SP1 RP2 fixed issues

Fixed Issues	Description
2416842	_had consuming over 99% CPU time. Multiple ha commands are hung in pollsys()
2411653	Add check for MAX message size in GAB
2407755	Application and Netlsnr Agents failing
2407653	In case of forceful unload of AMF module, Module reference count of 'vxfs'/'ext3' should be handled correctly.
2406748	We are able to register already online process for offline monitor with AMF.
2405780	Cable pull test fails when Mii is set to 0
2405391	LLT: The arp ack packet should include the nodename of the node.
2403851	AMF status is showing Module loaded but not configured.
2403782	Sybase agent scripts are setting incorrect path for cat command on linux.
2403633	ContainerInfo attribute should be allowed to be updated even when Group is not completely offline
2400485	Once vxfenconfig -c with mode A has returned EFAULT ("1036 Unable to configure..."), all subsequent runs of vxfenconfig -c with mode B fail with error EBADMSG ("1050 Mismatched modes...").
2400330	whyonlining does not behave as advertised in VCS 5.1SP1.
2399898	hagrp -switch of child group fails in 5.0MP3RP2 and later if 2 or more parent groups online on alternate.
2398807	VCS should be setting a soft limit for file descriptors in /opt/VRTSvcs/bin/vcsenv.
2394176	vxfenswap process hangs, "ps -ef" shows "vxfenconfig -o modify" on one node but not on other. "vxfenswap -a cancel" kills the stuck operation.
2386326	cannot configure fencing, vxfenadm prints same Serial Number for all LUNs which have more than 96 bytes of SCSI Inquiry data in page 0x83
2382592	Issue with displaying "ResourceInfo" Attribute of SRDF Resource using hares -display
2382493	Parent service group does not failover in case of online local firm dependency with child service group.

Table 1-9 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2382463	Had weight(1) is not added if we reach the boundary condition(10000) in System policy in CPS preffered fencing.
2382335	vxfentsthdw fails to choose the same fencing disk on two nodes.
2381083	Broadcast address 0.0.0.0 was set by IP-Agent
2372483	SambaServerAgent generated core dumps on FileStore 5.7.
2372072	User core for "hacf"
2366201	Enhanced Fencing to start when a majority of the coordination points are available.
2354932	hacli -cmd' triggers had coredump on 5.1SP1RP1 system
2330980	No notifications about resources should be sent to agents running on nodes already existing in SystemList of Group, when a node is added / deleted to SystemList.
2330045	RemoteGroup resource does not go offline when network fails.
2330041	VCS group dependencies do not online parallel parent group after upgrading SF 5.0MP3 RP2 to SF5.1SP1.
2318334	Oracle needs its database's $Oracle_home/lib library to be first in LD_LIBRARY_PATH before /usr/lib.
2301731	Panic in amf_lock() due to bad mutex during system shutdown.
2287061	When enabling the amf, cfsmount agent cannot start normally. The basic event registration with AMF driver is failing.
2276622	Cannot configure SCSI-3 fencing using RamSan DMP devices.
2271882	MonitorMethod attribute does not reflect IMF value without setting Listener attribute on Netlsnr Resource.
2253441	VCS should setup the right default netmask when NetMask attribute is not set
2528475	support IPMultiNIC/ IPMultiNICB type in preonline_ipc for VCS5.x.
2509515	The resource fails to go offline when Options attribute and class B address is used.

Table 1-9 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2483964	Monitor for Process resource faults right after online, though the process appears to be running correctly.
2483314	Oracle agent core dumps when large number of oracle instances are running. (Around 50)
2483044	had' crashed with SIGSEGV when asserting against gp->activecount()->gets32GL(nodename) == 0\", in "Resource.C" in check_failover function
2477372	LLT: reduce "lltd" CPU consumption by reducing the wakeup calls
2477296	Application service group did not fail over on node panic
2477280	Application resource is not failover when system reboot after Concurrency Violation
2439772	wac resource offline failed after network interruption on SFHA5.1RP2, Solaris 10
2438261	Failed to perform online migration from scsi raw to scsi dmp policy.
2426663	On OCPR from customized mode to scsi3 mode, vxfend does not terminate
2426572	Persistent resource is reported OFFLINE (not FAULTED) when system is added to group using hagrp -modify command
2423990	Application Agent is not working properly when nonexistent user is configured.
2382559	Online Migration fails with the message pI/O fencing does not appear to be configured on nodeq
2382460	Configuring fencing is successful with 3 disks even when single_cp=1 and formatting of warning messages required in vxfend_A.logo
2382452	Syntax errors while unconfiguring CP server using configure_cps.pl scripto.
2367721	Enable selinux permissive / enforcing for Virtual Fire Drill by modifying owner.vfd.
2366701	Query regarding usage of variable in VCS attributes
2366201	Allow Fencing to start when a majority of the coordination points are available.

Table 1-9 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2364875	Enhancing the Bundled agents to support the RHEL 6 environment.
2330047	VCS share agent hostname comparison is case sensitive.
2511385	Sybase online script marks the database as online before Database has recovered
2439695	VXFEN module gets loaded even though user chooses not to enable VXFEN.
2426572	Persistent resource is reported OFFLINE (not FAULTED) when system is added to group using hagrp -modify command
2411860	Various VCS service group switch failures
2407755	Application and Netlsnr Agents failing
2405514	Panic in amf_lock() due to bad mutex during system shutdown.
2400330	whyonlining does not behave as advertised in VCS 5.1SP1
2382452	Syntax errors while unconfiguring CP server using configure_cps.pl script
2372072	User core for "hacf"
2296172	Failover Policy does not take into consideration AutoFailover = 2 and SystemZones when the nodes within a SystemZone are brought down/rebooted.
2393939	Enhanced Apache agent version parsing to accommodate IBM HTTP server 7.0.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Cluster Server in 5.1SP1RP1 release.

Table 1-10 Veritas Cluster Server 5.1 SP1 RP1 fixed issues

Fixed Issues	Description
1949294	`fdsetup` can now correctly parse disk names containing characters such as "-".

