Veritas Storage Foundation™ and High Availability Solutions Release Notes

Solaris

5.1 Service Pack 1 Rolling Patch 4

Symantec™

Veritas Storage Foundation and High Availability Solutions Release Notes 5.1 Service Pack 1 Rolling Patch 4

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

Product version: 5.1 SP1 RP4

Document version: 5.1 SP1 RP4 Rev0

Legal Notice

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 website 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 website.

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

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/
forums/storage-and-clustering-documentation

About Symantec Connect

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

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

Contents

Chapter 2   Installing the products for the first time 164

Chapter 3   Installing the products using JumpStart 168

Chapter 4   Upgrading to 5.1 SP1 RP4 176

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- Introduction

- About the installrp and the uninstallrp scripts

- Overview of the installation and upgrade process

- Changes introduced in 5.1 SP1 RP4

- System requirements

- List of products

- Fixed Issues

- Known Issues

- Software limitations

- Documentation errata

- List of patches

- Downloading the 5.1 SP1 RP4 archive

Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Rolling Patch 4 release.

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

http://www.symantec.com/docs/TECH75362

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 RP2

- Storage Foundation and High Availability Solutions 5.1 SP1 RP3

- VirtualStore 5.1 SP1 PR3

This rolling patch is available as 5.1 SP1 RP4.

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 PR3 RP1

About the installrp and the uninstallrp scripts

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP4 provides an upgrade script.

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.

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

Command Line Option	Function
[-jumpstart *jumpstart_path]*	Use the -jumpstart option to generate finish scripts. You can use the finish scripts with the Solaris JumpStart Server to automate installation of all packages and patches for every product. You need to specify an available location to store the finish scripts as a complete path. The -jumpstart option is supported on Solaris only.
[-keyfile *ssh_key_file*]	Use the -keyfile option to specify a key file for SSH. When you use this option the -i *ssh_key_file* is passed to every SSH invocation.
[-patchpath *patch_path*]	Use the -patchpath option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by installrp.
[-rootpath *root_path*]	Use the -rootpath option to re-root the installation of all packages to the given path. On Solaris, -rootpath passes -R <root_path> to pkgadd.

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

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.

The uninstallrp script options

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP4 provides a new uninstallation script.

See Table 1-2 for release versions and products that support rolling back.

Symantec recommends that you use the new uninstallation script. The `uninstallrp` script uninstalls all the patches associated with packages installed, and starts the processes. Do not use the `uninstallrp` script for rolling back, because it removes the entire stack.

Table 1-2   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.
`[ -logpath log_path ]`	Use the `-logpath` option to select a directory other than `/opt/VRTS/install/logs` as the location where the uninstallrp 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.

Table 1-2   The command line options for the product upgrade script *(continued)*

Command Line Option	Function
[-tmppath *tmp_path*]	Use the -tmppath option to select a directory other than /var/tmp as the working directory for uninstallrp. This destination is where initial logging is performed and where packages 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.
[-rootpath *root_path*]	Use the -rootpath option to re-root the installation of all packages to the given path. On Solaris, -rootpath passes -R <root_path> to pkgadd or pkgrm.

Table 1-2   The command line options for the product upgrade script *(continued)*

Command Line Option	Function
`[-rsh	-redirect	-makeresponsefile	-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 `-makeresponsefile` option to generate a response file without doing an actual installation. Text displaying installation, uninstallation, start and stop operations are simulations. These actions are not being performed on the system.
	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 packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable.

Overview of the installation and upgrade process

Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

To install the Veritas software for the first time

1   Skip this step if you are upgrading to 5.1 SP1 RP4. If you are installing 5.1 SP1 RP4 on a fresh system or cluster:

  ■ Download Storage Foundation and High Availability Solutions 5.1 SP1 from http://fileConnect.symantec.com.

  ■ Extract the tar ball into a directory called `/tmp/sfha51sp1`.

- Check https://sort.symantec.com/patches to see if there are any patches available for the 5.1 SP1 Installer. Download applicable P-patches and extract them to the /tmp directory.

- Change the directory to /tmp/sfha51sp1:

  # **cd /tmp/sfha51sp1**

- Install the 5.1 SP1 software. Follow the instructions in the Installation Guide.

  # **./installer -require *complete_path_to_SP1_installer_patch***

2   Download SFHA 5.1 SP1 RP4 from https://sort.symantec.com/patches and extract it to a directory called /tmp/sfha51sp1rp4.

3   Check https://sort.symantec.com/patches to see if there are patches available for the 5.1 SP1 RP4 installer. Download applicable P-patches and extract them to the /tmp directory.

4   Change the directory to /tmp/sfha51sp1rp4:

  # **cd /tmp/sfha51sp1rp4**

5   Install 5.1 SP1 RP4:

  # **./installrp -require *complete_path_to_SP1RP4_installer_patch***

Changes introduced in 5.1 SP1 RP4

This section lists the changes in 5.1 SP1 RP4.

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 RP4.

Enhancement on VRTSaslapm upgrade and rolling back

During the upgrade, if the version of VRTSaslapm is earlier than 5.1SP1RP4, the installer uninstalls VRTSaslapm of the old version and reinstalls the latest one.

During the rolling back, VRTSaslapm doesn't revert to the earlier version that you roll back to. If you need to use VRTSaslapm of an earlier version, uninstall the current VRTSaslapm and reinstall VRTSaslapm of a specific version after rolling back.

Changes related to Veritas Cluster Server

Veritas Cluster Server includes the following changes in 5.1 SP1 RP4:

Enhancement to use different RouteOptions while adding and deleting the route

If you set the RouteOptions attribute to include a modifier that applies only to adding a route, then if you take the resource offline, the route delete command fails.

Symantec has added the DeleteRouteOptions to the `IPMultiNICB` resource configuration to fix this issue.

- If you set the DeleteRouteOptions attribute, VCS uses its value to implement the route delete command, and the value of the RouteOptions attribute to add a route.

- If you do not set the DeleteRouteOptions attribute, VCS uses the RouteOptions attribute to add or delete a route.

- If you do not set the RouteOptions value, VCS also ignores the DeleteRouteOptions value.

Zpool agent supports the `DeviceDir` attribute to reduce the time taken for zpool imports

The Zpool agent now supports the `DeviceDir` attribute, which specifies the directories that the `zpool import` command searches for devices or files. Each directory is passed to the `zpool import` command with the `-d` option. The directory path must exist and be absolute. No default value for this attribute. The value should be strings. For example: `DeviceDir = { "/dev/vx/dmp" }`

You can also use the DeviceDir attribute to reduce the time taken for the zpool import operations. To reduce the time taken for zpool imports, create links to the devices constituting the zpool in a directory and provide that directory path in the `DeviceDir` attribute.

For more information, see the `zpool(1M)` manual page.

Db2udb agent support is extended to DB2 10.1

The Db2udb agent support is extended to DB2 10.1.

Changes in Veritas Storage Foundation for Oracle RAC

This release introduces script-based installer support for configuring Highly Available IP (HAIP) addresses on SF Oracle RAC nodes running Oracle RAC 11.2.0.2 and later versions.

The Oracle Network Configuration menu now displays the following options:

```
1)  Configure private IP addresses (For HAIP Configuration)
    - For Oracle RAC 11.2.0.2 and later
2)  Configure private IP addresses (PrivNIC Configuration)
    - For Oracle RAC prior to 11.2.0.2
3)  Configure private IP addresses (MultiPrivNIC Configuration)
    - For Oracle RAC prior to 11.2.0.2
4)  Exit SF Oracle RAC Configuration
b)  Back to previous menu
```

System requirements

This section describes the system requirements for this release

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.

Supported Solaris operating systems

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

- SPARC

    - Solaris 10 Update 7, 8, 9, 10, and 11

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

    - Oracle VM for SPARC Version 1.3, 2.0, 2.1, 2.2, and 3.0

- x86-x64

■ Solaris 10 Update 7, 8, 9, 10, and 11

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

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

See "Required Solaris patches" on page 21.

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

http://www.symantec.com/docs/TECH75362

Required Solaris patches

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

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

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

Table 1-3   Solaris SPARC patches

Operating system	Oracle patch number
Solaris 9	114477-04 122300-29 (required for Live Upgrade) Patches required for FS, SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC: ■ For Solaris 9 Update 7 (to bring OS kernel to FS patch 122300 level): 117171-17 113073-14 ■ For Solaris 9 Update 7, Update 8 and Update 9: 112908-33 (required for 122300-10) 118558-39 (required for 122300-10) 122300-10
Solaris 10	119254-06 119042-02 113280-10 125731-02 128306-05 127111-01

The following patches (or a later revision of those patches) are required for Solaris x86-x64:

Table 1-4   Solaris x86-x64 patches

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

Database requirements

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

http://www.symantec.com/docs/DOC4039

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 in a cluster, use 1 GB of memory

  - More than eight nodes in a cluster, 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 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 5.1 SP1 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

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

- Veritas Volume Manager fixed issues

- Veritas File System fixed issues

- Veritas Storage Foundation Cluster File System fixed issues

- Veritas Storage Foundation for Oracle RAC fixed issues

- Veritas Cluster Server fixed issues

- Veritas Storage Foundation for Databases (SFDB) tools fixed issues

- Veritas Perl Redistribution fixed issues

Fixed issues related to Installation

This section describes the fixed issues related to the product installer in 5.1 SP1 RP4.

Table 1-5   fixed issues related to Installation

Fixed Issues	Description
3282042 2936047	The 144159-01 patch and the 143687-03 patch may get hang during its installation on Solaris SPARC 10.

Table 1-5   fixed issues related to Installation *(continued)*

Fixed Issues	Description
3243089	During a live upgrade the installation process takes more time than expected.
3230218	The Common Product Installer (CPI) based installation fails on Solaris version 11 because the nslookup command is missing.

Veritas Volume Manager fixed issues

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP4

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP4.

Table 1-6 describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP4.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues

Fixed issues	Description
3268905	After reboot, the non-root zpools created using DMP device go into FAULTED state and DMP device state is shown as UNAVAIL.
3264358	The vxdmpadm(1M) command reports "awk" errors when the dmp_native_support tuneable is set to "on".
3261485	vxcdsconvert failed with error "Unable to initialize the disk as a CDS disk".
3259926	vxdmpadm command fails to enable the paths when option '-f' is provided.
3225910	Device paths remain in the Enable state in DMP database after being disconnected from the array.
3194358	Continuous messages in syslog with EMC not-ready (NR) Logical units.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
3158323	In VVR environment, with multiple secondaries, if SRL overflows for rlinks at different times, it may result into vxconfigd(1M) daemon hang on the Primary.
3146715	Rlinks do not connect with NAT configurations on Little Endian Architecture.
3139983	DMP failed I/O without SCSI/HBA error.
3136272	Diskgroup import with "-o noreonline" option takes much longer time in 51SP1RP2 than on 5.0.1 for clone dg import.
3131071	VxVM patch installation in Solaris Alternate Boot Environment (ABE) results in data corruption.
3130353	Continuous disable/enable path messages are seen on the console for EMC NR (Not Ready) devices.
3116990	Syslog message is filled with lots of extra write protected message.
3102114	System crash during 'vxsnap restore' operation can lead to vxconfigd dump core after system reboot.
3101419	In CVR environment, when SRL overflows, IO of the rvg hangs for a long time.
3091916	SCSI i/o errors overflow the syslog.
3090667	System panics/hangs while executing vxdisk -o thin,fssize list as part of VOM SF discovery.
3090488	Memory leaks observed in device discovery code path of VxVM.
3087893	EMC TPD emcpower names are changing every reboot with VxVM.
3087250	In CVM, host joining a cluster might take little longer.

Table 1-6 Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
3077582	Interfaces to get and reset failio flag on disk after path failure recovered.
3076093	installrp command hung while upgradation from 5.1 to 5.1SP1 + RP3.
3067784	vxresize grow/shrink operation dumps core.
3065072	Data loss during import of clone diskgroup when some of the disks are missing and import options useclonedev and updateid are specified.
3058746	When the dmp disks of one RAID volume group is disabled, IO of the other volume group hangs.
3041014	Beautify error messages seen during relayout operation.
3038684	Restore daemon enables the paths of BCV NR devices.
3038382	The vxlufinish(1M) command runs 'fuser -k' on non-root file systems, which is unexpected.
3031796	Snapshot reattach operation fails if any other snapshot of the primary volume is not accessible.
3015181	IO can hang on all the nodes of a cluster when the complete non-A/A class of storage is disconnected.
3012929	vxconfigbackup gives errors when disk names are changed.
3006245	System panics infrequently while executing any snapshot operation.
3002498	When a disk is initialized with the "vxdisk -f init <DANAME>" command, vxconfigd(1M) dumps core.
2993667	Setting CDS attribute on a DG succeeds even with some in-accessible disks.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2992667	Change the framework for SAN of VIS from FC-switcher to direct connection, can not get the new DMP disk by "vxdisk scandisks".
2986596	DG imported with mix of standard and clone luns may lead to data corruption.
2979824	vxdiskadm bug results in exclusion of unintended paths.
2969844	Device discovery failure shouldn't cause DMP database to be destroyed completely.
2969335	Node leaving cluster while instant operation is in progress hangs in kernel and can not join back to cluster unless rebooted.
2966990	primary master meets IO hang when reboot primary slave.
2964547	"Cannot load module 'misc/ted'" message observed during system reboot.
2962010	Replication hangs upon resizing SRL (Storage Replicator Log).
2959733	Handling device path reconfiguration incase the device paths are moved across LUNs or enclosures to prevent vxconfigd coredump.
2959333	CDS flag is not listed for disabled CDS dgs.
2959325	vxconfigd dumps core while performing the DG(disk Group) move operation.
2957556	vxdisksetup' command fails when tpdmode is native and enclosure based naming scheme is on.
2957555	vxconfigd on CVM master node hangs in userland during vxsnap restore operation.
2952553	Refresh of a snapshot should not be allowed from a different source volume without force option.
2938710	vxassist dumps core while doing relayout.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2935771	In VVR environment, RLINK disconnects after master switch.
2933688	Continuing I/O even after DCPA bit is set, may result in data corruption.
2933476	Enhancement of vxdisk resize failure error messages.
2932214	After performing vxdisk resize on simple disk it went into "online invalid" state.
2929206	When turning on the dmp_native_support tunable with Solaris10 U10 and onwards, the Zettabyte File System (ZFS) pools are seen on the OS device paths but not on the dynamic multipathing (DMP) devices.
2921816	System panics while starting replication after disabling DCM volumes.
2921147	udid_mismatch flag is absent on a clone disk when source disk is unavailable.
2919720	vxconfigd core in rec_lock1_5().
2916911	vxconfigd(1M) daemon could send a VOL_DIO_READ request before the device is open.
2915063	Rebooting VIS array having mirror volumes, master node panicked and other nodes CVM FAULTED.
2912263	On the Solaris SPARC operating system, the vxdmpadm exclude command fails to exclude a controller specified in /etc/vx/vxvm.exclude file.
2911040	Restore from a cascaded snapshot leaves the volume in unusable state if any cascaded snapshot is in detached state.
2910367	When SRL on seondary site is disabled, secondary paniced.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2909668	In case of multiple sets of cloned disks of same source diskgroup, import of 2nd set clone disks fails if first set of clone disks were imported with updateid.
2899173	vxconfigd hang after executing command "vradmin stoprep".
2898547	vradmind on VVR Secondary Site dumps core, when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2898324	UMR errors reported by Purify tool in "vradmind migrate" command.
2893530	With no VVR configuration, when system is rebooted, it paniced.
2886333	vxdg join command should not allow mixing clone & non-clone disks in a DiskGroup.
2882566	No error is observed when a disk, which is removed from a disk group using the "vxdg rmdisk -k" command is added to another disk group.
2882312	Faulting SRL in middle of I/O load and issuing immediate read on data written during SRL fault can return old data.
2866059	Improving error messages hit during vxdisk resize operation.
2861011	The "vxdisk -g <dgname> resize <diskname>" command fails with an error for the Cross-platform Data Sharing(CDS) formatted disk.
2856579	When disk is resized from less than 1TB to greater than 1TB, "EFI PART" is missing in the primary label.
2855707	IO hangs with SUN6540 array during path fault injection test.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2845383	The site gets detached if the plex detach operation is performed with the site-consistency set to off.
2837717	"vxdisk(1M) resize" command fails if 'da name' is specified.
2825102	CVM reconfiguration and VxVM transaction code paths can simultaneously access volume device list resulting in data corruption.
2815441	After the vxresize operation vxvm volume can not be mounted.
2814891	vxconfigrestore doesn't work properly if SCSI page 83 inquiry returns more than one FPCH name identifier for single LUN.
2812161	Detach of primary rlink results in vxconfigd hang on secondary cluster nodes.
2787908	vxconfigd core dumped on start vxrecover with force option due to missing disks on slave and reconfig in parallel.
2779580	Secondary node gives configuration error (no Primary RVG) after reboot of master node on Primary site.
2746907	vxconfigd can hang on all nodes during a node join with heavy IO load on master.
2743926	DMP restore daemon fails to restart during system boot.
2735364	[VxVM]When destroying a VxVM diskgroup which contains CLONED devices, the VXVM clone_disk flag should be automatically turned off.
2692012	vxevac move error message needs to be enhanced to be less generic and give clear message for failure.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2685230	In CVR environment, IO hang on primary master, after rebooting master followed by slave node with logowner as slave.
2631369	When vxconfigd is started with -x nothreads (single- threaded) option, cluster reconfiguration such as node join and VxVM commands on shared diskgroup takes more time to complete.
2599887	PINNACLE DMP: disabled paths are not getting excluded from VM control.
2588771	System panics while disabling multi-controller enclosure.
2567618	VRTSexplorer coredumps with segmentation fault in checkhbaapi/print_target_map_entry.
2530536	DMP internal db repeatedly does reconfiguration when disable DMP path using CLI.
2518067	Interlock vxdmpadm disable operations with our DMP IO error handling code path.
2497074	"Configuration daemon error 441" while trying to stop a volume using vxvol(1M) on CDS EFI disks.
2488323	Write on volumes with links could hang if the volume has snapshots as well.
2477418	In VVR environment, logowner node on the secondary panics in low memory situations.
2438536	Reattaching a site after it was either manually detached or detached due to storage inaccessibility, causes data corruption.
2422535	Changes on the Veritas Volume Manager (VxVM) recovery operations are not retained after the patch or package upgrade.