Table 1-10 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Fixed Issues	Description
1949303	`fdsetup` no longer allows volume that are not part of the RVG, which fixes a possible cause of the RVGSnapshot agent failing.
2011536	Added IMF support for the db2udb agent.
2159991	Fixed an issue with messages in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2172181	Fixed an issue with AMF-related messages for the CAVF agent in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2179652	The monitor script of the db2udb agent can now handle empty attribute values.
2184205	Fixed an issue with HAD in which the parent service group did not fail over if the parent service group had an online local firm dependency with a child service group.
2194473	HAD no longer dumps core while overriding the static attribute to the resource level.
2205556	Fixed an issue with the offline EP of the DNS agent, which did not remove all A/AAAA records if OffDelRR=1 for multi-home records.
2205563	A clean EP now properly removes resource records when OffDelRR=1.
2205567	Fixed an issue in which having an attribute set to `master.vfd` caused the DNS agent to fail to query the DNS server.
2208675	There is now a return value check for broadcast ping in NIC/MultiNICA monitor, which fixes one possible cause of the MultiNic resource is going into the FAULTED state in IPv6 with the Link option configuration.
2208901	Fixed an issue with the RVGSnapshot agent.
2209337	Fixed an issue with VCSAPI where the RemoteGroup agent crashed if the VCSAPI log level was set to a non-zero value.
2214539	Fixed an issue in which rebooting a node sometimes set the intentonline of a group to 2, even if the group was online somewhere else. This caused the group to use the autostartlist and not perform a failover.
2217446	Fixed an issue that caused the installation of `VRTSvcsag` to fail.

Table 1-10 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Fixed Issues	Description
2218556	Fixed an issue in the `cpsadm` command in which it sometimes failed if LLT was not installed or configured on a single node cluster.
2218561	Fixed an issue in which MonitorTimeStats incorrectly showed 303 seconds intermittently.
2219955	Fixed an issue in which a split-brain condition occurred even when using VCS Steward.
2220749	Fixed an issue in which the Cluster Manager (Java Console) was not encrypting the `DBAPword` attribute of the Oracle Agent.
2241419	Fixed an issue in which halogin did not work in a secure environment where the root broker was not a VCS node.

Storage Foundation Manager fixed issues

This section describes Storage Foundation Manager fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Storage Foundation Manager: Issues fixed in 5.1 SP1 RP2

There are no Storage Foundation Manager fixed issues in this release.

Storage Foundation Manager: Issues fixed in 5.1 SP1 RP1

There are no Storage Foundation Manager fixed issues in this release.

Veritas Storage Foundation for Databases (SFDB) tools fixed issues

This section describes Veritas Storage Foundation for Databases (SFDB) tools fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP2

There are no SFDB fixed issues in 5.1 SP1 RP2.

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

Table 1-11 describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in this release.

Table 1-11 Storage Foundation for Databases fixed issues

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

Known issues

This section covers the known issues in this release.

Issues related to installation

This section describes the known issues during installation and upgrade.

Incorrect version listed after upgrading (2121881)

When you upgrade from Veritas product 5.1 SP1 to Veritas product 5.1 SP1 RP2, the previous version is incorrectly listed as 5.1.001.000

Incorrect error messages: error: failed to stat (2120567)

During installation, you may receive errors such as, "error: failed to stat /net: No such file or directory." Ignore this message. You are most likely to see this message on a node that has a mount record of `/net/x.x.x.x`. However, the `/net` directory is unavailable at the time of installation.

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

The installer displays the space that all the product RPM and patches needs. During migration some RPM are already installed and during migration some RPM 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.

Error messages in syslog (1630188)

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

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

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

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

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

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

Or

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

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

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

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

SFCFSHA upgrade shows partial upgrade warning

When you install 5.1 SFCFSHA and try to upgrade to SFCFSHA 5.1SP1 using the `./installsfcfs` command, you may receive a partial upgrade error message.

Workaround: Use the `./installer -upgrade` command instead of the `./installsfcfs` command.

installrp fails to install 5.1SP1RP2 when the root user shell is set to csh (2523643)

VCS installation fails if super user (root) logged-in is using C shell (csh). Currently the installer does not support c-shell (/usr/bin/csh).

Workaround: Change your super-user (root) shell to shell (/usr/bin/sh) and retry the installation.

Veritas Storage Foundation known issues

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

Oracle 11gR2 is not supported with the managed host client included in this release (2568403)

The Veritas Operations Manager (VOM) managed host client version 3.xxx that is included in this release is not supported with Oracle 11gR2.

Workaround: After installing this release, upgrade the managed host client to version 4.0 RU1 or later.

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

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

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

Or

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

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

db2exp may frequently dump core (1854459)

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

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

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

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

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

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

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

DBI1070I  Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

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

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

To communicate in a dual-stack environment

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

  127.0.0.1 swlx20-v6

  Or

  127.0.0.1 swlx20-v6.punipv6.com

  where *127.0.0.1* is the IPv4 loopback address.

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

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

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

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

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

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

Tools like `dbca` may hang during database creation.

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

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

After upgrading SF stack from 5.0 MP3 SP1 RP2 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.

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

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

Workaround

Reboot the system or to run the following command.

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

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

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

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

vxrestored daemon fails to restore disabled paths (1663167)

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

Workaround:

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

To enable the mpt_disable_hotplug_remove tunable

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

   ```
   options mptsas mpt_disable_hotplug_remove=0
   ```

2. Rebuild the initrd image:

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

3. Reboot the system.

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

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

Workaround:

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

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

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

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

Workaround:

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

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

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

Root disk encapsulation issue (1603309)

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

VxVM starts before OS device scan is done (1635274)

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

Workaround:

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

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

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

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

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

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

Workaround: There is no workaround for this issue.

Co-existence check might fail for CDS disks

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

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

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

Workaround: There is no workaround for this issue.

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

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

Changes in enclosure attributes are not persistent after an upgrade to VxVM 5.1 SP1 RP2 (2082414)

The Veritas Volume Manager (VxVM) 5.1 SP1 RP2 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 RP2, 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 RP2. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-12 shows the Hitachi arrays that have new array names.

Table 1-12 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 RP2. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

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

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

Work around:

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

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

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

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

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

When using EMC Powerpath with VxVM5.1 SP1 RP2on SLES11, set the fast_io_fail_tmo on the HBA port to any non-zero value that is less than the dev_loss_tmo value so as to avoid a panic in case a DDL device discovery is initiated by the `vxdisk scandisks` command or the `vxdctl enable` command immediately after loss of connectivity to the storage.

VxVM vxdg V-5-1-16063 is returned from 'vxdg rmdisk' when attempting to perform storage reclamation on Hitachi AMS 2500 array

See http://www.symantec.com/docs/TECH162709 for more information.