Table 1-6 Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2398954	Machine panics while doing I/O on a VxFS mounted instant snapshot with ODM smartsync enabled.
2270686	vxconfigd stuck on master during vxsnap prepare & node join operation.
2270593	Shared disk group went into disabled state when vxconfigd daemon was restarted on master node followed by node join.
2257733	pinnacle: vxconfigd: ddl_search_and_place() in volddl_claim.c causes memory leak.
2255018	vxplex core dump during relayout from concat to raid5.
2236443	[Enhancement]Diskgroup import failure should be made fencing aware, in place of VxVM vxdmp V-5-0-0 I/O error message.
2233225	vxassist grow doesn't synchronize the plexes of mirrored volume if volume is grown more than 1G or default stripe-mirror-col-split-trigger-pt.
2165920	vxrelocd(1M) daemon creates defunct(zombie) process.
2106530	vxresize(1M) fails for a data volume in rootdg if filesystem is mounted using block device reference as 'bootdg'.
2101093	System panic is observed in the dmp_signal_event() function.
2091520	The ability to move the configdb placement from one disk to another using "vxdisk set <disk> keepmeta=[always|skip|default]" command.
2054606	System panic occurs during DMP (Dynamic Multi-Pathing) driver unload operation.

Table 1-6   Veritas Volume Manager 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
1982965	vxdg import fails if da-name is based on naming scheme which is different from the prevailing naming scheme on the host.
1973983	vxunreloc fails when dco plex is in DISABLED state.
1901838	Incorrect setting of Nolicense flag can lead to dmp database inconsistency.
1871118	VxVM fails to recognize ASM disk and the disk is shown in 'error' state in 'vxdisk list' output.
1783763	vxconfigd hang seen during configuration change operation.
1540834	Modify vxdg(1M) man page to document 0-4 are reserved minor numbers and cannot be used in re-minoring operations.
1289985	vxconfigd core dumps upon running "vxdctl enable" command.

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP3 P2.

Table 1-7   Veritas Volume Manager 5.1 SP1 RP3 P2 fixed issues

Fixed issues	Description
2567618	VRTSexplorer core dumps in checkhbaapi/print_target_map_entry
1190117	The 'vxdisk -f init' command can corrupt/overwrite some of the public region data.
2875962	When an upgrade install is performed from VxVM 5.0MPx to VxVM 5.1(and higher) the installation script may give an message.
2398954	Machine panics while doing I/O on a VxFS mounted instant snapshot with ODM smartsync enabled.

Table 1-7   Veritas Volume Manager 5.1 SP1 RP3 P2 fixed issues *(continued)*

Fixed issues	Description
2441615	Cannot create zpool of correct size with LUNs of sizes greater than 2TB.
2970368	SRDF-R2 WD(write-disabled)devices are shown in error state and lots of path enable/disable messages are generated in /etc/vx/dmpevents.log file.
3008423	vxdisksetup fails and throws an error message
3002770	Excluding a controller from VxVM using the command "vxdmpadm exclude ctlr=" leads to panic
2622536	Under a heavy I/O load, write I/Os on VVR Primary logowner takes a very long time to complete.
3025713	VxVM commands "vxdg adddisk" and "vxdg rmdisk" takes long time (approximately 90 seconds) and I/Os hang during the command execution.
2277359	vxdisksetup -fi succeeds on ZFS LUN's but it may fail
2910043	Frequent swapin/swapout seen due to higher order memory requests
2834046	VxVM dynamically reminors all the volumes during DG import if the DG base minor numbers are not in the correct pool.
2859470	The EMC SRDF-R2 disk may go in error state when you create EFI label on the R1 disk.
2878876	vxconfigd, VxVM configuration daemon dumps core
2189812	While executing 'vxdisk updateudid' on a disk which is in 'online invalid' state causes vxconfigd to dump cor

Table 1-7 Veritas Volume Manager 5.1 SP1 RP3 P2 fixed issues *(continued)*

Fixed issues	Description
2915751	Solaris machine panics while resizing CDS-EFI LUN or CDS VTOC to EFI conversion case where new size of resize is greater than 1TB.
2919714	On a THIN lun, vxevac returns 0 without migrating unmounted VxFS volumes.
2928987	vxconfigd hung is observed when IO failed by OS layer.
2940446	I/O can hang on volume with space optimized snapshot if the underlying cache object is of very large size.
2406096	vxconfigd, VxVM configuration daemon, dumps core
2671241	When the DRL log plex is configured in a volume, vxnotify doesn't report volume enabled message.
2575051	In a CVM environment, Master switch or master takeover operations results in panic.
2223250	Node join fails if the recovery for the leaving node is not completed.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP3.

Table 1-8 describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP3.

Table 1-8 Veritas Volume Manager 5.1 SP1 RP3 fixed issues

Fixed issues	Description
925653	Node join fails for higher CVMTimeout value.
2915302	`./scripts/ddl/dmp_add_path.tc` failing on solaris sparc.
2910987	vxspec and `vxconfigd` cannot start after applying VxVM 5.1SP1RP3 patch. `vxconfigd` failed to start after rejecting VxVM 5.1SP1RP3 patch.

Table 1-8   Veritas Volume Manager 5.1 SP1 RP3 fixed issues *(continued)*

Fixed issues	Description
2866997	VxVM Disk initialization fails as an un-initialized variable gets an unexpected value after OS patch installation.
2858853	After master switch, `vxconfigd` dumps core on old master.
2838059	VVR Secondary panic in `vol_rv_update_expected_pos`.
2836798	In VxVM, resizing simple EFI disk fails and causes system panic/hang.
2836528	Unable to grow LUN dynamically on Solaris x86 using the `vxdisk resize` command.
2826125	VxVM script daemon is terminated abnormally on its invocation.
2818840	Enhance the `vxdmpasm` utility to support various permissions and "root:non-system" ownership can be set persistently.
2815517	`vxdg adddisk` allows mixing of clone & non-clone disks in a DiskGroup.
2807158	On Solaris platform, sometimes system can hang during VM upgrade or patch installation.
2801962	Growing a volume takes significantly large time when the volume has version 20 DCO attached to it.
2775960	In secondary CVR case, IO hang is seen on a DG during SRL disable activity on other DG.
2774406	System may panic while accessing data change map volume.
2763206	The `vxdisk rm` command dumps core when disk name of very large length is given.
2760181	Panic hit on secondary slave during logowner operation.
2756059	System may panic when large `cross-dg` mirrored volume is started at boot.
2754819	Diskgroup rebuild through `vxmake -d` loops infinitely if the diskgroup configuration has multiple objects on a single cache object.
2753954	When a cable is disconnected from one port of a dual-port FC HBA, the paths via another port are marked as SUSPECT PATH.
2742706	Panic due to mutex not being released in vxlo_open.

Table 1-8   Veritas Volume Manager 5.1 SP1 RP3 fixed issues *(continued)*

Fixed issues	Description
2739709	Disk group rebuild fails as the links between volume and vset were missing from the `vxprint -D -` output.
2739601	VVR: VRAS: repstatus output occasionally reports abnormal timestamp.
2735951	Uncorrectable write error is seen on subdisk when SCSI device/bus reset occurs.
2729911	IO errors are seen during controller reboot or array port disable/enable.
2715129	`Vxconfigd` hangs during Master takeover in a CVM (Clustered Volume Manager) environment.
2689845	Data disk can go in error state when data at the end of the first sector of the disk is same as MBR signature.
2688308	When re-import of disk group fails during master takeover, other shared disk groups should not be disabled.
2684558	vxesd dumps core on startup in libc.
2680343	Manual disable/enable of paths to an enclosure leads to system panic.
2664825	DiskGroup import fails when disk contains no valid UDID tag on config copy and config copy is disabled.
2657797	Starting a RAID5 volume fails, when one of the sub-disks in the RAID5 column starts at an offset greater than 1TB.
2656803	Race between vxnetd start and stop operations causes panic.
2653143	System panics while loading vxdmp driver during VxVM 5.1SP1 installation.
2648176	Performance difference on Master vs Slave during recovery via DCO.
2647975	Customer ran `hastop -local` and shared dg had splitbrain.
2637217	Document new storage allocation attribute support in vradmin man page for resizevol/resizesrl.
2627126	IO hang is seen due to IOs stuck at DMP level.
2627056	`vxmake -g DGNAME -d desc-file` fails with very large configuration due to memory leaks.

Table 1-8 Veritas Volume Manager 5.1 SP1 RP3 fixed issues *(continued)*

Fixed issues	Description
2626741	Using `vxassist -o` ordered and mediatype:hdd options together do not work as expected.
2626199	`vxdmpadm` list dmpnode printing incorrect path-type.
2620556	IO hung after SRL overflow.
2620555	IO hang due to SRL overflow & CVM reconfig.
2606709	IO hang when SRL overflow & reboot one node.
2606695	Machine panics in CVR (Clustered Volume Replicator) environment while performing I/O Operations.
2599526	IO hang is seen when DCM is zero.
2578336	Failed to online the cdsdisk.
2576602	`vxdg listtag` should give error message and display correct usage when executed with wrong syntax.
2575172	I/Os hung on master node after reboot the slave node.
2567618	VRTSexplorer coredumps in `checkhbaapi/print_target_map_entry`.
2566174	Null pointer dereference in volcvm_msg_rel_gslock() results in panic.
2561012	The offset of private(and/or public) region of disks are shown incorrect in the vxdisk list output which could lead to DG import problem as well as IO errors and system hang reported by VxFS or other applications.
2560843	Node join fails for higher CVMTimeout value.In VVR(Veritas Volume Replicator) setup I/Os can hang in slave nodes after one of the slave node is rebooted.
2560835	I/Os and `vxconfigd` hung on master node after slave is rebooted under heavy I/O load.
2556781	In cluster environment, import attempt of imported disk group may return wrong error.
2556467	Disabling all paths and reboot of the host causes losing of `/etc/vx/.vxdmprawdev` records.
2533015	Powerpath/SVM controlled disks are not tagged as SVM in vxdisk list output.

Table 1-8   Veritas Volume Manager 5.1 SP1 RP3 fixed issues *(continued)*

Fixed issues	Description
2526623	Memory leak detected in CVM code.
2516584	startup scripts use 'quit' instead of 'exit', causing empty directories in `/tmp`.
2513101	User data corrupted with disk label information.
2495332	`vxcdsconvert` fails if the private region of the disk to be converted is less than 1 MB.
2441937	`vxconfigrestore` precommit fails with awk errors.
2425259	vxdg join operation fails with VE_DDL_PROPERTY: Property not found in the list.
2413763	Uninitialized memory read results in a `vxconfigd` coredump.
2389554	The `vxdg listssbinfo` output is not correct.
2348199	vxconfig dumps core while importing a Disk Group.
2277558	`vxassist` outputs a misleading error message during snapshot related operations.
2257850	`vxdiskadm` leaks memory while performing operations related to enclosures.
2252680	`vxtask abort` does not appropriately cleanup the tasks.
2227678	Second rlink goes into DETACHED STALE state in multiple secondaries environment when SRL has overflowed for multiple rlinks.
2216951	`vxconfigd` dumps core because `chosen_rlist_delete()` hits NULL pointer in linked list of clone disks.
2149922	Record the diskgroup import and deport events in syslog.
2104887	vxdg import error message needs improvement for cloned diskgroup import failure.
2088426	Re-onlining of disks in DG during DG deport/destroy.
2024617	volpagemod_max_memsz setting in `/kernel/drv/vxio.conf` is not honoured by system boot.
2000585	`vxrecover` doesn't start remaining volumes if one of the volumes is removed during the `vxrecover` command run.

Table 1-8   Veritas Volume Manager 5.1 SP1 RP3 fixed issues *(continued)*

Fixed issues	Description
1903700	Removing mirror using `vxassist` does not work.
1675482	The `vxdg list` *dgname* command shows configuration copy in new failed state.
1431223	The `vradmin syncvol` and the `vradmin syncrvg` commands do not work if the remote diskgroup and vset names are specified when synchronizing vsets.
1291519	After multiple VVR migrate operations, `vrstat` fails to output statistics.

This section describes Veritas Volume Manager fixed issues in 51 SP1 RP2 P3.

Table 1-9   Veritas Volume Manager 51 SP1 RP2 P3 fixed issues

Fixed issues	Description
2771452	IO hung because of hung port deletion.
2741240	Invoking "vxdg join" operation during heavy IO load results in a transaction failure and leaves disks in an intermediate state.
2553942	LU issue with latest sol10u10Build image.
2370250	vxlufinish script is running 'fuser -k' on lofs file systems mounted under ABE, killing processes running under these FS in PBE.
2423701	Upgrade of VxVM caused change in permissions.
2439481	After doing live upgrade on encapsulated disk with mirror, mirror disk entry is not removed.
2729501	`vxdmpadm exclude vxvm path=<>` results in excluding unexpected set of paths.
2722850	DMP fail over hangs when the primary controller is disabled while I/O activity is ongoing.
2390998	System panicked during SAN reconfiguration because of the inconsistency in dmp device open count.
2617336	Solaris patch 147440- panics in vxioioctl.
2698860	`vxassist` mirror failed for thin LUN because `statvfs` failed.

Table 1-9   Veritas Volume Manager 51 SP1 RP2 P3 fixed issues *(continued)*

Fixed issues	Description
2700086	EMC BCV (NR) established devices are resulting in multiple dmp events messages (paths being disabled/enabled).
2700792	The VxVM volume configuration daemon may dump a core during the Cluster Volume Manager(CVM) startup.
2700486	`vradmind` coredumps when Primary and Secondary have the same hostname and an active Stats session exists on Primary.
2688747	Logowner local sequential I/Os starved with heavy I/O load on logclient.
2674465	Adding/removing new LUNs causes data corruption.
2675538	`vxdisk resize` may cause data corruption.
2553729	Disk groups do not get imported and 'clone_disk' flag is seen on non-clone disks after uprade of VxVM.
2666163	A small portion of possible memory leak in case of mix (clone and non-cloned) diskgroup import.
2649846	On SunCluster3. cldg create, cldg core dumps with umem allocator: redzone violation: write past end of buffer.
2643634	Message enhancement for a mixed(non-cloned and cloned) dg import.
2635476	Volume Manager does not recover a failed path.
2509291	"`vxconfigd`" daemon hangs if host side i/o paths are failing.
2061082	The `vxddladm -c assign names` command should work for devices with native support not enabled (VxVM labeled or TPD).
2605444	`vxdmpadm` disable/enable primary path (EFI labelled) in A/PF array results in all paths getting disabled.
2608849	VVR Logowner: local I/O starved with heavy I/O load from Logclient.
2621465	When detached disk after connectivity restoration is tried to reattach gives `Tagid conflict` error.
2419948	Race between the SRL flush due to SRL overflow and the kernel logging code, leads to a panic.
2589569	vxdisksetup on EFI disk is taking ~2- mins.

Table 1-9   Veritas Volume Manager 51 SP1 RP2 P3 fixed issues *(continued)*

Fixed issues	Description
2527289	Site consistency: Both sites become detached after data/dco plex failue at each site, leading to I/O cluster wide outage.
2495186	With TCP protocol used for replication, I/O throttling happens due to memory flow control.
2253970	Support per-disk maxiosize for private region I/Os.
2365486	In 2-nodes SFRAC configuration, after enabling ports systems panics due to improper order of acquire and release of locks.
2710579	Do not write backup labels for CDS disk - irrespective of disk size.

This section describes Veritas Volume Manager fixed issues in 51 SP1 RP2 P2.

Table 1-10   Veritas Volume Manager 51 SP1 RP2 P2 fixed issues

Fixed issues	Description
2185069	panic in vol_rv_mdship_srv_start().

This section describes Veritas Volume Manager fixed issues in 51 SP1 RP2 P1.

Table 1-11   Veritas Volume Manager 51 SP1 RP2 P1 fixed issues

Fixed issues	Description
2536667	Slave node panics when private region I/O and dg deport operation are executed simulatenously.
2524936	DG disabled after vold found the process file table is full.
2510523	The `ls -l` command hang during RMAN backup on VVR/RAC cluster.
2489350	Memory leak in VVR.
2484334	Panic in `dmp_stats_is_matching_group()`.
2483053	Primary Slave node runs out of memory, system hang on VRTSvxvm.
2438426	VxVM is failing to correctly discover ZFS LUNs presented via PP after excluding/including `libvxpp.so`.

Table 1-11   Veritas Volume Manager 51 SP1 RP2 P1 fixed issues *(continued)*

Fixed issues	Description
2432006	pending read count with kio cache is not decremented when read object is locked in transaction.
2431448	CVR:I/O hang while transitioning to DCM mode.
2428170	IO hung on Mirror volume and return error on DMP disk, but phydisk(/dev/sdbw) is OK.
2420386	Data corruption creating data in a vxfs filesystem, while being grown with vxresize on efi thinrclm disks.
2419803	Secondary Site panics in VVR (Veritas Volume Replicator).
2419486	Data corruption occurs on changing the naming scheme.
2390431	VVR: system crash dring DCM flush not finding the parent_sio volsiodone+.
2375011	User should be able to enable or disable dmp_native_support irrespective of active rpool.
2344186	Volume recovery is not clearing the need sync flag from volumes with DCO in BADLOG state. Thus nodes are unable to join the cluster.
2235382	IO hung in DMP while restoring a path in presence of pending IOs on local A/P class LUN.
2169726	CLONE : Disk group is imported using a Non-cloned and cloned disks, it can lead to data corruption.
2148851	`vxdisk resize` failed to resize the disk which is expanded physically from array console.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP2

Table 1-12 describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP2.