On 5.1SP1RP2, changing path attributes by using 'vxdmpadm' gives an error message "Failed to save path information persistently" (2433012)

On Linux system with "selinuxenabled" set to 1, setting path attributes using "vxdmpadm setattr path *pathname* pathtype=*attr*", CLI will set the path attribute but these changes may not remain persistent across reboots. Following error displays if CLI fails to save the changes persistently:

```
VxVM vxdmpadm ERROR V-5-1-14526 Failed to save path information
persistently
```

Workaround: Reset these attributes using vxdmpadm CLI, if system is rebooted.

Veritas File System known issues

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

Storage Checkpoints can exceed the quota limit (2102201)

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

- The Storage Checkpoint quota has been enabled.
- The Storage Checkpoint quota is not exceeded.
- A file content modification operation, including removing a file, needs to push some or all blocks of the file to the Storage Checkpoint.
- Number of blocks that need to be pushed to the Storage Checkpoint is enough to exceed Storage Checkpoint quota hard limit.

Workaround: There is no workaround for this issue.

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

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

Panic occurs when VxFS module parameter vxfs_hproc_ext is set to 1 and you attempt to mount a clone promoted file system (2163931)

A system panic occurs if the following two conditions are met:

- The VxFS module parameter `vxfs_hproc_ext` is set to 1.
- A clone is promoted as a primary using the `fsckpt_restore` command, and then you attempt to mount the promoted file system.

Workaround: There is no workaround for this issue.

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

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

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

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

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

```
# lsmod | grep freevxfs
```

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

```
# rmmod freevxfs
```

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

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

Workaround: Force unmount the file system:

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

Possible write performance degradation with VxFS local mounts

Some applications that allocate large files without explicit preallocation may exhibit reduced performance with the VxFS 5.1 release compared to the VxFS 5.0 MP3 release due to a change in the default setting for the tunable `max_seqio_extent_size`. One such application is DB2. Hosting DB2 data on a single file system extent maximizes the potential for sequential pre-fetch processing. When DB2 detects an application performing sequential reads against database data, DB2 begins to read ahead and pre-stage data in cache using efficient sequential physical I/Os. If a file contains many extents, then pre-fetch processing is continually interrupted, nullifying the benefits. A larger `max_seqio_extent_size` value reduces the number of extents for DB2 data when adding a data file into a tablespace without explicit preallocation.

The `max_seqio_extent_size` tunable controls the amount of space that VxFS automatically preallocates to files that are allocated by sequential writes. Prior to the 5.0 MP3 release, the default setting for this tunable was 2048 file system blocks. In the 5.0 MP3 release, the default was changed to the number of file system blocks equaling 1 GB. In the 5.1 release, the default value was restored to the original 2048 blocks.

The default value of `max_seqio_extent_size` was increased in 5.0 MP3 to increase the chance that VxFS will allocate the space for large files contiguously, which tends to reduce fragmentation and increase application performance. There are two separate benefits to having a larger `max_seqio_extent_size` value:

- Initial allocation of the file is faster, since VxFS can allocate the file in larger chunks, which is more efficient.
- Later application access to the file is also faster, since accessing less fragmented files is also more efficient.

In the 5.1 release, the default value was changed back to its earlier setting because the larger 5.0 MP3 value can lead to applications experiencing "no space left on device" (ENOSPC) errors if the file system is close to being full and all remaining space is preallocated to files. VxFS attempts to reclaim any unused preallocated space if the space is needed to satisfy other allocation requests, but the current implementation can fail to reclaim such space in some situations.

If your workload has lower performance with the VxFS 5.1 release and you believe that the above change could be the reason, you can use the `vxtunefs` command to increase this tunable to see if performance improves.

To restore the benefits of the higher tunable value

1. Increase the tunable back to the 5.0 MP3 value, which is 1 GB divided by the file system block size.

   Increasing this tunable also increases the chance that an application may get a spurious ENOSPC error as described above, so change this tunable only for file systems that have plenty of free space.
2. Shut down any application that are accessing any large files that were created using the smaller tunable setting.
3. Copy those large files to new files, which will be allocated using the higher tunable setting.
4. Rename the new files back to the original names.
5. Restart any applications were shut down earlier.

umount can hang when inotify watches are used (1590324)

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

Workaround: There is no workaround for this issue.

NFS issues with VxFS checkpoint (1974020)

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

```
Stale NFS file handle
```

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

There is no workaround at this time.

Veritas Cluster Server known issues

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

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

Parent service groups fail to restart after a child service group that has recovered from a fault restarts (2330038)

A child service group that has recovered from a fault will not be able to restart its faulted parent service group, if the parent service group's fault is detected before the child service group's fault.

Workaround:

Set the child service group's `OnlineClearParent` attribute to 1. When the child service group recovers from a fault and comes online, VCS clears the fault of the parent service group. This allows the VCS to bring the parent service group online.

hacmd -display G or A or O or E or S or C dumps core (2415578)

VCS ha commands do not work when object names are single character long.

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.

Application Agent does not handle a case when user is root, envfile is set and shell is csh. (2490299)

The Application Agent uses the `system` command to execute the Start/Stop/Monitor/Clean Programs for root user. This executes Start/Stop/Monitor/Clean Programs in sh shell, due to which there is an error when root user has csh shell and `EnvFile` is written accordingly.

Workaround:

Do not set csh as shell for root user. Use sh as shell for root instead.

Forcefully un-configuring AMF does not change the monitor method of agent to TRADITIONAL (2564376)

If you forcefully unconfigure AMF when a script based agent is registered with AMF, the monitor method of the agent is not changed to `TRADITIONAL`. It remains `IMF`.

Workaround: Restarting the agent will resolve the issue.

Forcefully un-configuring AMF causes the engine log to be flooded with error messages (2535690)

If you forcefully unconfigure AMF when a script based agent is registered with AMF, the `getnotification` thread continuously polls and displays error messages in the engine log.

Workaround: Restarting the agent will resolve the issue.

NFS resource goes offline on its own and errors out when restarted (2490415)

If multiple agent processes are running because an agent process is restarted multiple times by _had, then only one of the agent process survives and other agent processes go offline on its own. Even though the agent process is running, _had does not recognize it and hence does not perform any resource operations.

Workaround: Kill the agent process to recover from this situation. Refer to the engine log for further actions (if required) to restart the agent.

HAD dumps core when hagrp -clear is executed on a group in OFFLINE|FAULTED state and a resource in the fault path is waiting to go online (2536404)

This issue occurs if you have a resource dependency, such as r1 -> r2 -> r3. While resources r2 and r3 are online and you initiate bringing resource r1 online, before the OnlineTimeout occurs, resources r2 and r3 suffer a fault. Resource r2 faults first, and then r3 faults. After the fault of both resources is detected, the group is becomes in an OFFLINE|FAULTED state and resource r1 is stuck waiting to become online. If you execute the `hagrp -clear` command to clear the fault, then HAD dumps core on all nodes due to assertion failure.