Table 1-12   Veritas Volume Manager 5.1 SP1 RP2 fixed issues

Fixed issues	Description
2480600	I/O permanent hung on master node when IO size larger than 512K, and 32+ threads write in parallel

Table 1-12 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
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 2 TB
2425551	IO hung for 6 mintues when reboot the slave node, if there is I/O on both master and slave
2424833	while autosync and deport is ongoing the primary logowner hits ted assert nmcom_send_msg_tcp
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
2411698	VVR:iohang: On I/O to both master and slave
2410845	Lots of 'reservation conflict' messages seen 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 array.info file can cause vxdmpadm commands to hang
2396293	I/Os loaded, sanboot failed with vxconfigd core dump
2387993	While testing including/excluding libvxpp.so vxconfigd goes into disabled mode

Table 1-12   Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2386120	Enhancement request to add diagnostic logging to help triage a CVM master takeover failure situation
2385680	VVR: vxio panic in 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
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
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
2337353	vxdmpadm include vxvm dmpnodename= 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

Table 1-12 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2317703	Vxesd/Vxconfigd leaks file descriptors
2316297	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 5.1SP1 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
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_parse.y 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
2232789	supporting NetApp Metro Cluster
2228531	cvm master vxconfigd process hung in vol_klog_lock()
2205108	SVS: 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

Table 1-12 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
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
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	dmp_deamon applying 1m continuous memory paging which is too large
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 5.1RP1
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
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
2369177	DDL: do_diskio function should be able to handle offset greater than 2TB
2346470	Excluding and including a LUN in a loop using vxdmpadm, triggers a huge memory leak
2339251	Newfs fails for volumes greater than 2Tb due to DKIOCGMEDIAINFO and DKIOCGETEFI ioctls failing

Table 1-12   Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2320669	SDS to VxVM conversion failed on Solaris
2291176	vxrootadm does not set dump device correctly with LANG=ja
2270880	On Solaris 10 SPARC, disk capacity will be truncated to 2TB for EFI labeled disk with size greater than 2TB
2268408	suppressing a powerpath disk's path using vxdiskadm 17-2 causes the disk to go in error state
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated
2244210	vxdisksetup failed with default parameters passed through CLI on EFI setup
2230716	SVM migration to VxVM fails to convert due to /etc/lvm/md.cf not being cleared
2226304	Cannot create 1TB+ ufs file system with Solaris 9
2216515	vxunreloc may corrupt boot disk, if original offsets are used
2215262	Netapp iSCSI LUN goes into error state while initializing via VEA GUI
2215256	Support of Oracle F5100 flash array on an Oracle X4470 server with new HBA
2196480	Initialization of VxVM cdsdisk layout fails on a disk of size less than 1 TB
2179259	DMP SCSI bypass needs to be enhanced to handle I/O greater than 2TB
2146833	vxmirror operation fails with error "Device has UFS FS on it"
2108152	vxconfigd goes to DISABLED state upon system reboot with NR (Not Ready) devices in configuration
2064490	Ensure vxcdsconvert works for greater than 1 TB CDS disks
2063152	vxbootadm failed with xml related error
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 1-13   Veritas Volume Manager 5.1 SP1 RP1 fixed issues

Fixed issues	Description
2160199	Master takeover fails as the upcoming Master could not import shared DG. SUN Bug ID is 6988420.
2148682	while shipping a command node hangs in master selection on slave nodes and master update on master node
2181631	Striped-mirror volume cannot be grown across sites with -oallowspansites with DRL
2133503	Renaming enclosure results in dmpevents.log reporting Mode for Enclosure has changed from Private to Private
2200670	vxattachd does not recover disks if disk group is not imported
2191693	'vxdmpadm native list' command is not displaying any output nor error
2080730	vxvm/vxdmp exclude file contents after updation should be consistent via vxdiskadm and vxdmpadm
2129989	EVA ASL should report an error message if pref_bit is not set for a LUN
2129477	vxdisk reclaim command fails after resize operation.
2194492	VxVM-ASM co-existence enablement
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core
2166682	checks needed to make sure that a plex is active before reading from it during fsvm mirror read interface
2088007	Possibility of reviving only secondary paths in DMP
2188590	An ilock acquired by a SLAVE node for a read on a DCL object can lead to IO hang when the node becomes MASTER before reading is done
2015467	Performance improvement work for NetBackup 6.5.5 on SF 5.1 VxVM mapping provider
1829285	vxconfigd coredumps while assigning unique native name to a disk

Table 1-13   Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2226813	VVR: rlinks remain disconnected with UDP protocol if data ports are specified
2227923	renaming of enclosure name is not persistent
2172488	FMR: with dco version 0 restore operation doesn't sync the existing snapshot mirrors
2183984	System panics due to race condition while updating DMP I/O statistics
2038928	creation of pre 5.1SP1 (older) version diskgroup fails
1970560	When vxconfigd is idle (which is not shipping the command) slave dies and command shipping is in progress, vxconfigd core dumped on Master
2082450	In case of failure, vxdisk resize should display more meaningful error message
1869002	Introduction of Circular buffer at vold level for master-slave communication.
1426480	VOLCVM_CLEAR_PR ioctl does not propogate the error returned by DMP to the caller
2105547	tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault
1959513	Propagate -o noreonline option of disk group import to slave nodes
1940052	vxconfigd hung on Master after removing the hba alias from zone and node leave followed by join
2199496	Data Corruption seen with "site mirror" Campus Cluster feature
2234844	asm2vxfs conversion fails
2215216	vxkprint does not report TP related values
2141820	vxrootadm:boot entry not created after second split on Solaris opetron
2211283	system goes in maintainence mode after removal of 51SP1 patch with boot disk encapsulated
2148738	vxdmpadm invoked by vxvm-sysboot is killed during system boot
2062190	vxrootadm split/join operation fails when there is a rvg present in the rootdg/backupdg

Table 1-13   Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2130744	1TB disks are seen in error state on sol9u7 and vxdisksetup also fails.
2121183	installmp/rp/patchadd results in a partially installed state/unusable
2164906	Need to enhance vxlustart script to support new option to luupgrade command in Sol10U9
2209391	'vxdisk init' failure modifies disk label from sliced to EFI type for devices >1TB on Solaris 10

Veritas File System fixed issues

This section describes Veritas File System fixed issues in 5.1 SP1 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2, and 5.1 SP1 RP1.

Veritas File System: Issues fixed in 5.1 SP1 RP4

This section describes Veritas File System: fixed issues in 5.1 SP1 RP4.

Veritas File System: Issues fixed in 5.1 SP1 RP4 describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP4.

Table 1-14   Veritas File System 5.1 SP1 RP4 fixed issues

Fixed issues	Description
3272896	Internal stress test on the local mount hits a deadlock.
3270357	The fsck (1m) command fails to clean the corrupt file system during the internal 'noise' test.
3259634	A CFS that has more than 4 GB blocks is corrupted because the blocks containing some file system metadata gets eliminated.
3257314	On systems installed with the SFORA/SFRAC stacks, when the DBED operations like the dbdst_obj_move (1M) command are run, the operation may fail.
3257142	The fscdsconv(1M) man page does not show examples on Solaris x86 machines.

Table 1-14 Veritas File System 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
3240635	In a CFS environment, when a checkpoint is mount using the mount(1M) command the system may panic.
3240403	The fidtovp()system call may cause panic in the vx_itryhold_locked () function.
3214816	When you create and delete the inodes of a user frequently with the DELICACHE feature enabled, the user quota file becomes corrupt.
3206266	During an internal noise test, the "f:vx_purge_nattr:1" assertion fails.
3192985	Checkpoints quota usage on CFS can be negative.
3156701	The system panics because of incorrect credentials reference counts caused by vx_mkdir().
3121933	The pwrite(2) fails with the EOPNOTSUPP error.
3101418	Oracle startup displays ORA-01513 errors.
3092114	The information output displayed by the "df -i" command may be inaccurate for cluster mounted file systems.
3073372	On larger CPUs/memory configurations with partition directory feature enabled operations such as find, ls may seem to be slower.
3072036	Read operations from secondary node in CFS can sometimes fail with the ENXIO error code.
3046983	Invalid CFS node number in ".__fsppadm_fclextract", causes the DST policy enforcement failure.
3013950	Solaris 11 update 1 validation encounters the following test assert "f:vx_info_init:2" during internal testing.

Table 1-14   Veritas File System 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
3008451	In a Cluster File System (CFS) environment, shutting down the cluster may panic one of the nodes with a null pointer dereference.
3003679	When running the fsppadm(1M) command and removing a file with the named stream attributes (nattr) at the same time, the file system does not respond.
2997334	On a Veritas File System (VxFS) file system during the internal umount test the "vx_detach_fset:10k" assertion fails.
2972299	The initial and subsequent reads on the directory with many symbolic links is very slow.
2972183	The fsppadm(1M) enforce command takes a long time on the secondary nodes compared to the primary nodes.
2966277	Systems with high file-system activity like read/write/open/lookup may panic the system.
2963763	When the thin_friendly_alloc() and deliache_enable() functionality is enabled, VxFS may hit a deadlock.
2926684	In rare cases the system may panic while performing a logged write.
2923105	Removal of the VxFS module from the kernel takes a longer time.
2912089	The system becomes unresponsive while growing a file through vx_growfile in a fragmented file system.
2908391	It takes a longer time to remove checkpoints from the Veritas File System (VxFS) file system with a large number of files.
2906018	The vx_iread errors are displayed after successful log replay and mount of the file system.

Table 1-14   Veritas File System 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2905820	If the file is being read via the NFSv4 client, then removing the same file on the NFSv4 server may hang if the file system is VxFS.
2899907	On CFS, some file-system operations like vxcompress utility and de-duplication fail to respond.
2895743	Accessing named attributes for some files stored in CFS (Cluster File Systems) seems to be slow.
2858683	Reserve extent attributes changed after vxrestore, only for files greater than 8192bytes.
2838471	Cannot change the group or owner of the file by any of the mount option.
2834192	Unable to mount the file system after the full fsck(1M) utility is run.
2827751	High kernel memory allocation is observed when Oracle Disk Manager (ODM) is used with non-VxVM devices.
2822984	The extendfs(1M) command fails when it attempts to extend a file system that is greater than 2 terabytes.
2806466	A reclaim operation on a file system that is mounted on a Logical Volume Manager (LVM) may panic the system.
2756779	The read and write performances are slow on Cluster File System (CFS) when it runs applications that rely on the POSIX file-record using the fcntl lock.
2750860	Performance of the write operation with small request size may degrade on a large file system.
2735912	The performance of tier relocation using the fsppadm(1M) enforce command degrades while migrating a large number of files.

Table 1-14   Veritas File System 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2732427	A Cluster mounted file-system may hang and become unresponsive.
2703747	The Cluster File System (CFS) failover may take over 20 minutes to complete.
2700814	The system panics while running the internal fsmig conformance testing.
2695390	Accessing a vnode from cbdnlc cache hits an assert during internal testing.
2667658	The 'fscdsconv endian' conversion operation fails because of a macro overflow.
2594774	The "vx_msgprint" assert is observed several times in the internal Cluster File System (CFS) testing.
2439261	When the vx_fiostats_tunable value is changed from zero to non-zero, the system panics.
2370627	fsck(1M) core dumps while running internal tests.

Veritas File System: Issues fixed in 5.1 SP1 RP3

This section describes Veritas File System fixed issues in 5.1 SP1 RP3.

Table 1-15 describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP3.

Table 1-15   Veritas File System 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2841059	full fsck fails to clear the corruption in attribute inode 15.
2839871	process hung in vx_extentalloc_delicache.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2597347	fsck segmentation fault bc_rgetblk ().

Table 1-16   Veritas File System 5.1 SP1 RP2 P2 fixed issues

Fixed issues	Description
2753944	VxFS hang in `vx_pd_create`.
2715028	`fsadm -d` hang during `vx_dircompact`.
2712392	thread hung in vxfs looks like to be 2689750 but customer is on 5.1SP1.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2677231	Internal Test quota getting hang.
2670022	Duplicate file names can be seen in a directory.
2651922	Performance degradation of 'll' and high SYS% CPU in `vx_ireuse()`.
2566875	A write(2) operation exceeding the quota limit fails with an EDQUOT error.
2426648	The fsck(1M) operation on a Veritas File Systems (vxfs) can fail with an EINVAL error while validating the inode attributes.
2291410	Internal quotas test hit assert of " BAD TRAP: type=31 rp=2a100909490 addr=0 mmu_fsr=0 occurred in module "unix" due to a NULL pointer dereference".
2086902	system crash when spinlock was held too long.

This section describes Veritas File System fixed issues in 5.1 SP1 RP2 P1.

Table 1-17   Veritas File System 5.1 SP1 RP2 P1 fixed issues

Fixed issues	Description
2631276	QIO does not work in a partitioned directory.
2624262	Dedup:fsdedup.bin hit oops at vx_bc_do_brelse.
2600168	The -p option of cp_vxfs command does not work correctly in solaris.
2599590	Expanding or shrinking a DLV5 file system using the fsadm(1M)command causes a system panic.
2588593	usage of volume in the output of df command do not back to beginning after created files and deleted files.
2561334	using flockfile() instead of adding new code to take lock on .__fsppadm_enforcesq file descriptor before writing into it.

Table 1-17   Veritas File System 5.1 SP1 RP2 P1 fixed issues *(continued)*

Fixed issues	Description
2536269	The access time for a zero size file may not updated correctly.
2528819	VxFS thread create warning messages.
2526174	Wrong offset calculation affects replication functionality.
2515459	mount command still hanged even with the fix of e1466351.
2515380	ff_vxfs ERROR: V-3-24347: program limit of 30701385 exceeded.
2350956	fsck fails with the following message ASSERT(devid == 0 || (start == VX_HOLE && devid ==VX_DEVID_HOLE)) failed.
2332314	Writing to a clone using the FDD fails for async I/Os.
2326037	Write operation on a Cluster mounted filesystem may fails with ENOENT.
2271797	On disk and the in core structures may be out of sync in case of clone writes.
2246127	Mount should perform read ahead on IAUs.
1590963	Requirement for vx_maxlink tunable on Linux.

Veritas File System: Issues fixed in 5.1 SP1 RP2

Table 1-18 describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP2.

Table 1-18   Veritas File System fixed issues

Fixed issues	Description
2340953	cfs.stress.enterprise hit an assert f:vx_iget:1a.
2481984	file system will hang if customer creates 400 shares
2247387	LM stress.S3 test hit an assert "vx_ino_update:2"
2486589	threads blocked behind vx_ireuse_steal
2431674	panic in vx_common_msgprint() via vx_inactive()
2480935	fsppadm: ERROR: V-3-26626: File Change Log IOTEMP and ACCESSTEMP index creation failure for /vx/fsvm with message Argument list too long

Table 1-18   Veritas File System fixed issues *(continued)*

Fixed issues	Description
2384861	CFS stress+reconfig test hit assert "f:vx_do_filesnap:1b".
2333907	LM Conformance->fsmig->migvops test get fail.
2413172	There is a priority 1 issue reported by AXA Rosenburg for Filestore replication and issue seems related to VxFS
2399228	TRuncate up size updates can be missed
2412604	it does not work when set homedir user softlimit numspace quota after generate data
2422574	Reboot one node and the node can't mount file system , after turn on the homedir quota on
2283893	Add functionality of free space defragmentation through fsadm.
2372093	new fsadm -C hung
2387609	User quota corruption
2371710	user quota information corrupts on 5.1SP1
2384831	vxfs panic in iput() from vx_softcnt_flush() ,after filesystem full fsck,and run reboot
2407895	System panic in atomic_add_32_nv because of bad address.
2146573	qdetails performance downgraded on Aug 16th.
2399178	fsck : pass2c needs performance enhancements
2374887	Accessing FS hung. FS marked full fsck after reboot of node.
2283315	cfs-stress_S5 hits assert of "f:vx_reorg_emap:10 via vx_extmap_reorg"
2368737	RCQ processing code should set FULLFSCK flag if it finds a corrupt indirect block.
1956458	fsckpt_fbmap for changed blocks failed with ENXIO due to inode mapped to hole in ILIST of down stream checkpoint
2337470	In the process of shrink fs, the fs out of inodes, fs version is 5.0MP4HF*
2332460	vxedquota slow on some systems
2300682	Question about IOTemp on fsppadm query

Table 1-18 Veritas File System fixed issues *(continued)*

Fixed issues	Description
2272072	Threads stuck in vx_rwsleep_rec_lock_em
2290800	investigation on ilist HOLE
2192895	VxFS 5.0MP3RP4 Panic while set/get acls - possible race condition
2172485	Metadata was not updated correctly after write() with O_SYNC flag.
2059611	Panic in vx_unlockmap() due to NULL ml_tranp
2282201	vxdump core dumped whilst backing up layout 7 local VxFS file system
2316094	There was discrepancy between vxi_bcache_maxkbyte and vx_bc_bufhwm.
2419989	ncheck -i does not limit output to the specified inodes when using -o device/block/sector
2074806	dm_punch_hole request does not invalidate pages
2296107	Operation not applicable appear on fsppadm query result
2246579	Panic at getblk() when growing a full filesystem with fsadm
2061177	fsadm -de' command erroring with 'bad file number' on filesystem(s) on 5.0MP3RP1
1475345	write() system call hangs for over 10 seconds on VxFS 3.5 on 11.23
2271878	WARNING message with fsmigadm start command
2251223	df -h after removing files takes 10 seconds
2253617	LM stress aborted due to "run_fsck : Failed to full fsck cleanly".
2220300	vx_sched' is hogging CPU resources.
2226762	vx_ntran overflow causes a lot of buffer flushes
1949445	hang due to large number of files in a directory
1296491	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2230351	After installation of SF 5.1SP1, Solaris 10 system PANICed 8 times due to \"BAD TRAP: type=31 rp=2a1001c6e40 addr=28 mmu_fsr=0 occurred in module \"vxfs\" due to a NULL pointer dereference\"

Table 1-18   Veritas File System fixed issues *(continued)*

Fixed issues	Description
2169324	5.1SP1 sol_sprac 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 1-19   Veritas File System fixed issues

Fixed issues	Description
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.
2030119	fsppadm core dumps when analysing a badly formatted XML file, is resolved
2162822	During online migration from ufs to vxfs, df command returns a non-zero return value.
2169273	During online migration, nfs export of the migrating file system leads to system panic
2177253	A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for vxupgrade purposes
2178147	Linking a IFSOC file now properly calls `vx_dotdot_op`(), which fixes the cause of a corrupted inode.
2184528	`fsck` no longer fails to repair corrupt directory blocks that have duplicate directory entries.
2178147	Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system
2221623	Fixed a performance loss due to a delxwri_ilist spin lock with the default values for `vx_idelxwri_timelag`.
2228311	Support for mostly online migration from ZFS to VxFS

Veritas Storage Foundation Cluster File System fixed issues

This section describes the Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2, and 5.1 SP1 RP1.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP4

This section describes Veritas Storage Foundation Cluster File System: fixed issues in 5.1 SP1 RP4.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP4on page 61 describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP4.

Table 1-20   Veritas Storage Foundation Cluster File System 5.1 SP1 RP4 fixed issues

Fixed issues	Description
2720034	The vxfsckd(1M) daemon does not restart after being killed manually.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP3.

Table 1-21 describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP3.

Table 1-21   Veritas Storage Foundation Cluster File System 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2925918	Mount point getting hanged after starting async conversion of a ckpt to 'nodata'.

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP2 P2.

Table 1-22   Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 P2
                        fixed issues

Fixed issues	Description
2824895	vcscvmqa "cfsumount" test getting fail.
2745357	Performance enhancements are made for the read/write operation on Veritas File System (VxFS) structural files.
2684573	Enhancement request for force option of the `cfsumount` command.
2674639	VxFS returning error 61493 (VX_EFCLNOSPC) on CFS.
2669724	CFSMountAgent core dump due to assertion failure in `VCSAgThreadTbl::add()`.
2161660	While performing the inode update operations on a cluster mounted file system, the system may panic.

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP2 P1.

Table 1-23   Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 P1
                        fixed issues

Fixed issues	Description
2565400	Poor read performance with DSMC (TSM) backup on CFS filesystems.
2433934	Performance discrepency between CFS and standalone VxFS using NFS.
2630954	The fsck(1M) command exits during pass2.

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

Table 1-24 describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP2.

Table 1-24   Veritas Storage Foundation Cluster File System fixed issues

Fixed issues	Description
2146573	qdetails performance downgraded

Table 1-24   Veritas Storage Foundation Cluster File System fixed issues
(continued)

Fixed issues	Description
2407895	System panic in atomic_add_32_nv because of bad address.

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 1-25   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
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
2180476	System panic in vx_iupdat_clustblks() due to an unprotected inode getting corrupted.
2241123	glmdump failure with the error "/opt/VRTSglm/sbin/glmdump: syntax error at line 340: `count=$' unexpected" is resolved.

Veritas Storage Foundation for Oracle RAC fixed issues

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP4

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP4.

Table 1-26   Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP4 fixed issues

Fixed issues	Description
3090447	The CRSResource agent does not support the C shell (csh) environment.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP3.

Table 1-27   Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2740150	SFRAC CPI does not set OfflineWaitLimit attribute for CSSD agent resource.
2853859	The MultiPrivNIC agent plumbs private network interfaces on the global zone.
2746940	Sometimes during bootup, some of the drivers fail to add in the system because of add_drv/rem_drv race between our modules which are independent of each other.
2850538	System panic in std_devstrat from ODM stack.

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP2 P1.

Table 1-28   Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 P1 fixed issues

Fixed issues	Description
2603511	Database operations can fail on nodes running Oracle RAC 11.2.0.3 and later. The following message is reported in the system logs: `ODM ERROR V-41-4-1-105-22 Invalid argument`

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

Table 1-29 describes the incidents that are fixed in Veritas Storage Foundation for Oracle RAC in 5.1 SP1 RP2.

Table 1-29   Veritas Storage Foundation for Oracle RAC fixed issues

Fixed issues	Description
2374977	Oracle instance crashed; failure occurred at: vcsipc_dosnd
2390892	memory leak in vcsmm_set_cluster_proto
2374970	Update the CRSResource agent to avoid using -n option for srvctl on 11gR2
2429449	The cssd agent explicitly uses hard-coded string "cssd" as resource name.

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 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2, and 5.1 SP1 RP1.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP4

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP4.

Table 1-30   Veritas Cluster Server 5.1 SP1 RP4 fixed issues

Fixed issues	Description
3271592	The Veritas Cluster Server (VCS) logs error messages in the engine log file from the Zone resource.
3266168	During Veritas Cluster Server (VCS) patch upgrade, the file "/opt/VRTSvcs/bin/vcsenv" is overwritten.
3261625	Fail to register a process with Asynchronous Monitoring Framework (AMF) for offline monitoring inside zone.

Table 1-30   Veritas Cluster Server 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
3222541	The amfstat(1M) command prints erroneous column headings in the "Registered Reapers" section.
3207663	Incorrect user privileges are set in case of incorrect use of the '-group' option in command "hauser -addprive.
3162084	If the storage devices of a ZFS storage pool are not in the default directory, Zpool agent fails to import the pool.
3154104	For Application agent, an error message is logged when the StartProgram or StopProgram returns a non-zero value. This gives incorrect implication of the program failure.
3153987	In the Application agent, the clean operation is reported successful even when the CleanProgram returns a non-zero value.
3135560	When a non-global zone is created or attached with the -U option, an error message is displayed.
3106493	Cluster communication fails if some components like Global Atomic Broadcast (GAB) are stopped and quickly restarted.
3090710 2429272	High Availability Daemon (HAD) starts and stops before the VxFEN driver configuration completes.
3079893	Veritas Cluster Server (VCS) fails to bring the service group online when the resource faults.
3076468	Incorrect configuration for one zone results in failure report for other zone resources.
3042545	Support for protocol 30 is required in 5.1SP1RP3 setups to enable RU to 6.0.1.
3042450	A parent service group which is frozen and configured with online local hard dependency is brought offline when its child service group faults.
3028644	The Veritas Cluster Server (VCS) notifier process dumps core if there are any issues with the SNMP configuration.
3013962	Added support for DB2 Version 10.1 in DB2 agent.
3002932	The VxFEN module fails to stop on restarting a node that leads to a failure of removing the keys.

Table 1-30   Veritas Cluster Server 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2999773	When the NFSRestart resource is brought offline, it forcefully stops the automountd process.
2987868	When a resource faults, a service group does not fail over as the TargetCount becomes less than the CurrentCount.
2979745	The MultiNICA agent is unable to detect the loss of network connectivity.
2964772	If you take an NFSRestart resource offline, the NFSRestart agent may unexpectedly stop the NFS processes in a local container (Zones on Solaris or WPARs on AIX).
2941155	Veritas Cluster Server (VCS) does not mark a group as OFFLINE on a failed cluster when a cluster failure is declared in a Global Cluster Option (GCO) environment.
2941092	The DeleteRouteOptions attribute for IPMultiNICB agent is not present in IPMultiNICB agent type definition.
2937673	While executing the amfstat command, the AMF driver causes a panic in the machine.
2848009	Asynchronous Monitoring Framework (AMF) panics the system when an agent is exiting.
2831044	The Sybase agent offline script fails to stop the Sybase processes.
2736627	The remote cluster remains in INIT state, and the Internet Control Message Protocol (ICMP) heartbeat status is UNKNOWN.
2713877	The system panics when an event is registered with Asynchronous Monitoring Framework (AMF).
2650264	The command "hares -display <resname>" fails if a resource is part of a global service group.
2564373	The MonitorMethod attribute does not change from IMF to Traditional, if the AMF module is unconfigured forcefully.
2558232 2558234	The Zone Agent fails to correctly determine the zone state.
2556350 2536404	Veritas Cluster Server (VCS) generates core when the command "hagrp -clear" is executed on a group in "OFFLINE|FAULTED" state.

Table 1-30   Veritas Cluster Server 5.1 SP1 RP4 fixed issues *(continued)*

Fixed issues	Description
2535680 2535690	The Veritas Cluster Server (VCS) engine log is flooded with error messages indicating that imf_getnotification failed.
2486413	Global Atomic Broadcast (GAB) errors are observed in the engine log while running a single node and a standalone Veritas Cluster Server (VCS) cluster where GAB is disabled.
2423680	The Veritas Cluster Server (VCS) commands do not work when VCS object (i.e. Group, Resource, or Cluster) name is G, A, O, E, S, or C.
2398802	VCS should be setting a soft limit for file descriptors in /opt/VRTSvcs/bin/vcsenv
2279845	Veritas Cluster Server (VCS) does not restart the application (configured in a parent service group) running inside the container (configured in a child service group) after the container recovers from a fault.
2210717	When a non-critical resource of a service group faults, the service group remains in the STARTING|PARTIAL state.

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP3 P1.

Table 1-31   Veritas Cluster Server 5.1 SP1 RP3 P1 fixed issues

Fixed issues	Description
2937673	AMF driver panics the machine when amfstat is executed

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP3 P2.

Table 1-32   Veritas Cluster Server 5.1 SP1 RP3 P2 fixed issues

Fixed issues	Description
3002932	vxfen module fails to stop when you manually restart a node by issuing the shutdown -r now command
3042545	Rolling upgrade of Veritas Cluster Server (VCS) from release 5.1SP1RP3 to 6.0 or 6.0.1 might result in a failure.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP3.

Table 1-33   Veritas Cluster Server 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2896402	The resource unregistration gets executed with wrong state when you run the `agrp -online/-offline` or `hares -online/-offline` command.
2869154	Sybase/SybaseBk resource is unable to offline (goes into Unknown state) if Zone state is "down" .
2855755	VxFEN might fail to start or online coordination point replacement (OCPR) might fail if a CP server used as a coordination point for the first time and not reachable that time.
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility hagrp gives incorrect output with the "-clear", "-flush", "-state" options.
2831283	System got panic on GAB with below: panic string: `BAD TRAP: type=31 rp=2a10d4cf530 addr=28 mmu_fsr=0` occurred in module "gab" due to a NULL pointer dereference .
2818567	LLT ARP flood issue.
2804891	lltconfig on boot up core dump and unable to send packets using sendto().
2788059	System did not panic when "PanicSystemOnDGLoss" is set.
2779951	Release reference on event pointer if exec() fails.
2746816	Remove the syslog() call from SIGALRM handler.
2746802	VCS engine should not clear the MigrateQ and TargetCount when failover service group is probed on a system.
2741299	CmdSlave dumped core with SIGSEGV.
2735410	The High Availability Daemon (HAD) core dumps and gets restarted.
2732228	VCS is unable to shut down with the init script.
2731133	When NFSRestart resource is brought offline, it forcefully stops automountd process.
2730979	In IPMP mode, when if_mpadm command to diable interface fails, IPMultiNICB agent may report resource as faulted.

Table 1-33 Veritas Cluster Server 5.1 SP1 RP3 fixed issues *(continued)*

Fixed issues	Description
2729867	Global group did not failover to remote site after HAD gets killed and the primary site node crashed.
2729816	Service group failover failure caused by ToQ not getting cleared when OnlineRetryLimit larger than 0.
2728802	Apache agent should work correctly even if Mountpoint for httpd directory is not present on the failover node.
2714464	Using only spaces in an attribute value may cause issues with the related VCS agent.
2710892	Node is unable to join fencing cluster after reboot, due to snapshot mismatch.
2705391	Enhancement to use different RouteOptions while adding and deleting the route.
2701840	The Netlsnr resource running within a local zone will not go offline when the latter is halted outside of VCS.
2696510	VCS and VxFen may fail to start after reboot.
2692173	The Child service group can be online on the same node with parent group when -nopre is used for an online remote firm dependency.
2689300	Convert 5054 ERROR message to debug log.
2684818	If a pure local attribute like PreOnline is specified before SystemList in main.cf then it gets rejected when HAD is started.
2680428	When you configure an IPMultiNICB resource for a Solaris zone, agent fails to plumb the options.
2660011	Restart of an agent moves a critical resource to FAULTED state and hence the group, even if value of ManageFaults attribute is set to NONE at service group level.
2639181	The clean entry point for Mount agent fails to un-mount a file system of type nfs.
2632569	Unable to stop to load amf driver even if amf in SMF is disabled.
2593173	DiskGroup agent do not detect serial split-brain situation.
2567582	Apache agent should remove the check for Solaris 8 and Solaris 9.

Table 1-33   Veritas Cluster Server 5.1 SP1 RP3 fixed issues *(continued)*

Fixed issues	Description
2564477 2561563	Oracle agent incorrectly reports the global resource as online when the resource inside the local zone is online and the Sid's are same.
2561722 2411882	The imf_register entry point failure count gets incremented even when we imf_unregister entry point fails.
2558988	CurrentLimits not getting updated when a node faults.
2536670	Oracle resource configured inside zone should go in offline state when the Zone is in 'DOWN' state.
2531558	graceful shutdown of node should not trigger race condition on peer.
2526681	The Zone agent and the hazonesetup utility can not use the same VCS user name across multiple zones.
1919382	Mount agent fails to detect the mounted file system with trailing "/".

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP2

Table 1-34 describes the incidents that are fixed in Veritas Cluster Server in 5.1 SP1 RP2.

Table 1-34   Veritas Cluster Server 5.1 SP1 RP2 fixed issues

Fixed Issues	Description
2636874	AMF calls VxFS API with spinlock held.
2423740	Offline registration for a process with AMF succeeds even when the process is running.
2423819	System cannot unload the AMF driver.
2533303	Fixed Oracle agent so that Oracle resource configured inside locale zone goes into offline state when the Zone is in 'DOWN' state.
2528475	Fixed preonline_ipc to support IPMultiNIC/ IPMultiNICB type in preonline_ipc
2528317	Fixed issue with agent framework where Application agent logs the message "Resource(app_res): Output of the completed operation (monitor)" continuously.
2516807	Fixed issue with Application agent to support physical to virtual failover.

Table 1-34   Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2514514	Added support for benchmarking tool ab2 as well as ab to Apache.
2513928	Fixed issue with Mount agent while supporting physical to virtual failover.
2512840	Fixed the issue with Oracle agent to support physical to virtual failover.
2511385	Fixed issue with Sybase where online script marks the database as online before Database has recovered
2508114	Fixed issue with Sybase agents to support physical to virtual failover.
2491635	Fixed issue with DiskGroup which throws ERROR message when 5.1 SF is used.
2485202	Fixed issue with Application agent where monitor fails when 80th character of psargs is whitespace
2483044	Fixed issue with VCS where 'had' crashed with SIGSEGV when asserting against gp->activecount()->gets32GL(nodename) == 0\", in "Resource.C" in check_failover function
2477372	Reduced lltd CPU consumption by reducing the wakeup calls
2477296	Fixed issue with VCS where Application service group did not fail over when the node panic
2477280	Fixed issue with VCS where Application resource failed to failover when system reboot after concurrency violation
2477268	Fixed Zone agent to support physical to virtual failover.
2476897	Fixed issue with IP agent while supporting physical to virtual failover.
2439895	Fixed issue with lltconfig which reports its own cluster node as part of duplicate cluster
2439772	Fixed issue with VCS where wac resource offline failed after network interruption
2439695	Fixed issue with VCS where VXFEN module gets loaded even though user chooses not to enable VXFEN.
2438261	Fixed issue with vxfen where it fails to perform online migration from scsi raw to scsi dmp policy.

Table 1-34 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2434782	Fixed issue with VCS to allow ContainerInfo attribute to be updated even when Group is not completely offline
2433347	Fixed issue with Application Agent where it looks for Home Directory in Global Zone for a particular user even if it configured for local zones.
2427464	Enhanced Multinic Agent to support for dual stack IPv4 and IPv6 (issue with IPv6)
2426663	Fixed issue with vxfen where vxfend does not terminate on OCPR from customized mode to scsi3 mode
2426572	Fixed issue with VCS where persistent resource is reported OFFLINE (not FAULTED) when system is added to the group using hagrp -modify command
2423990	Fixed issue with Application Agent which fails to work correctly when nonexistent user is configured.
2423838	Enhanced IP Agent to log ERROR message for a failure instead of DEBUG message.
2416956	Fixed Application agent not to call StopProgram when CleanProgram is not defined.
2416842	Fixed issue with VCS where had consuming over 99% CPU time. Multiple ha commands are hung in pollsys()
2411860	Fixed issue with VCS where service group switch fails
2411858	Fixed issue with agent framework where various VCS service groups fails to switch
2411653	Added check for MAX message size in GAB
2407872	Fixed shell errors in Application agent "sh: setenv: not found" error during Application online or offline
2407755	Fixed issue with agent framework where async-signal function call removed from signal handler and between fork and exec, which is casuing agents to fail.
2407653	Fixed issue with AMF where in case of forceful unload of AMF module, module reference count for 'vxfs'/'ext3' handled correctly.
2407617	Fixed VCS SMF dependency to be clean after install

Table 1-34 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2407612	Fixed AMF SMF dependency to be clean after install
2406748	Fixed issue with AMF where AMF allowing to register already online process for offline monitor with AMF.
2405391	Fixed LLT to include the nodename in the arp ack packet
2403851	Fixed issue with AMF status where it is showing module loaded but not configured.
2403782	Enhanced Sybase agent scripts to open the file and read instead of using cat command.
2403633	Enhanced agent framework to allow ContainerInfo attribute to be updated even when Group is not completely offline if not already set.
2400485	Fixed issue with vxfen vxfenconfig -c where with mode A has returned EFAULT, all subsequent runs of vxfenconfig -c with mode B fail with error EBADMSG
2400330	Fixed issue with preonline trigger where whyonlining does not behave as documented
2399898	Fixed issue with VCS where hagrp -switch of child group fails if two or more parent groups online on other nodes with 'online-global-soft' dependency.
2398807	Fixed VCS to set soft limits for file descriptors in /opt/VRTSvcs/bin/vcsenv
2394176	Fixed issue with vxfenswap process hangs, "ps -ef" shows "vxfenconfig -o modify" on one node but not on other.
2386326	Fixed issue with vxfen where vxfenadm prints same Serial Number for all LUNs which have more than 96 bytes of SCSI Inquiry data in page 0x83
2382592	Fixed issue with displaying "ResourceInfo" Attribute of SRDF Resource using hares -display
2382583	Fixed issue with CP Agent where it does not show coordination point information in engine log when CP server is not accessible.
2382559	Fixed issue with vxfen where online migration fails while migrating coordination points from CP server to disks.
2382493	Fixed issue with VCS where parent service group does not failover in case of online local firm dependency with child service groups

Table 1-34 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2382463	Fixed issue with VCS where had weight(1) is not added if we reach the boundary condition(10000) in System policy in CPS preffered fencing.
2382460	Fixed issuw with vxfen where the message 'configuring fencing is successful with 3 disks even when single_cp=1'o displays.
2382452	Fixed issue syntax errors with CPS while unconfiguring CP server using configure_cps.pl script
2382384	Fixed VXFEN SMF dependency to be clean after install
2382380	Fixed GAB SMF dependency to be clean after install
2382335	Fixed issue with vxfentsthdw which fails to choose the same fencing disk on two nodes.
2373431	Fixed LLT SMF dependency to be clean after install
2372072	Fixed issue with 'hacf' which dumps core if it could not get current working directory
2371652	Fixed issue with IPMultiNIC agent where it's resources unable to probe when underlying MultiNICA resource is in a global service group
2367721	Fixed owner.vfd file for Oracle to compare only the uid and gid of the user, instead of the complete output string of id command.
2366701	Query regarding usage of variable in VCS attributes
2366201	Enhanced Fencing to start when a majority of the coordination points are available.
2358616	Fixed issue with MultiNICB for nxge interface in base mode when IgnoreLinkStatus is disabled.
2354932	Fixed issue with 'hacli -cmd' which triggers had coredump on a system
2330980	Fixed issue with VCS where HAD sends a resource snapshot to agents running on existing nodes when a node is added to or removed from SystemList.
2330045	Fixed issue with RemoteGroup agent where it's resource during network failure cannot be recovered without bouncing entire cluster.
2330041	Fixed issue with VCS where the group dependencies do not online parallel parent group

Table 1-34   Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2323312	Fixed issue with Application Agent where agent fails to monitor when MonitorProcesses attribute is configured with strings more than 79 characters.
2318334	Fixed issue with Oracle where it needs database's $Oracle_home/lib library to be first in LD_LIBRARY_PATH before /usr/lib
2317067	Fixed issue with Application Agent where the application resource unable to come online inside zones if pid files are specified.
2301731	Fixed panic issue in amf_lock() due to bad mutex during system shutdown.
2298775	Fixed issue with hazonesetup command to set localized zone name for parallel zone service group configuration.
2296172	Fixed issue with VCS where Failover Policy does not take into consideration AutoFailover = 2 and SystemZones when the nodes within a SystemZone are brought down/rebooted.
2276622	Fixed issue with vxfen to configure SCSI-3 fencing using RamSan DMP devices.
2275376	Enhanced Zone agent to used default BootState (multi-user) if not specified.
2271882	Fixed issue with Netlsnr where MonitorMethod attribute does not reflect IMF value without setting Listener attribute on Netlsnr Resource.
2253349	Enhanced IP agent to log a warning message when netmask changed outside of VCS
2393939	Enhanced Apache agent version parsing to accommodate IBM HTTP server 7.0.
1907650 1967408	Coordination Point agent does not provide detailed log message for inaccessible CP servers.