Workaround: Flush the pending online operation using the `hagrp -clear` command before clearing the fault.

The hares -display command fails if the resource is part of a global service group (2358600)

The `hares -display` command incorrectly processes the response received from the had process. Due to the processing error, `hares -display` does not show the resource details.

Workaround: Use the -localclus or -clus option with `hares -display`.

Excessive delay messages in one-node configurations of Storage Foundation High Availability (SFHA) (2486415)

The GAB error, "`Excessive delay between successive calls to GAB heartbeat`" appear in the engine log while running a single node or standalone VCS cluster where GAB is disabled.

GAB heartbeat log messages are logged as informational when there is delay between heartbeats (HAD being stuck). When HAD runs in -onenode, GAB does not need to be enabled. When HAD is running in -onenode, for self-checking purposes, HAD simulates heartbeat with an internal component of HAD itself. These log messages are logged because of delay in simulated heartbeats.

Workaround: Log messages are for informational purpose only. When HAD is running in -onenode, no action is needed on excessive delay between heartbeats.

Agent does not retry resource registration with IMF to the value specified in RegisterRetryLimit key of IMF attributes (2411882)

Agent maintains internal count for each resource to track the number of times a resource registration is attempted with IMF. This count gets incremented if any

attempts to register a resource with IMF fails. At present, any failure in un-registration of resource with IMF, also increments this internal count. This means that if resource un-registration fails, next time agent may not try as many resource registrations with IMF as specified in RegisterRetryLimit. This issue is also observed for multiple resources.

Workaround: Increase the value of RegisterRetryLimit if un-registration and registration of one or more resources is failing frequently. After increasing the RegisterRetryLimit, if resource registration continues to fail, report this problem to Symantec.

DBED commands such as sfua_db_config and dbed_ckptcreate fail

DBED commands such as `sfua_db_config` and `dbed_ckptcreate` fail with the following error message:

```
VXDBA_PRODUCT exec_remote ERROR V-81-7700 Can not connect to the
vxdbd. It might be down. Check the status and restart it if it is
not up.
```

However, the output of the `ps` command shows that vxdbd is running:

```
# ps -ef | grep vxdbd
root  14572   1   0 05:07:35      -  0:13 /opt/VRTSdbed/common/bin/vxdbd -d
```

Workaround: Set the environment variable EAT_HOME_DIR to the value /opt/VRTSdbed/eat before running the command:

```
# EAT_HOME_DIR=/opt/VRTSdbed/eat

# export EAT_HOME_DIR
```

installation of Storage Foundation 5.1SP1RP2 fails on SLES10 SP4

Installation of Storage Foundation 5.1SP1RP2 fails on SLES10 SP4 with the error:

```
The following required OS rpms were not found on host:

compat-libstdc++-5.0.7-22.2.x86_64
```

Workaround: Install the compat-libstdc++-32bit RPM package. It is available on the SLES10 SP3 installation media.

Oracle agent incorrectly reports the global resource as online when the resource inside the local zone is online and the Sid's are same (2561563)

Oracle agent incorrectly reports the resource configured for Oracle instance running in global container as online, if the resource configured for Oracle instance running in local container also has same value for Sid attribute and the instance in local container is online.

The above issue is also applicable for ASMInst and Netlsnr agents.

For Netlsnr agent the above issue appears when the Home and listener attributes of the resources running in global and local container are same.

The issue does not appear for Oracle and ASMInst agents when multiple local containers have resources configured with the same value of Sid attribute.

The issue does not appear for Netlsnr agent when multiple local containers have resources configured with the same value of Home and listener attributes.

NFS mount does not work whenever the NFS daemon is started on port other than 2049 (2477799)

On RHEL 6.0 and 6.1, if the paths are exported by a NFS service (nfsd) which is running on any port other than the default (2049), the mount command fails to mount the exported paths. If NFS service is running on the default port (2049), then you are able to mount the exported paths successfully.

This issue occurs only on RHEL 6.0 and 6.1. It does not exist on RHEL 5.6, 5.7 and SLES 10, 11.

Veritas Volume Replicator known issues

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

vradmin syncvol command compatibility with IPv6 addresses (2075307)

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

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

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

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

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname could not be
imported on bunker host hostname. Operation failed with error 256
and message VxVM VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote
server unreachable... Timestamp VCS ERROR V-16-2-13066 (hostname)
Agent is calling clean for resource(RVGPrimary) because the resource
is not up even after online completed.
```

Workaround:

To resolve this issue

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

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

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

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

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

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

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

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

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround: The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

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

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

- the host name, the DNS server name and domain name are specified to the YaST tool.
- IPv6 address is assigned by the Dynamic Host Configuration Protocol (DHCP).
- the "Write Hostname to /etc/hosts" option is selected (this is selected by default).

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

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

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

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

Workaround: The following procedure resolves this issue.

To resolve this issue

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

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

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

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

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

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

environment, vradmin createpri may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

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

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

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

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

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

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

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

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

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

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

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

Workaround:

To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

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

Workaround: There is no workaround for this issue.

vxassist relayout removes the DCM (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

1 Pause or stop the applications.

2 Wait for the RLINKs to be up to date. Enter the following:

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

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

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

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

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

Workaround:

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

Database fails over during Flashsnap operations (1469310)

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

Workaround:

There is no workaround for this issue.

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

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

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

Workaround

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

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

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

```
# vxdg join snapshop_disk_group_name
          primary_disk_group_name
```

2 Start all the volumes in primary disk group.

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

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

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

Repeat this step for all the volumes.

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

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

Workaround:

There is no workaround for this issue.

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

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

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

```
dbed_vmclonedb started at 2009-08-26 11:32:16
Editing remote_login_passwordfile in initclone2.ora.
Altering instance_name parameter in initclone2.ora.
```

```
Altering instance_number parameter in initclone2.ora.
Altering thread parameter in initclone2.ora.
SFORA dbed_vmclonedb ERROR V-81-4918 Database clone2 has not been
correctly recovered.
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.20569/recover.log.
```

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

Workaround

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

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

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

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

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

4. Start the database.

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 RP2 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.

Failure to set the MTU (Maximum Transmission Unit) size in LLT over UDP environments causes issues with the PrivNIC/MultiPrivNIC agents (2557144)

If the MTU size field is not set explicitly when you configure the PrivNIC/MultiPrivNIC agents in an LLT over UDP environment, the agents may

fail to plumb the private IP addresses during their operations or may configure incorrect MTU size on the LLT interfaces.

The agents use the `lltstat -l` command to retrieve MTU size information for LLT interfaces. In an LLT over UDP environment, the command retrieves 8192 as the MTU size. When the PrivNIC/MultiPrivNIC agents use this size information to plumb the IP addresses, the operation may fail causing the agents to fault. However, even if the plumbing operation succeeds, the incorrect MTU configuration may still cause issues in the cluster later.

Workaround:

To update the PrivNIC/MultiPrivNIC resource configuration in an LLT over UDP environment

1. Retrieve the MTU size of the network interfaces configured under PrivNIC/MultiPrivNIC agents:

   For AIX: `# lsattr -El en1`

   For HPUX: `# lanadmin -m 1`

   For Linux: `# ifconfig eth1`

   For Solaris: `# ifconfig ce0`

2. Set the MTU attribute for the PrivNIC/MultiPrivNIC resource:

   ```
   # haconf -makerw
   ```

   Run the following command for all the network interfaces configured under PrivNIC/MultiPrivNIC agents:

   ```
   # hares -modify resource_name MTU -add interface_name mtu_size
   ```

   Where:

   resource_name is the name of the PrivNIC/MultiPrivNIC resource

   interface_name is the name of the network interface for which the MTU size is set

   mtu_size is the MTU size retrieved in step 1.

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

The MultiPrivNIC agent configures IP addresses on virtual interfaces though the UseVirtualIP parameter is set 0 (2566142)

The `MultiPrivNIC` agent configures the IP addresses on virtual interfaces though the `UseVirtualIP` parameter is set to 0 while configuring the `MultiPrivNIC` resource in the `main.cf` file. However, no functional issues are observed in the cluster.

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 RP2 release.

Veritas Volume Manager software limitations

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

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

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

Workaround:

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

DMP settings for NetApp storage attached environment

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

Table 1-13

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

The change is persistent across reboots.

To change the tunable parameters

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

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

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

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

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

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

Veritas File System software limitations

The following are software limitations in the 5.1 SP1 RP2 release of Veritas Storage Foundation.

There are no Veritas Storage Foundation software limitations in the 5.1 SP1 RP2 release.

Linux I/O Scheduler for Database Workloads (1446361)

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

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

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

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

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

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

To:

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

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

Reboot the system once the appropriate file has been modified.

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

Veritas Volume Replicator software limitations

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

Replication in a shared environment

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

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

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

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 5.1 SP1 RP2 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 Storage Foundation for Databases tools software limitations

The following are software limitations in this release of Veritas Storage Foundation for Databases.

Oracle Data Guard in an Oracle RAC environment

Database snapshots and Database Checkpoints are not supported in an environment where there are both Data Guard and Oracle RAC. But separately, either Data Guard or Oracle RAC supports Database snapshots and Database Checkpoints.

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 5.1 SP1 RP2: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 5.1 SP1 RP2.

Documentation errata

The following section provides documentation updates.

Veritas Cluster Server Administrator's Guide (2444653)

In the "VCS environment variables" section, the definition of the variable `VCS_GAB_RMTIMEOUT` should be "Timeout in milliseconds for HAD to register with GAB."

The value of `VCS_GAB_RMTIMEOUT` is specified in milliseconds. The minimum value is 200000 milliseconds or 200 seconds. If `VCS_GAB_RMTIMEOUT` is less than the minimum value then VCS overrides and sets it to 200000 milliseconds.

Veritas Installation Guides (2521411)

The `ignorepatchreqs` option is no longer valid and does not work.

List of RPMs

This section lists the RPMs for 5.1 SP1 RP2.

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

Table 1-14 RPMs for Red Hat Enterprise Linux 5

Name	Version	Arch	Size in bytes
VRTSamf	5.1.132.000	x86_64	767643
VRTScavf	5.1.132.000	i386	177095
VRTScps	5.1.132.000	i686	15383964
VRTSdbac	5.1.132.000	x86_64	928008
VRTSdbed	5.1.132.000	i686	6167029
VRTSfssdk	5.1.132.000	x86_64	350561
VRTSgab	5.1.132.000	x86_64	939315
VRTSglm	5.1.132.000	x86_64	106311
VRTSllt	5.1.132.000	x86_64	906720
VRTSlvmconv	5.1.132.000	i686	70558
VRTSob	3.4.312	i686	33422315

Table 1-14 RPMs for Red Hat Enterprise Linux 5 *(continued)*

Name	Version	Arch	Size in bytes
VRTSodm	5.1.101.000	x86_64	236749
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcs	5.1.132.000	i686	46725855
VRTSvcsag	5.1.132.000	i686	684656
VRTSvcsdr	5.1.132.000	x86_64	198949
VRTSvcsea	5.1.132.000	i686	3814205
VRTSvxfen	5.1.132.000	x86_64	450643
VRTSvxfs	5.1.132.000	x86_64	9357549
VRTSvxvm	5.1.132.000	x86_64	27138970

Table 1-15 RPMs for Red Hat Enterprise Linux 6

Name	Version	Arch	Size in bytes
VRTSamf	5.1.132.000	x86_64	536196
VRTScavf	5.1.132.000	i386	161560
VRTScps	5.1.132.000	x86_64	12056540
VRTSdbac	5.1.132.000	x86_64	791136
VRTSfssdk	5.1.132.000	x86_64	290924
VRTSgab	5.1.132.000	x86_64	777720
VRTSglm	5.1.132.000	x86_64	111700
VRTSllt	5.1.132.000	x86_64	740312
VRTSlvmconv	5.1.132.000	i686	64212
VRTSob	3.4.312	i686	33422315
VRTSodm	5.1.132.000	x86_64	252344
VRTSsfmh	3.1.830.0	i686	27303764
VRTSvcs	5.1.132.000	i686	30664348
VRTSvcsag	5.1.132.000	i686	529760

Table 1-15 RPMs for Red Hat Enterprise Linux 6 *(continued)*

Name	Version	Arch	Size in bytes
VRTSvcsdr	5.1.132.000	x86_64	225808
VRTSvcsea	5.1.132.000	i686	3392124
VRTSvxfen	5.1.132.000	x86_64	904128
VRTSvxfs	5.1.132.000	x86_64	5760456
VRTSvxvm	5.1.132.000	x86_64	17036144

Table 1-16 RPMs for SUSE Linux Enterprise Server 10

Name	Version	Arch	Size in bytes
VRTSamf	5.1.132.000	x86_64	4152755
VRTScavf	5.1.132.000	i386	151296
VRTScps	5.1.132.000	i686	14430937
VRTSdbac	5.1.132.000	x86_64	1728211
VRTSdbed	5.1.132.000	i586	5515449
VRTSfssdk	5.1.132.000	x86_64	307367
VRTSgab	5.1.132.000	x86_64	2062011
VRTSglm	5.1.132.000	x86_64	177403
VRTSllt	5.1.132.000	x86_64	1554395
VRTSlvmconv	5.1.132.000	i586	61695
VRTSob	3.4.312	i686	33422315
VRTSodm	5.1.101.000	x86_64	358835
VRTSsfmh	3.1.830	i686	27303764
VRTSvcs	5.1.132.000	i586	42271927
VRTSvcsag	5.1.132.000	i586	548232
VRTSvcsdr	5.1.132.000	x86_64	404562
VRTSvcsea	5.1.132.000	i586	3597404
VRTSvxfen	5.1.132.000	x86_64	2198415

Table 1-16 RPMs for SUSE Linux Enterprise Server 10 *(continued)*

Name	Version	Arch	Size in bytes
VRTSvxfs	5.1.132.000	x86_64	11331284
VRTSvxvm	5.1.132.000	x86_64	25711499

Table 1-17 RPMs for SUSE Linux Enterprise Server 11

Name	Version	Arch	Size in bytes
VRTSamf	5.1.132.000	x86_64	1155799
VRTScavf	5.1.132.000	i386	154750
VRTScps	5.1.132.000	i686	11569670
VRTSdbac	5.1.132.000	x86_64	1150418
VRTSdbed	5.1.132.000	i586	5515449
VRTSfssdk	5.1.132.000	x86_64	268072
VRTSgab	5.1.132.000	x86_64	1297527
VRTSglm	5.1.132.000	x86_64	166520
VRTSllt	5.1.132.000	x86_64	1155493
VRTSlvmconv	5.1.132.000	i586	62772
VRTSob	3.4.312	i686	33422315
VRTSodm	5.1.132.000	x86_64	275632
VRTSsfmh	3.1.830	i686	27303764
VRTSvcs	5.1.132.000	i686	30408990
VRTSvcsag	5.1.132.000	i686	495338
VRTSvcsdr	5.1.132.000	x86_64	330201
VRTSvcsea	5.1.132.000	i686	2630250
VRTSvxfen	5.1.132.000	x86_64	1398658
VRTSvxfs	5.1.132.000	x86_64	6166483
VRTSvxvm	5.1.132.000	x86_64	18203807

Downloading the 5.1 SP1 RP2 archive

The patches that are included in the 5.1 SP1 RP2 release are available for download from the Symantec website. After downloading the 5.1 SP1 RP2 rolling patch, use gunzip and tar commands to uncompress and extract it.

For the 5.1 SP1 RP2 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

http://www.symantec.com/docs/TECH75506

Chapter 2

Installing the products for the first time

This chapter includes the following topics:

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

Installing the Veritas software using the script-based installer

This section describes how to install a 5.1 SP1 RP2 Veritas Storage Foundation and High Availability Solutions product for the first time on a host. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 5.1 SP1 *Installation Guide* and *Release Notes* for your product for more information.

See "Upgrading to 5.1 SP1 RP2" on page 98.

To install the Veritas software for the first time

1. Download Storage Foundation and High Availability Solutions 5.1 SP1 from http://fileConnect.symantec.com.
2. Extract the tar ball into a directory called /tmp/sfha51sp1.
3. Check http://sort.symantec.com/patches to see if there are any patches available for the 5.1SP1 Installer. Download applicable P-patches and extract them to the /tmp directory.

4 Change the directory to /tmp/sfha51sp1:

```
# cd  /tmp/sfha51sp1
```

5 Run the installer to install SFHA 5.1SP1. See the Installation Guide for instructions on installing the 5.1 SP1 version of this product.

```
# ./installer -require complete_path_to_SP1_installer_patch
```

6 Download SFHA 5.1 SP1 RP2 from http://sort.symantec.com/patches.

7 Extract it to a directory called /tmp/sfha51sp1rp2.

8 Check http://sort.symantec.com/patches to see if there are patches available for the 5.1SP1RP2 installer. Download applicable P-patches and extract them to the /tmp directory.

9 Change the directory to /tmp/sfha51sp1rp2:

```
# cd  /tmp/ sfha51sp1rp2
```

10 Invoke the `installrp` script to install 5.1SP1 RP2:

```
# installrp -require complete_path_to_SP1RP2_installer_patch
```

11 If you did not configure the product after the 5.1 SP1 installation, the installer prompts you to configure the product during RP installation. If you do not want to configure the product now, answer **n** when prompted. To configure the product in the future, run the product installation script from the 5.1 SP1 installation media or from /opt/VRTS/install directory 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 RP2 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.

See “Upgrading to 5.1 SP1 RP2” on page 98.

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 5.1 SP1 RP2 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 RP2.

2. On the Select a task and product page, select the installation from the **Task** drop-down list, and click **Next**.
3. On the **Select a task and product** page, select **Install SP1 RP2** from the **Task** drop-down list, and click **Next**
4. Choose minimal, recommended, or all packages. Click **Next**.
5. Stop all applications accessing the filesystem. Unmount all mounted file systems before installation.
6. Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
7. After the validation completes successfully, click **Next** to install 5.1 SP1 RP2 patches on the selected system.
8. After the installation completes, you must choose your licensing method.

   On the license page, select one of the following tabs:

   - Keyless licensing

     Note: The keyless license option enables you to install without entering a key. However, in order to ensure compliance you must manage the systems with a management server.

     For more information, go to the following website:

     http://go.symantec.com/sfhakeyless

     Complete the following information:
     Choose whether you want to enable Veritas Volume Replicator.
     Choose whether you want to enable Global Cluster option.
     Click Register.

   - Enter license key
     If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

9 The installer prompts you to configure the cluster.

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

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

10 Select the checkbox to specify whether you want to send your installation information to Symantec.