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 1-35 Veritas Cluster Server 5.1 SP1 RP1 fixed issues

Fixed Issues	Description
1949294	`fdsetup` can now correctly parse disk names containing characters such as "-".
1949303	`fdsetup` no longer allows volume that are not part of the RVG, which fixes a possible cause of the RVGSnapshot agent failing.
2011536	Added IMF support for the db2udb agent.
2159991	Fixed an issue with messages in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2172181	Fixed an issue with AMF-related messages for the CAVF agent in the `engine_A.log` file after configuring Veritas Storage Foundation for Oracle RAC on a Japanese language system.
2175599	Fixed an issue with fencing configuration on a 64-node cluster.
2179652	The monitor script of the db2udb agent can now handle empty attribute values.
2184205	Fixed an issue with HAD in which the parent service group did not fail over if the parent service group had an online local firm dependency with a child service group.
2194473	HAD no longer dumps core while overriding the static attribute to the resource level.
2203430	Fixed an issue with haping.
2205556	Fixed an issue with the offline EP of the DNS agent, which did not remove all A/AAAA records if OffDelRR=1 for multi-home records.
2205563	A clean EP now properly removes resource records when OffDelRR=1.
2205567	Fixed an issue in which having an attribute set to `master.vfd` caused the DNS agent to fail to query the DNS server.
2208416	The zone agent now properly starts the network services when the `BootState` attribute is set to multi-user-server.
2208901	Fixed an issue with the RVGSnapshot agent.
2209337	Fixed an issue with VCSAPI where the RemoteGroup agent crashed if the VCSAPI log level was set to a non-zero value.
2210281	Fixed an issue with the agent framework in which MonitorTimeStats incorrectly showed 303 seconds intermittently.

Table 1-35   Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Fixed Issues	Description
2213725	Fixed an issue with the zone agent, in which some of the system services did not start if BootState was configured.
2214539	Fixed an issue in which rebooting a node sometimes set the intentonline of a group to 2, even if the group was online somewhere else. This caused the group to use the autostartlist and not perform a failover.
2218556	Fixed an issue in the `cpsadm` command in which it sometimes failed if LLT was not installed or configured on a single node cluster.
2230869	Fixed an issue with the `IPMultiNICB.xml` file in which the VCS GUI showed the `UseMpathd` attribute for IPMutliNICB resources.
2232633	Fixed an issue with the zone agent in which offline zone resource removed all VCS zone users.
2241419	Fixed an issue in which halogin did not work in a secure environment where the root broker was not a VCS node.

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 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP4

This section describes Veritas Storage Foundation for Databases (SFDB) tools fixed issues in 5.1 SP1 RP4.

Table 1-36 describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 5.1 SP1 RP4.

Table 1-36   Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP4 fixed issues

Fixed Issues	Description
3244283	Sometimes dbed_vmsnap -o snapshot fails on slave node of SFRAC cluster.
3247474	dbed_clonedb operation fails if we have deprecated parameters in the init file.

Table 1-36	Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP4 fixed issues *(continued)*

Fixed Issues	Description
3239626	dbed_vmclonedb fails if the primary database has multiple log_archive_dests set

Veritas Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP3

This section describes Veritas Storage Foundation for Databases (SFDB) tools fixed issues in 5.1 SP1 RP3.

Table 1-37 describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 5.1 SP1 RP3.

Table 1-37	Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP3 fixed issues

Fixed issues	Description
2848204	vxdbd sigsec in strcpy in dbed_ul_print_valist.
2848193	vxdbd coredumps in build_function_header on malloc failure.
2848176	vxdbd memory leak in build_function_header.
1957142	reverse_resync_abort and reverse_resync_commit fails.

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

Table 1-38 describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in 5.1 SP1 RP2.

Table 1-38	Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Fixed issues	Description
2395192	vxdbed looping doing read/write to IDLE sockets

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

There are no SFDB fixed issues in 5.1 SP1 RP1.

Veritas Perl Redistribution fixed issues

This section lists the Veritas Perl Redistribution issues fixed in this release.

Table 1-39   Fixed Issues related to Veritas Perl Redistribution

Incident	Description
3209518	CVE-2011-3597 Perl Digest improper control of generation of code.
3278056	For Veritas Storage Foundation (SF) 5.1, assess PERL vulnerability for CVE-2012-6329.

Known Issues

This section covers the known issues in this release.

- Issues related to installation

- Veritas Dynamic Multi-pathing known issues

- Veritas Storage Foundation known issues

- Veritas Cluster Server known issues

- Veritas Storage Foundation Cluster File System known issues

- Veritas Storage Foundation for Oracle RAC known issues

Issues related to installation

This section describes the known issues during installation and upgrade in this release.

SFRAC version is listed incorrectly when upgrading to 5.1 SP1 RP2 (2530303)

When you upgrade from SFRAC 5.1 SP1 RP1 to SFRAC 5.1 SP1 RP2, the previous SFRAC version is incorrectly listed as 5.1.100.000.

Workaround: This message can be safely ignored.

installrp fails to install 5.1 SP1 RP4 when the root user shell is set to csh (2523643)

The VCS installation fails, if superuser (root) login is using C shell (csh). Currently the installer does not support C shell (`/usr/bin/csh`).

Workaround: Change your superuser (root) shell to `/usr/bin/sh` and retry the installation.

Installation precheck can cause the installer to throw a license package warning (2320279)

If the installation precheck is attempted after another task completes (for example checking the description or requirements) the installer throws the license package warning. The warning reads:

```
VRTSvlic package not installed on system_name
```

Workaround:

The warning is due to a software error and can be safely ignored.

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

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

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

EULA changes (2161557)

The locations for all EULAs have changed.

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

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

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

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

The VRTSacclib package is deprecated (2032052)

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

- Fresh installs: Do not install VRTSacclib.

- Upgrade: Uninstall old VRTSacclib and install new VRTSacclib.

- Uninstall: Ignore VRTSacclib.

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

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

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

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

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

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

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

https://support.oracle.com

Live Upgrade fails when you try to upgrade to Solaris 10 9/10 or later(2168683)

When you try to upgrade to Solaris 10 9/10 or later, Live Upgrade fails. The Live Upgrade command, `luupgrade`, requires the `-k` *auto-registration-file* option, which Symantec's `vxlustart` script does not support.

To resolve this issue

1 Copy the luupgrade command that failed during the execution of the vxlustart command. For example:

```
# luupgrade -u -n dest.18864 \
-s /net/lyptus-new/image/solaris10/update9_GA63521
blocksminiroot filesystem is <lofs>
Mounting miniroot at
  </net/lyptus-new/image/solaris10/update9_GA/Solaris_10/Tools/Boot>
ERROR: The auto registration file <> does not exist or incomplete.
The auto registration file is mandatory for this upgrade.
Use -k <filename> argument along with luupgrade command.
cat: cannot open /tmp/.liveupgrade.11624.24307/.lmz.list
ERROR: vxlustart: Failed: luupgrade -u -n dest.18864
  -s/net/lyptus-new/image/solaris10/update9_GA
```

In this example, you would copy the **luupgrade -u -n dest.18864 -s/net/lyptus-new/image/solaris10/update9_GA** command.

2 Paste the command, and append the command with the -k *auto-registration-file* option. For example:

```
# luupgrade -u -n dest.18864 \
-s /net/lyptus-new/image/solaris10/update9_GA -k /regfile
```

/regfile is absolute path for the auto-registration file.

3 Mount the destination boot environment to /altroot.5.10. Do the following:

■ Display the source and destination boot environment. Enter:

```
# lustatus
```

■ Mount the boot environment. Enter:

```
# lumount destination_boot_environment /altroot.5.10
```

4 After luupgrade completes, install the Storage Foundation packages on the alternate root path.

If you are upgrading from Solaris 2.9 to 2.10, do the following in the order presented:

■ Remove the currently installed Storage Foundation packages. Enter:

```
# uninstallsf -rootpath /altroot.5.10
```

- Upgrade Storage Foundation to 5.1 SP1. Enter:

  ```
  # installsf -rootpath /altroot.5.10
  ```

5 Activate the destination boot environment. Do the following in the order presented:

- Display the source and destination boot environment. Enter:

  ```
  # lustatus
  ```

- Unmount the source and destination boot environment alternate root path. Enter:

  ```
  # luumount destination_boot_environment
  ```

- Activate the destination boot environment. Enter:

  ```
  # luactivate
  ```

6 If the system was encapsulated, manually encapsulate the destination boot environment after it is booted.

During Live Upgrade, installer displays incorrect message about VRTSaa package removal (1710504)

If you use Live Upgrade to upgrade Veritas product 5.0MP1 to Veritas product 5.1 SP1; the installer may display a message that the VRTSaa package failed to uninstall.

The issue is only seen in 5.1 SP1.

Workaround:

Verify whether the VRTSaa package was removed correctly from the alternate boot disk.

```
# pkginfo -R alternate_root_path -l VRTSaa
```

For example, run the following command

```
# pkginfo -R /altroot.5.10 -l VRTSaa
```

If the VRTSaa package was removed, you can ignore this error.

If the VRTSaa package was not removed, remove the package manually:

```
# pkgrm -R alternate_root_path VRTSaa
```

For example, run the following command

```
# pkgrm -R /altroot.5.10 VRTSaa
```

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

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

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

  ```
  # bootadm update-archive
  ```

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

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

See the `bootadm` manual page for Solaris.

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

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

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

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

Workaround: None.

Errors observed during partial upgrade of SFHA

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

The issue is only seen in 5.1 SP1.

Workaround

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

Manual upgrade of VRTSvlic package loses keyless product levels (2115662)

If you upgrade the `VRTSvlic` package manually, the product levels that were set using `vxkeyless` may be lost. The output of the `vxkeyless display` command will not display correctly.

To prevent this, perform the following steps while manually upgrading the VRTSvlic package.

To manually upgrade the VRTSvlic package

1.   Note down the list of products configured on the node for keyless licensing.

    ```
    # vxkeyless display
    ```

2.   Set the product level to NONE.

    ```
    # vxkeyless set NONE
    ```

3.   Upgrade the VRTSvlic package.

    ```
    # pkgrm VRTSvlic
    ```

    This step may report a dependency, which can be safely overridden.

    ```
    # pkgadd -d VRTSvlic.pkg
    ```

4.   Restore the list of products that you noted in step 1.

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

Issues with keyless licensing reminders after upgrading VRTSvlic (2141446)

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

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

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.

2. Set the product level to *NONE* with the command:

   ```
   # vxkeyless set NONE
   ```

3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic`:

   - Verify if the key has `VXKEYLESS` feature Enabled using the following command:

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

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

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

   ---

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

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

After Live Upgrade to Solaris 10 Update 10, boot from alternate boot environment may fail (2922557)

If your setup involves volumes in a shared disk group that are mounted as CFS in a cluster, then during Live Upgrade using the `vxlustart` command from any supported Solaris version to Solaris 10 Update 10, boot from an alternate boot environment may fail.

Workaround:

Run the `vxlufinish` command. Before rebooting the system, manually delete the entries of all the volumes of shared disks that are mounted as CFS in the `/altroot.5.10/etc/vfstab` directory.

After VCS 5.1 SP1 RP3 upgrades to 5.1 SP1 RP4, if a non-global zone is created or attached with the `-U` option, an error message is displayed (3135560)

After VCS 5.1 SP1 RP3 upgrades to 5.1 SP1 RP4, if a non-global zone is created or attached with the `-U` option, the following error message is displayed:

```
Installation of these packages generated errors: <VRTSllt VRTSgab
VRTSvxfen>
```

Workaround:

Since VRTSllt, VRTSgab, or VRTSvxfen installs only metadata inside non-global zones, this error can be safely ignored. This issue is not present on other supported upgrade paths.

5.1 SP1 RP4 configuration title shows as "Veritas Storage Foundation for Oracle RAC 5.1 SP1 PR2 Configure Program" (2908221)

The installer scripts under `/opt/VRTS/install` is still using base version and the configuration title still shows base version after you install 5.1 SP1 RPx patches.

Workaround

There is no workaround. The installer won't be updated when installing RP patches.

Veritas Dynamic Multi-pathing known issues

This section describes the Veritas Dynamic Multi-pathing known issues in this release.

The vxdmpasm command can not work on Solaris platform (2930396)

The `vxdmpasm` command does not work.

For example:

```
vxdmpasm enable user1 group1 600 emc0_02c8
expr: syntax error
/etc/vx/bin/vxdmpasm: test: argument expected
```

Workaround:

Use the following commands to set the ownership and permissions manually after each boot:

For example, for emc0_02c8 disk:

```
# chown user1:group1 /dev/vx/rdmp/emc0_02c8*
# chmod 0600 /dev/vx/rdmp/emc0_02c8*
```

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 available for I/O. The restore daemon probes the paths at the interval set with the tunable parameter dmp_restore_interval. If you set the dmp_restore_interval parameter to a high value, the paths are not available for I/O until the next interval.

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

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

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

Workaround:

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

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

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

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

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

Workaround:

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

To set the OnlineTimeout attribute for the HA resource type CVMCluster

1   Make the VCS configuration to be read/write:

    # **haconf -makerw**

2   Change the OnlineTimeout attribute value of CVMCluster:

    # **hatype -modify CVMCluster OnlineTimeout 600**

3   Display the current value of OnlineTimeout attribute of CVMCluster:

    # **hatype -display CVMCluster -attribute OnlineTimeout**

4   Save and close the VCS configuration:

    # **haconf -dump -makero**

Adding a DMP device or its OS device path as a foreign disk is not supported (2062230)

When DMP native support is enable, adding a DMP device or its OS device path as a foreign disk using the vxddladm addforeign command is not supported. Using this command can lead to unexplained behaviour.

ZFS pool creation on a DMP device fails when the LUN size is between 1 TB and 2TB (2010919)

Creating a ZFS pool on a DMP device using the whole disk of size > 1TB and < 2TB that contains a SMI SUN label fails. The issue is that zpool create on a whole disk changes the device label from SMI to EFI. This causes confusion between the OS device paths of the same DMP device due to a bug in the Sun SCSI layer. This is due to SUN BugID: 6912703.

The modified dmp tunable parameter 'dmp_monitor_osevent' doesn't persist the same value after upgrade to 5.1 SP1 RP4 and reboot (2936349)

If you modify the dmp tunable parameter dmp_monitor_osevent using the vxdmpadm settune command before upgrading to 5.1 SP1 RP4, the tunable value is not persistent after upgrade to 5.1 SP1 RP4.

Workaround:

There is no workaround for this issue.

The vxtune output value is not persistent across reboot and upgrades (2935245)

The `vxtune` output value is not persistent across reboots and upgrades. Even if all of the tunable parameter are changed, it's value will be changed back to previous value after system reboot or product upgrade.

Workaround:

Change the `vxtune` output value again after system reboot or product upgrade.

Veritas Storage Foundation known issues

This section describes the Veritas Storage Foundation known issues in this release.

- Veritas Storage Foundation known issues
- Veritas Volume Manager known issues
- Veritas File System known issues
- Veritas Volume Replicator known issues
- Veritas Storage Foundation for Databases (SFDB) tools known issues

vxlufinish fails on Solaris x86 during upgrade to Solaris 10 update 11 due to luumount or luactivate failure (2939321)

The `luumount` command and the `luactivate` command are internally used by vxlufinish for Live upgrade on Solaris. Behavior of these commands have changed in Solaris 10 update 11, resulting in Live Upgrade failure with following error message:

```
Generating boot-sign for ABE <dest.num3>
/bin/rmdir: directory "/tmp/.liveupgrade.num1.num2/.alt.luactivate": \
Directory is a mount point or in use Generating partition and slice \
information for ABE <dest.num3>
/tmp/.liveupgrade.num1.num2/.alt.luactivate
ERROR: The target boot environment <dest.num3> root device \
<device-path> is already mounted.
ERROR: The root slice <> of the target boot environment <dest.num3>\
 is not available.
rm: Unable to remove directory /tmp/.liveupgrade.num1.num2/.\
alt.luactivate:
Device busy
rm: Unable to remove directory /tmp/.liveupgrade.num1.num2: File exists
ERROR: vxlufinish Failed: luactivate dest.num3
```

Workaround:

This is an issue with Oracle Solaris 10 update 11 (x86). Symantec is following up with Oracle to resolve it.

Veritas Storage Foundation known issues

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

Adding a node fails when using the Web-based installer (2173672)

When you add a node using the Web-based installer you cannot proceed beyond starting GAB on new node if the cluster uses secure CPS.

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.

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.

An error message is received 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.

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

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

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

Live Upgrade may fail on Solaris 2.10 x86 (1984664)

This Live Upgrade issue is specific to Solaris 2.10 x86 operating system. If you run the VxVM `vxlustart` script and choose a disk for the destination boot environment (BE) that is missing the `fdisk` "SOLARIS System" partition, the Live Upgrade operation fails to properly setup the destination BE.

Workaround: Run the `fdisk` command on the destination BE disk before performing the Live Upgrade operation. If the fdisk table is missing or does not contain a "SOLARIS System" partition, it must be setup.

To setup the fdisk table

◆ Run the `fdisk` command on the root of the destination BE:

    # **fdisk /dev/rdsk/c1t1d0p0**
    No fdisk table exists. The default partition for the disk is:

      a 100% "SOLARIS System" partition

    Type "y" to accept the default partition,
    otherwise type "n" to edit the partition table.
    Please answer with "y" or "n": **y**

On Solaris x86 the disks must be initialized with a fdisk table and Solaris system partition before it can be used.

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

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

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

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

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

For example:

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

        stale //kernel/drv/sparcv9/vxportal
        stale //kernel/drv/vxportal.conf
        stale //kernel/fs/sparcv9/vxfs
        ...
        new   /kernel/drv/vxlo.SunOS_5.10
        new   /kernel/drv/vxlo.conf

```
changed /kernel/drv/vxspec.SunOS_5.9
changed /kernel/drv/vxspec.conf
```

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

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.

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

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

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

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

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

Veritas Volume Manager known issues

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

vxcdsconvert(1M) doesn't work on disk groups with thin Luns. (2747048)

vxcdsconvert(1M) converts thin Luns to thick Luns, so it doesn't work on disk groups with thin Luns.

Workaround:

There is no workaround for this issue.

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

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

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

Work-around:

To resolve the issue, perform the following steps.

To recover from the error

1   Stop the I/O.

2   Reboot the system with the following command:

```
# reboot -- r
```

3   Retry the operation.

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

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

Workaround: There is no workaround for this issue.

Node join can lead to hang if an upgrade of the cluster protocol version is in progress (2103567)

If you attempt to join a node to the cluster while Cluster Volume Manager (CVM) is upgrading the cluster protocol version, the system may hang. This issue occurs if the node is attempting to join the cluster after you issue the `vxdctl upgrade` command to upgrade the CVM cluster.

Work-around:

Avoid joining a new node to the cluster until the CVM cluster upgrade is completed.

Issues when removing the VxVM 5.1SP1 patch

If you back out the VxVM 5.1SP1 patch, note the following issues:

■ VxVM 5.1SP1 introduces a new disk group version. If you upgrade disk groups or create new disk groups with the new disk group version, VxVM 5.1 cannot access these disk groups if you back out the 5.1SP1 patch.

■ When you back out the VxVM 5.1SP1 patch, VxVM recreates the `/etc/vx/volboot` file with the default contents. If the existing VxVM 5.1 configuration had a modified `/etc/vx/volboot` file, these modifications are lost when you back out the patch.

■ The VxVM 5.1SP1 patch changes certain files in the `/etc/vx` directory, such as the `dmppolicy.info` file. After backing out the patch, the VxVM 5.1 may have issues in parsing these files.

Workaround:

Do not upgrade disk groups to version 160 until you are sure you do not need to back out the patch.

Veritas Volume Manager (VxVM) might report false serial split brain under certain scenarios (1834513)

VxVM might detect and report a false serial split brain when all of the following conditions are met:

■ One or more arrays that provide the shared storage for the cluster are being powered off

■ At the same time when the arrays are being powered off, an operation that requires an internal transaction is initiated (such as VxVM configuration commands)

In such a scenario, disk group import will fail with a split brain error and the vxsplitlines output will show 0 or 1 pools.

Workaround:

To recover from this situation

1   Retrieve the disk media identifier (dm_id) from the configuration copy:

> # **/etc/vx/diag.d/vxprivutil dumpconfig *device-path***

The dm_id is also the serial split brain id (ssbid)

2   Use the dm_id in the following command to recover from the situation:

> # **/etc/vx/diag.d/vxprivutil set *device-path* ssbid=*dm_id***

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

The attempted relayout operation on a disk group containing approximately more than 300 LUNs or disks may fail with the following error:

```
Cannot setup space
```

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

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

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

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

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

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

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

Work-around:

To resolve the issue, perform the following steps.

To recover from the error

1 Stop the I/O.

2 Reboot the system with the following command:

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

3 Retry the operation.

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

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

Workaround: There is no workaround for this issue.

Co-existence check might fail for CDS disks

In Veritas Volume Manager (VxVM) 5.1 SP1, 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.

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.

Node not able to join when recovery is in progress (2165829)

When a node leaves a cluster, there is an associated volume recovery for the leaving node. It is not allowed to join any node to the cluster during the recovery because the nodeid of the leaving node cannot be assigned to the joining node.

Workaround:

Retry node join after the recovery is completed.

Solaris machine panics while resizing CDS-EFI LUN (2915751)

While resizing a disk having CDS-EFI format or while resizing a disk from less than 1TB to >= 1TB, machine panics because of the incorrect use of device numbers. VxVM uses the whole slice number s0 instead of s7 which represents the whole device for EFI format. Hence, the device open failes and the incorrect disk maxiosize was populated. While doing an I/O, machine panics with divide by zero error.

Wokaround:

There is no workaround for this issue. You need to update to the VxVM version which has the available fix.

If vxconfigd is under heavy load, "vxassist settag" may make volume tagging information inconsistent (2484764)

If there are a lot of VxVM operations running, vxconfigd is under heavy load. If you execute the `vxassist settag` operations when vxconfigd is under stress, these operations will succeed, but the volume tagging information may be inconsistent. In such cases, you will not be able to use the tag for the further operations for that particular volume. And if you run the `vxassist listtag` operation, it will fail with error:

```
Inconsistent tag information found on disk
```

Workaround:

There is no workaround for this issue.

vxconfigd dumps core on all the nodes in Campus Cluster setup (2937600)

Campus Cluster Scenario (two sites A and B, with 2 nodes in each site):

1. Disabled site A storage from all the four nodes and also shutdown site A nodes.

2. Enabled site A storage and activated site A nodes.

3. Site B nodes panic.

After the reboot of the nodes in site A, when nodes try to join the cluster, vxconfigd dumps core.

Workaround:

There is no workaround for this issue right now.

Running Vxtrace against a volume shows response times as negative (1952197)

The vxtrace(1M) command displays negative or improper values in the output. For example:

```
# vxtrace -g <dgname> -o dev,disk,vol,plex,sd <volume name>
143 END write sd <subdisk> op 140 block 33872 len 16 time -5325
142 END write plex <plex> op 140 block 33872 len 16 time -5325
141 END write vol <volume> op 140 block 33872 len 16 time -5325
140 END write vdev <volume> op 140 block 33872 len 16 time 0
```

Workaround: There is no workaround for this issue.

Disk greater than 1TB goes into error state (3269099)

For a device larger than 1TB with multiple paths, if a path is labeled in EFI format by OS command such as format(1M) , vxdisk list command shows it as error in the output.

Workaround: There is no workaround for this issue because it is a Sun OS problem.

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 32 MB, then the public region data will be overridden.

Workaround:

A new "-o" option has been introduced since 5.1 SP1 RP3 P2.

The "-o retain" option for disk initialization keeps existing private or public offsets intact. To use the command, enter:

```
# vxdisk -f -o retain init da-name
```

Veritas File System known issues

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

File System hangs, when it has reached the space limit (3253210)

File System hangs and displays the following stack trace:

```
vx_svar_sleep_unlock
default_wake_function
__wake_up
```

```
vx_event_wait
vx_extentalloc_handoff
vx_te_bmap_alloc
vx_bmap_alloc_typed
vx_bmap_alloc
vx_bmap
vx_exh_allocblk
vx_exh_splitbucket
vx_exh_split
vx_dopreamble
vx_rename_tran
vx_pd_rename
```

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.

Truncate operation of a file with a shared extent in the presence of a Storage Checkpoint containing FileSnaps results in an error (2149659)

This issue occurs when Storage Checkpoints are created in the presence of FileSnaps or space optimized copies, and one of the following conditions is also true:

- In certain cases, if a FileSnap is truncated in the presence of a Storage Checkpoint, the `i_nblocks` field of the inode, which tracks the total number of blocks used by the file, can be miscalculated, resulting in inode being marked bad on the disk.

- In certain cases, when more than one FileSnap is truncated simultaneously in the presence of a Storage Checkpoint, the file system can end up in a deadlock state.

This issue causes the following error to display:

```
f:xted_validate_cuttran:10 or f:vx_te_mklbtran:1b
```

Workaround: In the first case, run a full `fsck` to correct the inode. In the second case, restart the node that is mounting the file system that has this deadlock.

When online migration is in progress, df command with no mount point or device argument fails with error 1 (2162822)

When online migration is in progress, the `df` command with no mount point or device argument fails with error 1.

The `df` comand also gives an error for every file system undergoing migration. The error is similar to the following example:

```
df: cannot statvfs /mntpt/lost+found/file_system:
  No such file or directory
```

Workaround: To avoid the error, specify a mount point or device.

Tunable not enabling the lazy copy-on-write optimization for FileSnaps (2164580)

The lazy copy-on-write tunable does not enable the lazy copy-on-write optimization for FileSnaps.

Workaround: There is no workaround for this issue.

vxfilesnap fails to create the snapshot file when invoked with the following parameters: vxfilesnap source_file target_dir (2164744)

The `vxfilesnap` command fails to create the snapshot file when invoked with the following parameters:

```
# vxfilesnap source_file target_dir
```

Invoking the `vxfilesnap` command in this manner is supposed to create the snapshot with the same filename as the source file inside of the target directory.

Workaround: You must specify the source file name along with the target directory, as follows:

```
# vxfilesnap source_file target_dir/source_file
```

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

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

Workaround: There is no workaround for this issue.

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

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

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

Workaround: You must reboot after the upgrade completes.

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 several minutes to recover from this state.

Workaround: There is no workaround for this issue.

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

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

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

Installing the VRTSvxfs 5.1 RP2 package on non-global zones can fail (2086894)

Installing the VRTSvxfs 5.1 RP2 patch on non-global zones can fail with the following error messages:

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

Workaround: The following procedure installs the VRTSvxfs 5.1 RP2 package on a non-global zone.

To install VRTSvxfs 5.1 RP2 on a non-global zone

1 Remove the `VRTSvxfs` 5.1 RP1 package.

2 Reinstall the `VRTSvxfs` 5.1 package.

3 Install the `VRTSvxfs` 5.1 RP2 package.

Storage Checkpoints can exceed the quota limit (2102201)

Under some circumstances, Storage Checkpoints can exceed the quota limit set by the `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 fails with the "Out of Buffer cache" error.

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

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

Workaround: There is no workaround for this issue.

Possible write performance degradation with VxFS local mounts (1837394)

Some applications that allocate large files without explicit preallocation may exhibit reduced performance with the VxFS 5.1 release and later releases 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.

Workaround: 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 applications 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 that were shut down earlier.

Installing the VRTSvxfs 5.1 SP1 RP2 package on non-global zones can fail (2406662)

Installing the VRTSvxfs 5.1 SP1 RP2 patch over 5.1 SP1 RP1 on non-global zones can fail with the following error messages:

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

Workaround: The following procedure installs the VRTSvxfs5.1 SP1 RP2 over VRTSvxfs 5.1 SP1 RP1 on a non-global zone.

To install VRTSvxfs 5.1 SP1 RP2 on a non-global zone

1   Remove the `VRTSvxfs` 5.1 SP1 RP1 package.

2   Reinstall the `VRTSvxfs` 5.1 SP1 package.

3   Install the `VRTSvxfs` 5.1 SP1 RP2 package.

Unable to grow the File System on svm volume (2340820 (2213282))

This issue will be seen only on Solaris10 update8 > > (sol10u8) and above.

If the underlying volume is Solaris Volume Manger (SVM) and the size of the volume is > 1 TB, then operations like resize of file system through fsadm will fail with the following error:

```
UX:vxfs fsadm: ERROR: V-3-20058: read_vtoc failed with return value -7
```

Workaround:

A ONEOFF patch 5.1SP1RP2ONEOFF[Sol_sparc 2.10=> 146907-70, Sol_x86 2.10=> 146908-70] is available which contains the fix for this issue.

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

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

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

Workaround: You must reboot after the upgrade completes.

NFS issues with VxFS checkpoint (2027492)

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 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 in the VCS engine log:

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

Workaround:

To resolve this issue

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

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

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

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

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

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

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

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

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround: The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

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

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

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:

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.

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.

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

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

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

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

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

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

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

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

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

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

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

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

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.

Live Upgrade fails when you try to upgrade to Solaris 10 9/10 or later

When you try to upgrade to Solaris 10 9/10 or later, Live Upgrade fails. The Live Upgrade command, `luupgrade`, requires the `-k auto-registration-file` option, which Symantec's `vxlustart` script does not support.

To resolve this issue

1   Copy the `luupgrade` command that failed during the execution of the `vxlustart`
    command. For example:

```
# luupgrade -u -n dest.18864 \
-s /net/lyptus-new/image/solaris10/update9_GA63521
blocksminiroot filesystem is <lofs>
Mounting miniroot at
  </net/lyptus-new/image/solaris10/update9_GA/Solaris_10/Tools/Boot>
ERROR: The auto registration file <> does not exist or incomplete.
The auto registration file is mandatory for this upgrade.
Use -k <filename> argument along with luupgrade command.
cat: cannot open /tmp/.liveupgrade.11624.24307/.lmz.list
ERROR: vxlustart: Failed: luupgrade -u -n dest.18864
  -s/net/lyptus-new/image/solaris10/update9_GA
```

In this example, you would copy the **luupgrade -u -n dest.18864
-s/net/lyptus-new/image/solaris10/update9_GA** command.

2   Paste the command, and append the command with the `-k`
    auto-registration-file option. For example:

```
# luupgrade -u -n dest.18864 \
-s /net/lyptus-new/image/solaris10/update9_GA -k /regfile
```

`/regfile` is absolute path for the auto-registration file.

3   Mount the destination boot environment to `/altroot.5.10`. Do the following:

    ■   Display the source and destination boot environment. Enter:

        ```
        # lustatus
        ```

    ■   Mount the boot environment. Enter:

        ```
        # lumount destination_boot_environment /altroot.5.1
        ```

4   After `luupgrade` completes and after mounting alternate boot environment,
    upgrade the Storage Foundation packages on the alternate root path using
    the following command:

```
# installsf -rootpath /altroot.5.10 -upgrade
```

If you are upgrading from Solaris 9 to 10, do the following in the order presented:

    ■   Remove the currently installed Storage Foundation packages. Enter:

```
# uninstallsf -rootpath /altroot.5.10
```

- Upgrade Storage Foundation to 5.1 SP1 RP4. Enter:

```
# installsf -rootpath /altroot.5.1
```

5   Activate the destination boot environment. Do the following in the order
     presented:

- Display the source and destination boot environment. Enter:

```
#  lustatus
```

- Unmount the source and destination boot environment alternate root path.
  Enter:

```
# luumount destination_boot_environment
```

- Activate the destination boot environment. Enter:

```
# luactivate
```

6   If the system was encapsulated, manually encapsulate the destination boot
     environment after it is booted.

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 (2184482)

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:

When upgrading from Veritas product version 5.0 to Veritas product 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 error message:

```
/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.
```

Workaround

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

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.

The the /opt/VRTSdbed/bin/dbdst_obj_move command may fail with error messages on 10gRAC env (2927308)

The `dbdst_obj_move` command may fail with FSPPADM error:

```
/opt/VRTS/bin/dbdst_obj_move -S $ORACLE_SID -H $ORACLE_HOME \
-v -t tab_part4 -s 0 -e 10 -c SLOW
FSPPADM err : Not enough space

SFORA dst_obj_adm ERROR V-81-6414 Internal Error at fsppadm_err
```

This error can be caused by the old filesystem layout version. To use the
dbdst_obj_move command, you need filesystem layout 8 or higher.

Workaround:

Upgrade the filesystem layout to version 8.

To upgrade the filesystem layout to version 8:

1   Use the following command to check the filesystem layout version:

    # **/opt/VRTS/bin/fstyp -v /dev/vx/dsk/oradatadg/oradatavol \
    | grep version**

2   Use the following command to upgrade the filesystem layout to version 8:

    # **/opt/VRTS/bin/vxupgrade -n 8 /oradata**

The dbed_vmclonedb -o recoverdb command may fail with error messages (2928666)

The dbed_vmclonedb -o recoverdb command may fail with following error
messages:

```
SFORA dbed_vmclonedb ERROR V-81-4882 An error occurred while reconfiguring
Oracle instance 'ckptcln'.
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is
at /tmp/dbed_vmclonedb.50528402/startup.log.
```

Also check the startup.log file if it contains the following information:

```
./home/oracle>cat /tmp/dbed_vmclonedb.47251692/startup.log
ORA-16019: cannot use LOG_ARCHIVE_DEST_1 with LOG_ARCHIVE_DEST or
LOG_ARCHIVE_DUPLEX_DEST
```

The error occurs if the log archive destination parameter of Oracle RAC instances
is configured individually using a similar command such as: :

```
SQL> alter system set log_archive_dest_1 = 'location=/arch MANDATORY' SID='RACID';
```

Workaround:

Use the following command to set log_archive_dest_1 for all of the instances at once:

```
SQL> alter system set log_archive_dest_1 = 'location=/arch MANDATORY' SID=*;
```

The `dbed_vmclonedb` command fails if the server_name argument has capitalized characters (3278357)

The `dbed_vmclonedb` command fails if the server_name argument has capitalized characters. You may receive similar error messages:

```
dbed_vmclonedb -S sid1 -o recoverdb, new_sid=newsid, server_name=
HOSTNAME -f snap1 -r /tmp/mnt
SFORA vxreptadm ERROR V-81-6512 Could not find database 'sid1' in
SFORA vxreptadm ERROR V-81-8851 Query key 27 failed, attempting rollback.
SFORA dbed_vmclonedb ERROR V-81-8842 Could not query repository.
```

Workaround: Make sure the server_name argument is in lower cases.

The `dbdst_obj_move` command fails with an error message for database table spread across multiple mounts (3260289)

The `dbdst_obj_move` command fails with the following error:

```
bash-2.05b$ dbdst_obj_move -S sdb -H $ORACLE_HOME -t test3 -c MEDIUM
FSPPADM err : UX:vxfs fsppadm: WARNING: V-3-26543: File handling failure
on /snap_datadb/test03.dbf with message -
SFORA dst_obj_adm ERROR V-81-6414 Internal Error at fsppadm_err
```

This issue is observed when a database table is spread across multiple mount-points (this issue can be ascertained using `dbdst_obj_view` command),and the user tries to move all extents of the table in one attempt.

Workaround: Specify a range of extents that belong to a common mount-point.The`dbdst_obj_move` command has an option to specify the range of extents. If the database table is spread across "n" mount-points, then you must run `dbdst_obj_move` "n" times with a different range of extents each time.

`dbed_vmclonedb -o umount` fails with an error message (3265937)

The `dbed_vmclonedb -o umount` command fails with the following error:

```
...
ERROR at line 1:
ORA-01187: cannot read from file  because it failed verification tests
ORA-01110: data file 201: '/data/temp01.dbf'
SFORA dbed_vmclonedb ERROR V-81-3228 Could not backup current control file.
```

This error may occur when you do the steps in the following sequence:

1   Clone and unmount a snapshot.

2   Start `reverse_resync` with `dbed_vmsnap -o reverse_resync_begin`.

3   Abort `reverse_resync` with `dbed_vmsnap -o reverse_resync_abort`.

4   Mount and unmounts the snapshot.

Workaround:

To solve this issue:

1   Mount the snapshot and start the clone database

2   Recreate the temp file on the clone database

    SQL > **alter database tempfile '/data/temp01.dbf' drop;**

    SQL > **alter tablespace temp add tempfile '/data/temp01.dbf';**

Clone recovery fails for RAC database (3271827)

During a RAC database clone, recovery may fail with error messages that are similar to:

```
$ dbed_vmclonedb -S racdb1  -f DDD -o
recoverdb,new_sid=smtfnew,server_name=l111027, -r /tmp/DDD
====================================================
dbed_vmclonedb started at 2013-07-03 23:32:12
Editing remote_login_passwordfile in initsmtfnew.ora.
All redo-log files found.
Altering instance_name parameter in initsmtfnew.ora.
Altering instance_number parameter in initsmtfnew.ora.
Altering thread parameter in initsmtfnew.ora.
Starting automatic database recovery.
SFORA dbed_vmclonedb ERROR V-81-4918 Database smtfnew has not been correctly
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is at /tmp/recover.log.7679.
oracle@l111027:~> cat /tmp/recover.log.7679
ALTER DATABASE OPEN RESETLOGS
*
ERROR at line 1:
ORA-01195: online backup of file 1 needs more recovery to be consistent
ORA-01110: data file 1: '/tmp/DDD/oradata/racdb/system01.dbf'
```

The automatic recovery applies log files to recover the database, which may adopt an incorrect order. The wrong order fails the recovery, but leaves the database in a mounted state.