```
Would you like to send the information about this installation
to Symantec to help improve installation in the future?
```

Click **Finish**.

Upgrading Veritas product with the Veritas Web-based installer

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

To upgrade Veritas product

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

2 Start the Web-based installer.

3 Select **Upgrade**.

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

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

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

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

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

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

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

Chapter 3

Upgrading to 5.1 SP1 RP2

This chapter includes the following topics:

- Prerequisites for upgrading to 5.1 SP1 RP2
- Downloading required software to upgrade to 5.1 SP1 RP2
- Supported upgrade paths
- Upgrading to 5.1 SP1 RP2
- Verifying software versions

Prerequisites for upgrading to 5.1 SP1 RP2

The following list describes prerequisites for upgrading to the 5.1 SP1 RP2 release:

- For any product in the Veritas Storage Foundation stack, you must have the 5.1 SP1 (or later) installed before you can upgrade that product to the 5.1 SP1 RP2 release.
- Each system must have sufficient free space to accommodate patches.
- The full list of prerequisites can be obtained by running `./installrp -precheck`
- Make sure to download the latest patches for the installer.
  See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

Downloading required software to upgrade to 5.1 SP1 RP2

This section describes how to download the latest patches for the installer.

To download required software to upgrade to 5.1 SP1 RP2

1. Download SFHA 5.1 SP1 RP2 from http://sort.symantec.com/patches.
2. Extract it to a directory, say /tmp/sfha51sp1rp2.
3. Check http://sort.symantec.com/patches to see if there are patches available for the 5.1 SP1 RP2 installer. Download applicable P-patches and extract them to the /tmp directory.
4. When you run the `installrp` script, use the `-require` option and specify the location where you downloaded the patches.

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 5.1 SP1 to 5.1 SP1 RP2
- 5.1 SP1 RP1 to 5.1 SP1 RP2
- 5.1SP1 PR2 to 5.1 SP1 RP2 (Applies to RHEL 6.0 only)
- 5.1SP1PR2P1 to 5.1 SP1 RP2(Applies to RHEL6.1 only)
- 5.1SP1 PR3 to 5.1 SP1 RP2

Upgrading to 5.1 SP1 RP2

This section describes how to upgrade from 5.1 SP1 (or later) to 5.1 SP1 RP2 on a cluster or a standalone system.

- Performing a full upgrade to 5.1 SP1 RP2 on a cluster
  Use the procedures to perform a full upgrade to 5.1 SP1 RP2 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System (SFCFS), Veritas Storage Foundation for Oracle RAC (SFRAC), Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC), or Symantec VirtualStore (SVS) installed and configured.
- Upgrading to 5.1 SP1 RP2 on a standalone system
  Use the procedure to upgrade to 5.1 SP1 RP2 on a system that has SF installed.
- Performing a rolling upgrade using the installer
  Use the procedure to upgrade your Veritas product with a rolling upgrade.

See “Installing the Veritas software using the script-based installer” on page 91.

Performing a full upgrade to 5.1 SP1 RP2 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.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.1 SP1 RP2:

- Performing a full upgrade to 5.1 SP1 RP2 on a Veritas Cluster Server
- Performing a full upgrade to 5.1 SP1 RP2 on an SFHA cluster
- Performing a full upgrade to 5.1 SP1 RP2 on an SFCFS cluster
- Performing a full upgrade to 5.1 SP1 RP2 on an SFCFS RAC cluster
- Performing a full upgrade to 5.1 SP1 RP2 on an SF Oracle RAC cluster

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

Performing a full upgrade to 5.1 SP1 RP2 on a 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 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

2 Log in as superuser.

Note: Upgrade the Operating System and reboot the systems if required.

---

3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

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

4 Resolve any issues that the precheck finds.

5 Start the upgrade:

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

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

Performing a full upgrade to 5.1 SP1 RP2 on an SFHA cluster

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

To perform a full upgrade to 5.1 SP1 RP2 on an SFHA cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

2 Log in as superuser.

3 From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

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

Performing a full upgrade to 5.1 SP1 RP2 on an SFCFS cluster

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

To perform a full upgrade to 5.1 SP1 RP2 on an SFCFS cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

2 Log in as superuser.

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

```
# mount | grep vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# mount | grep vxfs
```

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

```
# umount /filesystem
```

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

---

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

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

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

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

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

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

---

6 Stop activity to all VxVM volumes.

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

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

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

Verify that no volumes remain open:

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

8 If required, apply the OS kernel patches.

9 From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

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

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

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

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

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

13 If you stopped any RVGs in step 5, restart each RVG:

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

14 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

15 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

Performing a full upgrade to 5.1 SP1 RP2 on an SFCFS RAC cluster

To prepare for a full upgrade to 5.1 SP1 RP2 on an SFCFS RAC cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

2 Log in as superuser.

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

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

5 Stop Oracle Clusterware.

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

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

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

7 Unmount all file systems:

```
# umount /filesystem
```

8 Check if the VEA service is running:

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

9 If the VEA service is running, stop it:

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

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

11 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.

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

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

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

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

14 Disable the startup scripts before upgrading the operating system.

```
# chkconfig vcs off
# chkconfig vxodm off
# chkconfig vxfen off
# chkconfig vxgms off
# chkconfig vxglm off
# chkconfig gab off
# chkconfig llt off
```

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

To upgrade Storage Foundation Cluster File System for Oracle RAC

1 From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

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

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

4 Shut down and reboot the systems.

5 Upgrade Oracle RAC, if required.

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

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

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

  - Back up libodm10.so file.

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

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

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

- For Oracle 11g:

- Change to the $ORACLE_HOME/lib directory:

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

- Back up libodm11.so file.

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

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

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

To bring the upgraded cluster online and restore components

1. If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.
2. If necessary, reinstate any missing mount points in the /etc/fstab file on each node.
3. If any VCS configuration files need to be restored, stop the cluster, restore the files to the /etc/VRTSvcs/conf/config directory, and restart the cluster.
4. Restart all the volumes by entering the following command for each disk group:

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

5. Remount all VxFS file systems on all nodes:

   ```
   # mount /filesystem
   ```

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

Performing a full upgrade to 5.1 SP1 RP2 on an SF Oracle RAC cluster

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

To upgrade to 5.1 SP1 RP2 on an SF Oracle RAC cluster

1. Make sure you have downloaded the latest software required for the upgrade.

   See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

2. Log in as superuser.

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

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

```
# haconf -makerw
```

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

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

6 Make the configuration read-only:

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

7 If Oracle Clusterware is not controlled by VCS, enter the following command on each node of the cluster to stop Oracle Clusterware:

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

8 Stop VCS.