Workaround: Manually apply the necessary (archived) logs to recover the database to the desired point of time, and open the database with the "RESETLOGS" option.

Veritas Storage Foundation and High Availability known issues

This section describes the Veritas Storage Foundation and High Availability known issues in this release.

Use of Live Partition Mobility on an SFHA or SFCFSHA node with SCSI-3 fencing enabled for data disks causes service groups on that node to fail (2619600)

After you execute Live Partition Mobility (LPM) on an SFHA or SFCFSHA node with SCSI-3 fencing enabled for data disks, I/O fails on devices or disks with reservation conflicts. Reservation conflicts cause the associated service groups on the node to fail. Hence, the service groups failover to other available nodes.

Workaround:

After LPM completes migration for the node, manually make the service groups online on that node.

Veritas Cluster Server known issues

This section describes the Veritas Cluster Server known issues in 5.1 SP1 RP4, 5.1 SP1 RP3, 5.1 SP1 RP2, 5.1 SP1 RP1 and 5.1 SP1.

- Operational issues for VCS
- Issues related to the VCS engine
- Issues related to the bundled agents
- Issues related to the VCS database agents
- Issues related to the agent framework
- Issues related to Live Upgrade
- Issues related to VCS in Japanese locales
- Issues related to global clusters
- Issues related to LLT
- Issues related to I/O fencing
- Issues related to Symantec Product Authentication Service with VCS
- Issues related to Veritas Cluster Server agents for Veritas Volume Replicator

- Issues related to IMF

- Issues related to AMF

NFS cluster I/O fails when storage is disabled

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

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

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

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.

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.

pkgchk returns errors after installing VCS patches (2556541)

The `pkgchk` command returns errors for VCS patches as some of the files get modified during VCS configuration.

- Package: VRTSllt

```
# pkgchk -n VRTSllt
ERROR: /etc/default/llt
permissions <0744> expected <0644> actual
group name <sys> expected <root> actual
modtime <02/05/09 11:29:32 AM> expected <08/29/11 09:18:57 PM> actual
```

- Package: VRTSgab

```
# pkgchk -n VRTSgab
ERROR: /etc/default/gab
permissions <0744> expected <0644> actual
group name <sys> expected <root> actual
modtime <02/05/09 11:30:37 AM> expected <08/29/11 09:18:59 PM> actual
```

- Package: VRTSvxfen

  ```
  # pkgchk -n VRTSvxfen
  ERROR: /etc/default/vxfen
  permissions <0544> expected <0644> actual
  group name <sys> expected <root> actual
  modtime <02/05/09 11:28:11 AM> expected <08/30/11 04:52:11 AM> actual
  file cksum <21904> expected <21902> actual
  ```

- Package : VRTSvcs

  ```
  # pkgchk -n VRTSvcs
  ERROR: /etc/default/vcs
  permissions <0744> expected <0644> actual
  group name <sys> expected <root> actual
  modtime <02/05/09 03:01:55 PM> expected <08/29/11 09:19:01 PM> actual
  ```

- Package: pkgchk -n VRTSvcsea

  ```
  # pkgchk -n VRTSvcsea
  ERROR: /etc/VRTSvcs/conf/config/Db2udbTypes.cf
  File type <s> expected <f> actual
  ```

Workaround: There is no workaround for this issue. These errors are safe to ignore.

halting zone outside of VCS takes zone into down state (2557287)

If zone is halted outside of VCS control, zone goes into down state.

Workaround:

When zone goes to down state, an administrator intervention is required.

Unmount any mounted file systems inside of zone and bring the zone to steady state (configured/installed/running) using the **zoneadm –z *zonename* halt** command outside of VCS.

the VRTSllt, VRTSgab and VRTSvxfen patches fails to install in non global zone (2562424)

The VRTSllt, VRTSgab and VRTSvxfen patches fails to install during the Live Upgrade from 5.1SP1RP1 to 5.1SP1RP2 on the system where non global zone is running.

Workaround: This message is safe to ignore.

warning for patch "143290-05" during the Live Upgrade from 5.1SP1RP1 to 5.1SP1RP3 (2562430)

If the local zone is not in running state, the following message displays for patch "143290-05" during the Live Upgrade from 5.1SP1RP1 to 5.1SP1RP3:

```
Executing postpatch script...

WARNING: Unable to find bmcmap
```

Workaround: This message is safe to ignore.

hacf -verify displays a syntax error if the main.cf file is empty and VCS goes into STALE_ADMIN_WAIT state (2245296)

hacf -verify displays a syntax error because the command fails to parse an empty main.cf file. As hacf -verify fails to parse the main.cf file, VCS is unable to load the configuration and goes into STALE_ADMIN_WAIT state.

Workaround:

Include the appropriate type files (for example, types.cf) from the main.cf file and load the configuration by executing the hasys -force command.

IP resource does not go into UNKNOWN state when the Device attribute is configured with an invalid value (3267621)

IP resource does not go into UNKNOWN state when the Device attribute is configured with an invalid value. For example, interface name in uppercase or non-existent interface name.

Workaround:

For proper functioning of the IP resource, configure the Device attribute as per the interface name. You must use the interface name displayed in the OS command output.

NFS client reports I/O error because of network split-brain (3257399)

When network split-brain occurs, the failing node may take some time to panic. Thus, the service group on the failover node may fail to come online, as some of the resources (like IP resource) are still online on the failing node or disk group on the failing node may get disabled, but the IP resource on the same node continues to be online.

Workaround:

Configure the preonline trigger for the service group containing the DiskGroup resource on each system in the service group:

1. Enter the following command and copy the preonline_ipc trigger from /opt/VRTSvcs/bin/sample_triggers/VRTSvcs to /opt/VRTSvcs/bin/triggers/. Note that you must copy the preonline_ipc trigger as preonline.

```
# cp
/opt/VRTSvcs/bin/sample_triggers/VRTSvcs/preonline_ipc/opt/VRTSvcs/bin/triggers/preonline
```

2. Enter the following command and enable PreOnline trigger for the service group:

```
# hagrp -modify group_name PreOnline 1
```

Multiple CFSmount resources are in a single service group they may not all come online after a reboot (2164670)

In some cases when multiple CFSmount resources are in a single service group they may not all come online after a reboot. You will need to manually bring them online after a reboot.

Workaround

Create a resource dependency between the various CFSmount resources.

Operational issues for VCS

Issues with Zpool command

If all the paths to the storage are disabled, the Zpool command does not work. This is a behavior of the Zpool command. The Solaris Zpool agent uses pool commands to import and export Zpools. The command runs into a loop in and tries to export the Zpool when all the paths to the storage are lost. The loop continues till the storage paths are restored and Zpool is cleared.As a result, the offline and clean of Zpool Agent fails and the service group cannot fail over to the other node.

Resolution: You must restore the storage paths and run a Zpool clear command. Thus, the pending commands suceed and the service group fails over to another node.

Issues with configuration of resource values

If you configure a resource that has more than 425 values in its **ArgListValues**, the agent managing that resource logs a message such as:

```
VCS WARNING V-16-2-13806 Thread(1437547408) ArgListValues overflow;

Cannot append values more than upper limit of (425).
```

Normally, the number of values in **ArgListValues** for a resource must not exceed 425. However, in case of a keylist, association or vector type of attribute appears in the ArgList for a resource-type. Since these attributes can take multiple values, there is a chance for the resource values in **ArgListValues** to exceed 425.

Issues with bunker replay

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

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

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

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

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

The CmdServer process may not start in IPv6 environments in secure clusters

In an IPv6 environment on secure clusters, the CmdServer process may not start. In addition, security may not function correctly. If it does not start on a particular node, modify that node's /etc/hosts file so that the localhost resolves to ::1.

Workaround: In the /etc/hosts file, add the following:

```
::1        localhost
```

Issues related to the VCS engine

Missing host names in engine_A.log file

The GUI does not read the engine_A.log file. It reads the engine_A.ldf file, gets the message id from it, and then queries for the message from the bmc file of the appropriate locale (Japanese or English). The bmc file does not have system names present and so they are read as missing. [1736295]

Systems with multiple CPUs and copious memory shut-down time may exceed the ShutdownTimeout attribute (1919919)

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

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

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

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

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

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

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

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

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

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

Agent framework can reject `hares -action` **command**

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

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

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

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

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

Workaround

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

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

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

The hacf -cmdtocf command generates a broken main.cf file (1728738)

The `hacf -cmdtocf` command used with the `-dest` option and removes the `include`statements from the types files.

Workaround:

Add the include `statements` in the `main.cf` files that are generated using the `hacf -cmdtocf` command.

Issues related to the bundled agents

Solaris mount agent fails to mount Linux NFS exported directory

The Solaris mount agent mounts the mount directories. At this point, if it tries to mount a Linux NFS exported directory, the mount fails showing the following error:

```
nfs mount: mount: <MountPoint>: Not owner
```

This is due to system NFS default version mismatch between Solaris and Linux.

The workaround for this is to configure `MountOpt` attribute in mount resource and set `vers=3` for it.

Example

```
root@north $ mount -F nfs south:/test /logo/
nfs mount: mount: /logo: Not owner
root@north $
Mount nfsmount (
                MountPoint = "/logo"
                BlockDevice = "south:/test"
                FSType = nfs
                MountOpt = "vers=3"
              )
```

VolumeSet messages do not display in the engine log on Japanese locale

The VolumeSet messages do not get displayed in the engine log on Japanese locale.

Workaround: None.

Entry points that run inside a zone are not cancelled cleanly

Cancelling entry points results in the cancellation of only the `zlogin` process. The script entry points that run inside a zone are forked off using the `zlogin` command. However, the `zlogin` command forks off an `sh` command, which runs in the context of the Solaris zone. This shell process and its family do not inherit the group id of the `zlogin` process, and instead get a new group id. Thus, it is difficult for the agent framework to trace the children or grand-children of the shell process, which translates to the cancellation of only the `zlogin` process. [1179695]

Workaround: Oracle must provide an API or a mechanism to kill all the children of the `zlogin` process that was started to run the entry point script in the local-zone.

The zpool command runs into a loop if all storage paths from a node are disabled

The Solaris Zpool agent runs `zpool` commands to import and export zpools. If all paths to the storage are disabled, the zpool command does not respond. Instead, the zpool export command goes into a loop and attempts to export the zpool. This continues till the storage paths are restored and zpool is cleared. As a result, the offline and clean procedures of Zpool Agent fail and the service group cannot fail over to the other node.

Workaround: You must restore the storage paths and run the zpool clear command for all the pending commands to succeed. This will cause the service group to fail over to another node.

Sybase Agent is unable to distinguish between server processes with the same server name, when one process is in the global zone and another is in local. (3270401)

The Sybase agent is unable to distinguish between two Sybase Sever processes with the same server name. As a result, for the two Sybase Server processes, if one is in the global zone and the other in local, when one comes online, both of them are reported online.

Workaround:

There is no workaround for this issue.

Issues related to the VCS database agents

Concurrency violation due to process startup on failover node is not detected when detail monitoring is set for Oracle resources [2917558]

Inside a failover service group, when the administrator starts an Oracle resource on a node and if the Oracle instance is online on any other node within the cluster, the instance would come up. However, the database does not get mounted. In such circumstances, this startup attempt is detected by basic monitoring. If detail monitoring is enabled, this startup attempt does not get detected.

Workaround: No workaround.

VCS agent for Oracle: Health check monitoring does not work with Oracle 10.2.0.4

The health check monitoring in Oracle agent does not work with Oracle 10.2.0.4 due to incompatibility of the health check APIs provided by Oracle. [2101570]

Resolution: Disable health check monitoring by setting the MonitorOption attribute to 0 (zero).

VCS agent for Oracle: Make sure that the ohasd has an entry in the init scripts

Make sure that the ohasd process has an entry in the init scripts so that when the process is killed or the machine is rebooted, this automatically restarts the process.[1985093]

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

VCS agent for Oracle: Intentional Offline does not work

Due to issues with health check monitoring, Intentional Offline does not work for VCS agent for Oracle.

The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups

The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups.

Workaround:

Have a copy of the pfile/spfile in the default $GRID_HOME/dbs directory to make sure that this would be picked up during the ASM Instance startup.

VCS agent for ASM: Health check monitoring is not supported for ASMInst agent

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

Issues related to the agent framework

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

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

Resolution: No resolution.

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

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

Workaround

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

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

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

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

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

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

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

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

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

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

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

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

Agent performance and heartbeat issues

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

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

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

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

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

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

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

Issues related to Live Upgrade

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

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

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

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

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

Issues related to VCS in Japanese locales

This section covers the issues that apply to VCS 5.1 in a Japanese locale.

The hahbsetup command displays messages and warnings in English

The `hahbsetup` command incorrectly displays messages and warnings in English. [1652562]

The GetSafeMsg() returns error when one of the parameters is already localized

Some log messages that use the date-string parameter may not print correctly in non-English locales. [1715258, 1825966]

The gcoconfig script displays error messages in English

The gcoconfig script incorrectly displays English error messages. [1416136]

The hares -action command displays output in English

The `hares -action` command incorrectly displays output in English. [1786747]

Listener startup message on ja_JP.PCK locale may not be displayed correctly in the engine log (2075471)

Description: When you start the listener resource in VCS on ja_JP.PCK locale, the `lsnrctl` start command output is not shown correctly in the engine log. This happens because the converted output from the command generates special characters which are not handled by `halog` command.

Resolution: You can ignore this message as it does not have any functionality impact for the listener resource.

Issues related to global clusters

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

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

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

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

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

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

Issues related to LLT

LLT port stats sometimes shows recvcnt larger than recvbytes

With each received packet, LLT increments the following variables:

- recvcnt (increment by one for every packet)

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

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

This does not impact the LLT functionality.

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

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

VRTSllt, VRTSgab and VRTSvxfen patches do not rollback from local zones after uninstalling this patch from global zones (3249002)

When patches for the VRTSllt, VRTSgab or VRTSvxfen package are removed from a system with non-global zones running, the `pkginfo -l` *package_name* command output inside the non-global zone shows that the patches are still installed on the non-global zones.

This issue occurs because the `patchrm` command fails to rollback the package metadata information inside the non-global zone, even though patches are removed from the non-global zone.

Workaround:

The wrong output can be safely ignored, since the above packages are not required inside a non-global zone for functionality.

Issues related to I/O fencing

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

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

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

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

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

- The fencing setup uses customized mode with one or more CP servers.

- The application cluster has more than eight nodes.

- The node weight for a single node (say sys1 with node id 0) is more than the sum total of node weights for the rest of the nodes.

- A network fault occurs and the cluster partitions into two with the single node (sys1) on one part and the rest of the nodes on the other part.

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

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

Reconfiguring Veritas High Availability product with I/O fencing fails if you use the same CP servers

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

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

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

CP server cannot bind to multiple IPs (2085941)

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

Resolution: No known resolution for this issue.

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

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

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

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

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

To avoid the vxfen stop service timeout error

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

    # **hastop -all**

2   Reboot the systems:

    # **shutdown -i6 -g0 -y**

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

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

Issues related to Symantec Product Authentication Service with VCS

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

The vcsat and cpsat commands may appear to be hung

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

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

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

Workaround:

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

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

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

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

Veritas High Availability product may report AT error during system reboot

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

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

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

Issues related to Veritas Cluster Server agents for Veritas Volume Replicator

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

Issues related to IMF

Failure messages of resource un-registration with IMF appear in agent or engine logs after performing online or offline operations on the resource (2909184)

When a resource is registered with IMF for monitoring, any online or offline operation triggers un-registration of the resource from IMF. During such operations, agent may record an error message in the agent or engine logs stating that the un-registration failed. This issue is also observed for multiple resources.

Workaround:

There is no workaround. These failure messages are false positives and no
resolution is required. Agent registers resources with IMF again after some time.

Issues related to AMF

Issues with the amfstat output (2926158)

The `amfstat` output displays an extra column in the Registered Reapers list and
the `amfstat -n` output displays the header twice.

Workaround:

This issue does not have any effect on the functionality of AMF. It has been fixed
in VCS 6.0 and onwards.

AMF driver panics the system if VxFS module is unloaded while AMF has registrations related to file systems of type `VxFS` for resources of type `Mount` or `CFSMount`(3145047)

If you unload the Veritas File System (VxFS) module when Asynchronous Monitoring
Framework (AMF) has the registrations related to file systems of type `VxFS` for
resources of type `Mount` or `CFSMount`, the pointers by which AMF interacts with the
VxFS module change. But AMF cannot be aware of the change. As a result, when
AMF tries to communicate with VxFS using the old pointers, the system panics.

Workaround:

To solve this issue, disable Intelligent Monitoring Framework (IMF) for the Mount
and CFSMount agent resources before applying VxFS patches, unload and reload
the VxFS module, and re-enable IMF after the operations.

1   Disable IMF at the type level:

   ■   Check the current Mode value of IMF attribute for the Mount and CFSMount
       type at the type level by the following command and note them.

       # **hatype -display Mount   CFSMount -attribute IMF**

   ■   Disable IMF at the type level by the following command:

       # **hatype -modify Mount   IMF -update Mode 0**

       # **hatype -modify CFSMount   IMF -update Mode 0**

2   Disable IMF at the resource level if IMF is overridden at the resource level:

   ■   Check the current Mode value of the IMF attribute for resources of Mount
       and CFSMount if values is are overridden at the resource level and note
       them.

```
# hares -display  -attribute IMF -type CFSMount Mount
```

Note: If no resource is displayed with its IMF attribute overridden, you will see the following message:

```
VCS WARNING V-16-1-10554 No resource exists with attribute of
IMF
```

In this case, you can skip the step of disabling IMF at the resource level.

■ Disable IMF at the resource level for each resource for which value of IMF attribute is overridden at the resource level by the following command:

```
# hares -modify res_name IMF -update Mode 0
```

3   Make sure the monitor method is changed to Traditional for all resources of type Mount and CFSMount:

```
# hares -display -attribute MonitorMethod -type Mount CFSMount
```

4   Make sure that there is no registration of type VxFS in the amfstat output in following two sections:

```
AMF Status Report
....
....
Mount OFFLINE Monitors
==========================
Mount ONLINE Monitors
==========================
```

5   Apply VxFS patch or unload and reload the VxFS module

6   Enable IMF at the type level. Restore the Mode value at the type level back to the value before it's set to 0 in Step 1.

```
# hatype -modify Mount  IMF -update Mode original_val
```

```
# hatype -modify CFSMount  IMF -update Mode original_val
```

7 Restore the Mode value at resource level back to the value before it's set to 0 in Step 2 if any resource has overridden IMF attributes.

   # **hares -modify** *res_name* **IMF -update Mode** *original_val*

8 Make sure the monitor method is changed to IMF for all resources of type Mount and CFSMount:

   # **hares -display -attribute MonitorMethod -type Mount CFSMount**

Note: The resource will move to IMF-based monitoring after the next monitor completed. For example, after OfflineMonitorInterval for resources which are in the offline state or after MonitorInterval for resources which are in the online state.