```
# hastop -all
```

9 If required, apply the OS kernel patches.

See *Oracle's* documentation for the procedures.

10 From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. If ssh key authentication is configured then enter:

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

If ssh is not configured then enter:

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

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

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

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

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

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

```
# haconf -makerw
```

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

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

15 Make the configuration read-only:

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

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

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

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

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

18 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

19 If Oracle Clusterware is not controlled by VCS, enter the following command on each node to start Oracle Clusterware.

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

20 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 RP2 on a standalone system

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

To upgrade to 5.1 SP1 RP2 on a standalone system

1. Make sure you have downloaded the latest software required for the upgrade.

   See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

2. Log in as superuser.
3. Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
4. If required, apply the OS kernel patches.
5. Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

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

6. Unmount all Storage Checkpoints and file systems:

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

7. 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.

8. 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.

9 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
```

10 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
```

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

```
# ./installrp nodename
```

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

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

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

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

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

15 Remount all VxFS file systems and Storage Checkpoints:

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

16 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 RP2 on a system that has encapsulated boot disk

You can use this procedure to upgrade to 5.1 SP1 RP2 on a system that has encapsulated boot disk.

Note: Upgrading with encapsulated boot disk from 5.1 SP1 to 5.1 SP1 RP2 requires multiple reboots.

To upgrade to 5.1 SP1 RP2 on a system that has encapsulated boot disk

1. Manually unmount file systems and stop open volumes.
2. If required, manually break the mirror.
3. Upgrade to 5.1 SP1 RP2 using `installrp` command.
4. After upgrading, reboot the system to have the new VM drivers take effect.
5. If the mirrors of boot disk are split manually in step2, re-join the mirrors manually when systems reboot after upgrading to 5.1 SP1 RP2.

Performing a rolling upgrade using the installer

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

- About rolling upgrades
- Prerequisites for a rolling upgrades

About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime.

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

You can perform a rolling upgrade for the following products:

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

- Storage Foundation for Oracle RAC

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

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.
- For SFCFS for Oracle RAC, stop Oracle Clusterware before upgrading Kernel packages on any node.
- Make sure you have downloaded the latest software required for the upgrade. See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 97.

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.

Upgrading the operating system

This section describes how to upgrade the operating system on a Storage Foundation node where you plan to upgrade to 5.1 SP1 RP2. This section includes the following topics:

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

Upgrading RHEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 SP1 RP2.

To upgrade to a later version of RHEL 5

1 Stop Storage Foundation.

2 Upgrade to the latest update version of RHEL 5.

3 Upgrade to 5.1 SP1 RP2.

4 Start Storage Foundation.

Upgrading OEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 SP1 RP2.

To upgrade to a later version of OEL 5

1 Stop Storage Foundation.

2 Upgrade to the latest update version of OEL 5.

3 Upgrade to 5.1 SP1 RP2.

4 Start Storage Foundation.

Upgrading SLES 10

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 5.1 SP1 RP2.

To upgrade to a later version of SLES 10

1 Stop Storage Foundation.

2 Upgrade to the latest update version of SLES 10.

3 Upgrade to 5.1 SP1 RP2.

4 Start Storage Foundation.

Verifying software versions

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

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

Chapter 4

Removing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- About removing Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2
- Uninstalling Veritas Storage Foundation Cluster File System 5.1 SP1 RP2
- Uninstalling Veritas Storage Foundation for Oracle RAC
- Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP2

About removing Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2

Symantec recommends that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the previous release.

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

Uninstalling Veritas Storage Foundation Cluster File System 5.1 SP1 RP2

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

Preparing to uninstall Veritas Storage Foundation Cluster File System

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

To prepare to uninstall Veritas Storage Foundation Cluster File System

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

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

3. Back up the following configuration files:

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

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

     ```
     # df -v /
     ```

     The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:
   - Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes mirrootvol-01, and mirswapvol-01 that are configured on a disk other than the root disk:

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

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

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

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

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

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

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

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

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

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

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

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

7 Unmount all Storage Checkpoints and file systems:

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

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

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

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

To verify that no volumes are open:

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

9 Stop VCS:

```
# hastop -all
```

Uninstalling Veritas Storage Foundation Cluster File System

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

To uninstall Veritas Storage Foundation Cluster File System

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

2 Navigate to the directory that contains the uninstallation program:

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

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

```
# ./uninstallsfcfs node1 node2
```

4 Confirm the uninstallation:

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

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

5 Reboot the nodes:

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

After uninstalling the Veritas Storage Foundation Cluster File System, refer to the *Veritas Storage Foundation Cluster File System 5.1 SP1 Installation Guide* to reinstall the 5.1 SP1 software.

Uninstalling Veritas Storage Foundation for Oracle RAC

The following procedure uninstalls Veritas Storage Foundation for Oracle RAC (SFRAC).

Note: This procedure will remove the complete SFRAC stack from all nodes.

To uninstall Veritas Storage Foundation for Oracle RAC

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

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

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

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

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

- For 10gR2 or 11gR1:

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

- For 11gR2:

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

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

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

```
# fuser -c mount_point
```

4 Unmount CFS file systems that are not under VCS control

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

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

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

  # **umount *mount_point***

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

  On any one node execute following command to stop VCS:

  # **hastop -all**

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

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

  # **fuser -c *mount_point***

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

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

  # **mount -v | grep vxfs**

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

  # **umount *mount_point***

8 Remove SF for Oracle RAC.

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

  # **cd /opt/VRTS/install**

- Start the uninstallsfrac program:

  # **./uninstallsfrac**

9 After uninstalling the SFRAC, refer to the *Veritas Storage Foundation for Oracle RAC 5.1 SP1 Installation and Configuration Guide* document to reinstall the SFRAC 5.1 SP1 software.

Uninstalling Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 RP2

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

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

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

3. Back up the following configuration files:

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

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

   For example, to stop Oracle Clusterware:

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

5. Stop VCS:

   ```
   # hastop -all
   ```

6. Verify that port h is not open:

   ```
   # gabconfig -a
   ```

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

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

8 Unmount all file systems:

```
# umount /filesystem
```

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

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

To verify that no volumes are open:

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

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

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

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

2 Navigate to the directory that contains the uninstallation program:

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

3 Start the uninstallation program:

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

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

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

The installer checks the RPMs installed on the system.

5 Confirm the uninstallation at the following prompt:

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

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

6 Reboot the nodes:

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

After uninstalling the Veritas Storage Foundation Cluster File System for Oracle RAC, refer to the *Veritas Storage Foundation Cluster File System for Oracle RAC 5.1 SP1 Installation and Configuration Guide* to reinstall the 5.1 SP1 software.