Veritas Storage Foundation Cluster File System known issues

This section describes the Veritas Storage Foundation Cluster File System known issues in this release.

Miscalculated file set usage (2123429)

When file set quotas are enabled, it may be possible for VxFS to get into a state where it thinks a very large number of blocks are allocated to checkpoints. This issue can be seen using the fsckptadm command:

```
# fsckptadm getquotalimit /mnt1
Filesystem    hardlimit    softlimit    usage    action_flag
/mnt1         10000        10000        18446744073709551614
```

This could cause writes to checkpoints to fail. It could also trigger the removal of removable checkpoints.

Workaround

If this occurs, disabling and re-enabling file set quotas causes VxFS to recalculate the number of blocks used by checkpoints:

```
# fsckptadm quotaoff /mnt1
# fsckptadm quotaon /mnt1
# fsckptadm getquotalimit /mnt1
Filesystem    hardlimit    softlimit    usage    action_flag
/mnt1         10000        10000        99
```

installer –makeresponsefile detects the wrong product (2044525)

If you generate a response file to upgrade SFCFS or SFCFSHA using the `./installer -makeresponsefile` command, and then choose `G` (Upgrade a Product) option, the installer detects it as SFCFS RAC.

You can safely ignore that the installer detects it as SFCFS RAC.

CVMVolDg agent may fail to deport CVM disk group (2081221)

The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the `CVMDeportOnOffline` attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.

Workaround: If multiple CVMVolDg resources are configured for a shared disk group, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

Veritas Storage Foundation for Oracle RAC known issues

This section describes the Veritas Storage Foundation for Oracle RAC known issues in this release.

- Oracle RAC issues
- Veritas product issues

Oracle RAC issues

This section lists the known issues in Oracle RAC.

During installation or system startup, Oracle Grid Infrastructure may fail to start

After successful installation of Oracle RAC 11g Release 2 Grid Infrastructure, while executing the `root.sh` script, `ohasd` may fail to start. Similarly, during system startup, Oracle Grid Infrastructure may fail to start though the VCS engine logs may indicate that the cssd resource started Oracle Grid Infrastructure successfully.

The following message may be displayed on running the `strace` command:

```
# /usr/bin/strace -ftt -p pid_of_ohasd.bin
14:05:33.527288 open("/var/tmp/.oracle/npohasd", O_WRONLY <unfinished ...>
```

For possible causes and workarounds, see the Oracle Metalink document: 1069182.1

Oracle Grid Infrastructure installation may fail with the Veritas product installer

When you run the `installsfrac -configure` command to install Oracle Grid Infrastructure for Oracle RAC 11g Release 2, the installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

Workaround: Export the `OUI_ARGS` environment variable, before you run the Veritas product installation program:

```
export OUI_ARGS=-ignoreInternalDriverError
```

For more information, see the Oracle Metalink document: 970166.1

Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software

The Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software.If the failure indicates that the OCR and vote device locations are not shared, ignore the message.

Prerequisite check fails during Oracle RAC 11g installation

During Oracle RAC 11g installation, the Oracle Universal Installer (OUI) runs a product-specific prerequisite check that fails with the following message:

```
Checking the installed UDLM package...
Check complete. The overall result of this check is: Failed <<<<
Problem: The 11g compatible version of UDLM package is not present.
Recommendation: Remove the existing UDLM package using pkgrm and
install the 11g compatible UDLM package. The udlm/ directory at
the root of the 11g install media contains the required UDLM package.
README.udlm is bundled with the package binaries and contains
detailed installation instructions.
```

Workaround: Select the option "Manual" on the "Product-Specific Prerequisite Checks" page of the OUI to verify the flagged items manually.

Enabling ODM in Oracle RAC 11 Release 2 installations causes errors

Enabling ODM in Oracle RAC 11 Release 2 installations causes the following error:

```
'ODM ERROR V-41-4-1-253-12 Not enough space'
Oracle instance may also crash with same error.
```

The error is observed if the DISM (Dynamic Intimate Shared memory) feature is enabled. In Solaris, the Oracle database uses DISM if it is available on the system, and if the value of the `sga_max_size` initialization parameter is larger than the size required for all SGA components combined. [1913013]

Workaround: Make sure that the file `ORACLE_HOME/bin/oradism` is owned by the root user with "execute" and "setuid" permissions. If the problem persists after correcting the permissions, uncomment the `sga_max_size` and `memory_target` `init.ora` parameters.

Oracle VIP Configuration Assistant fails with an error message

During Oracle RAC 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

```
The given interface(s), "" is not public.
Public interfaces should be used to configure virtual IPs.
```

This message appears only when the VIP is not from the regular public IP range (for example, 200.). [1182220]

Workaround: Invoke the `vipca` utility manually as the superuser.

```
# export DISPLAY=nebula:0.0
# $CRS_HOME/bin/vipca
```

Oracle Cluster Verification utility displays a warning message

During the final stage of Oracle RAC 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification utility.

For example:

```
Utility
============================================================
OUI-25031: Some of the configuration assistants failed. It is
strongly recommended that you retry the configuration
assistants at this time. Not successfully running any "
Recommended" assistants means your system will not be correctly
configured.
1. Check the Details panel on the Configuration Assistant Screen
to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button
to retry them.
============================================================
```

Workaround: You may safely ignore this message if the cluster is operating satisfactorily.

Oracle Database Configuration Assistant displays an error

The Database Configuration Assistant utility displays the following error:

```
SGA size cannot be greater than maximum shared memory
segment size (0).
```

Workaround: Ignore this message and manually configure the database memory parameters for Oracle. In the "Memory" tab of the Oracle Database Creation Assistant (DBCA), select a Custom and Manual shared memory management configuration and enter the appropriate values.

Changing the Veritas agent for Oracle error handling

The Veritas agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file, oraerror.dat, which provides a list of Oracle errors and actions to address the errors.

For a description of the actions:

See the *Symantec High Availability Agent for Oracle Installation and Configuration Guide.*

Currently, the file specifies the NOFAILOVER action for the following Oracle errors: ORA-00061, ORA-02726, ORA-6108, ORA-06114

The NOFAILOVER action means that the agent sets the state of the resource to OFFLINE and freezes the service group. If you want to change this behavior, you can stop the agent, edit oraerror.dat, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

Veritas product issues

This section lists the known issues in Veritas product for this 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 10gR2 Clusterware: `oracle`; For Oracle 11gR2 Grid Infrastructure: `grid`).
   oraInventory_group_name is the name of the oraInventory group.
   Return to the former session and proceed with the installation.

Verification of Oracle binaries incorrectly reports as failed during Oracle Grid Infrastructure installation

The verification of Oracle binaries may incorrectly report as failed during the Oracle Grid Infrastructure installation using the Veritas product installer. The message is erroneously reported due to a break in passwordless SSH communication. The SSH communication fails because execution of the `root.sh` script changes the owner of the operating system root directory to the grid user directory.

`vxlustart` **utility fails to recognize mount points configured under VCS**

During Live Upgrade, the `vxlustart` utility fails to recognize the mount points that are configured under VCS. As a result, the data under the Oracle database and Oracle Clusterware mount points that are configured as CFS mount points under VCS get copied into the alternate boot environment.

Workaround:

1. Create a list of the CFS mount points configured under VCS before invoking the `vxlustart` utility:

   ```
   # for i in `hatype -resources CFSMount`; \
   do  hares -display  $i -attribute MountPoint | awk ' \
   NR != 1 { print $4}'; done > /tmp/sfracmnt
   ```

   Verify that the file `/tmp/sfracmnt` contains the correct list of directories configured as CFS mount points under VCS.

2. Run the `vxlustart` utility using the `-X` option along with other command options. The `-X` option indicates that the utility ignores the list of directories contained in the `/tmp/sfracmnt` file, thus preventing the files from being copied to the alternate boot environment.

   ```
   # ./vxlustart -v -u targetos_version \
   -s osimage_path -d disk_name -X /tmp/sfracmnt
   ```

   For details, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

3. Create the mount points manually on the alternate boot environment as follows:

   ```
   # for i in `cat /tmp/sfracmnt` ; \
   do mkdir -p /altroot.5.10/$i; done
   ```

4. Update the permissions, user name, and group name of the mount points (created on the ABE) to match that of the existing directories on the primary boot environment.

Error message displayed on rebooting the nodes after Live Upgrade of the cluster

When you reboot the nodes after completing Live Upgrade of the cluster, the following error message is displayed in the `/var/adm/messages` file:

```
vcsmm: [ID 269007 kern.notice]
VCS RAC ERROR V-10-1-15050 vcsmm_process_request: vrfsm is not running
```

[2090637]

You may ignore the message and proceed with the startup.

Messages scroll out of view on clusters with three or more nodes

On clusters with three or more nodes, messages scroll out of view during installation or configuration activities that print a large number of messages on screen. For example, when you run the installation and configuration checks using the **SF Oracle RAC Installation and Configuration Checks** option in the Veritas product installer menu on a three-node cluster, the messages run off the screen after the terminal window displays the first page of messages. These messages can not be viewed or retrieved.

Workaround: For any failures that may result during the checks, see the log file `/opt/VRTS/install/logs`.

Long messages run off the screen if the screen width is less than 100 characters

Messages that exceed 80 characters escape from view if the screen width of your terminal window is less than 100 characters. For example, when you run the installation and configuration checks using the **SF Oracle RAC Installation and Configuration Checks** option in the Veritas product installer menu, long messages run off the side on terminal window sizes less than 100.

Workaround: Set the terminal window size to a value greater than 100.

Installing all SFHA products on alternate root produces incorrect failure messages during initial reboot

For Solaris 10, while installing all the SFHA products on alternate root, the following failure messages appear during the initial reboot of the alternate root disk. These messages may also appear when you use the `JumpStart` command to install the SFHA products.

```
Reading ZFS config: done.
May 20 10:44:06 svc.startd[7]: svc:/system/llt:default: Method
"/lib/svc/method/llt start" failed with exit status 1.
May 20 10:44:06 svc.startd[7]: svc:/system/llt:default: Method
"/lib/svc/method/llt start" failed with exit status 1.
May 20 10:44:07 svc.startd[7]: svc:/system/llt:default: Method
"/lib/svc/method/llt start" failed with exit status 1.
May 20 10:44:07 svc.startd[7]: system/llt:default failed: transitioned to
maintenance (see 'svcs -xv' for details)
```

Workaround: Ignore these messages as the products are not yet configured.

Warning message displayed on taking cssd resource offline if LANG attribute is set to "eucJP"

When you take the cssd resource offline using the `hares -offline cssd` command and the LANG attribute is set to "eucJP", the following message may be observed in the `hamsg engine_A` command output [2123122]:

```
VCS INFO V-16-2-13716 Could not find message V-16-2-13716
```

You may ignore the message.

Deporting issues with shared disk groups

If you manually deport a shared disk group, the CVMVolDg agent does not automatically reimport it as a shared disk group. You must manually reimport it as a shared disk group.

Stopping cluster nodes configured with I/O fencing

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect or "split brain."

For more information, see *Veritas Cluster Server User's Guide*.

I/O fencing uses SCSI-3 Persistent Reservation keys to implement data protection. The software places keys on I/O fencing coordinator and data disks. The administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator disks and data disks to prevent possible difficulties with subsequent cluster startup. Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordinator and data disks. Depending on the order of reboot and subsequent startup events, the cluster might warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown command instead of the reboot command to perform a graceful reboot for systems.

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

Stopping VCS does not unregister port f from GAB membership

In an Veritas product cluster with all the CFS resources under VCS control, when you stop VCS, all the CFS resources must go down cleanly and CFS must unregister port f from GAB membership. Oracle RAC 10g Clusterware does not clean up all its processes when it is stopped. Now, when you stop VCS, all the CFS resources go down. However, due to the left over Oracle processes, CFS does not unregister port f from GAB membership.

Workaround: Perform the following steps to bring down port f.

To bring down port f

1   Kill all the Oracle processes.

```
# kill -9 'ps -u oracle|awk '{print $1}'`
```

2   Verify that all CFS file systems are unmounted.

```
# mount | grep cluster
```

3   Unregister port f from GAB membership.

```
# fsclustadm cfsdeinit
```

GAB and LLT kernel modules are still loaded after uninstalling

After a successful removal of Veritas product, the GAB and LLT kernel modules are still loaded. [605017]

Workaround: Run the command "modunload" for GAB and LLT on all nodes:

```
# modunload -i <gms_module_id>
# modunload -i <gab_module_id>
# modunload -i <llt_module_id>
```

To find the module ID, use modinfo.

For example, to find the module ID for gms:

```
# modinfo | grep gms
```

DBED features are not integrated with GCO

DBED features are not integrated with Global Cluster Option (GCO). After GCO migration, be aware that DBED features will not be functional. [1241070]

Issue with format of the last 8-bit number in private IP addresses

The PrivNIC/MultiPrivNIC resources fault if the private IP addresses have a leading 0 in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1. or X.0X.X.1 or 0X.X.X.1, where X is an octet of the IP address. [1164506]

When you configure private IP addresses for Oracle Clusterware, ensure that the IP addresses have a format as displayed in the following two-node example:

■   On sys1: 192.168.12.1

■   On sys2: 192.168.12.2

Confirm the correct format by viewing the PrivNIC or MultiPrivNIC resource in the `/etc/VRTSvcs/conf/config/main.cf` file.

When master node loses access to complete storage, detached sites remain in RECOVER state even after reattaching and recovering the sites

In a campus cluster environment, if the master node loses access to complete storage, all but one of the sites is detached and the DCO volumes may get detached if the `dgfailpolicy` is set to `dgdisable`. If the detached sites are reattached and recovered, the site still remains in RECOVER state. [1828142]

Workaround: Change the status of the site as described in the following procedure to resolve the issue.

To change the status of the site

1   Log onto the CVM master node.

2   Reattach the detached sites:

    # **vxdg -g** *dg_name* **reattachsite** *site_name*

    The site remains in RECOVER state.

3   Restore DCO volumes by unpreparing and preparing the volumes.

    Unprepare the volumes:

    # **vxsnap -g** *dg_name* **-f unprepare vol_name**

    Prepare the volumes:

    # **vxsnap -g** *dg_name* **prepare** *vol_name* **drl=on**

4   Reattach the detached sites:

    # **vxdg -g** *dg_name* **reattachsite** *site_name*

5   Verify that the state of the detached sites is now ACTIVE:

    # **vxprint**

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 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`**

MultiPrivnic agent reports error messages in the VCS engine logs when zone goes down (2624304)

When a zone is down, the `MultiPrivNIC` agent fails to get active device information. The agent prints this message without specifying the message ID.

As a result, the following warning message displays:

`VCS WARNING V-16-1-11328 Invalid message ID specified`

Workaround: You may ignore this message.

11.2.0.1 crsd.bin Fails in clsCcICIscWait (2933706)

The CRS processes `crsd.bin, evmd.bin, ohasd.bin` are dead.

Workaround:

Apply the 11814167 patch for this bug. For more information see Metalink ID 1326008.1 on the ORACLE support site.

File system check daemon fails to restart after abnormal termination (2689195)

The file system check daemon (`vxfsckd`) fails to update the `vxfsckd-pid` file with the new process ID (pid) of the vxfsckd process after abnormal termination. As a result, the CFSfsckd agent fails to detect the status of the vxfsckd daemon.

Workaround: Perform the following steps to resolve the issue on the node where the `vxfsckd` resource faults:

1. Log into the node as the root user.

2. Kill all `vxfsckd` processes:

   ```
   # kill -9 `ps -ef|grep vxfsckd|awk '{print $2}'`
   ```

3. Remove the `vxfsckd-pid` file:

   ```
   # rm /var/adm/cfs/vxfsckd-pid
   ```

4. Bring the `vxfsckd` resource online:

   ```
   # hares -online vxfsckd_resname -sys node_name
   ```

vxdmpasm cannot create the "auto:ASM" TYPE for the ASM disk (2944387)

The `vxdmpasm` command cannot create the "auto:ASM" TYPE for the ASM disk.

Workaround:

There is no workaround for this issue.

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

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

Workaround:

1. Log into each node in the cluster as the root user.

2. Perform the following operations:

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

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

     where:

     oracle_base is the name of the Oracle base directory.

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

     oraInventory_group_name is the name of the oraInventory group.
     Complete the other preparatory tasks before you invoke the installation program. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

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

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

     where:

     oracle_base is the name of the Oracle base directory.

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

     oraInventory_group_name is the name of the oraInventory group.
     Return to the former session and proceed with the installation.

Software limitations

This section covers the software limitations of this release.

Veritas Storage Foundation software limitations

There are no Veritas Storage Foundation software limitations in the 5.1 SP1 RP4 release.

Veritas File System software limitations

There are no Veritas File System software limitations in the 5.1 SP1 RP4 release.

Veritas Volume Replicator software limitations

There are no Veritas Volume Replicator software limitations in the 5.1 SP1 RP4 release.

Veritas Storage Foundation for Databases tools software limitations

There are no Veritas Storage Foundation for Databases toolssoftware limitations in the 5.1 SP1 RP4 release.

Veritas Cluster Server software limitations

This is the Veritas Cluster Server software limitations in the 5.1 SP1 RP4 release.

The Ldom agent do not gracefully detect guest domain migration

If the guest domain managed by the LDom agent is migrated to another host, it will not be gracefully detected by the agent and the resource will be marked FAULTED.

Workaround:

Upgrade to VCS version 6.0 or higher for the Ldom agent capability to gracefully responds to migration.

Documentation errata

The following section provides documentation updates.

Veritas Storage Foundation and High Availability Solutions Virtualization Guide (2909333)

The following line should be added before the `hazonesetup` command syntax of step 1, the "Modifying the service group configuration" section of *5.1 SP1 Veritas Storage Foundation and High Availability Solutions Virtualization Guide*.

Zone should be running when using the `hazonesetup` utility to configure Zones in VCS.

List of patches

This section lists the patches and packages for 5.1 SP1 RP4.

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

`./installrp -listpatches`

Table 1-40   Patches and packages for Solaris SPARC

Patch ID	Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
142629-19	VRTSvxvm	DMP, SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	200 MB	X	X
142631-07	VRTSdbed	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	48 MB	X	X
142633-13	VRTSvxfs	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	37 MB	X	
142634-13	VRTSvxfs	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	47 MB		X
143270-10	VRTSodm	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	1.2 MB	X	
143271-10	VRTSodm	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	1.1 MB		X
143273-09	VRTScavf	SFCFS, SFCFSHA, SFRAC, SVS	829 KB	X	
143274-09	VRTScavf	SFCFS, SFCFSHA, SFRAC, SVS	830 KB		X
143281-05	VRTSllt	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	1.9 MB	X	
143282-05	VRTSllt	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	1.6 MB		X
143283-04	VRTSgab	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	2.2 MB	X	
143284-04	VRTSgab	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	1.6 MB		X

Table 1-40   Patches and packages for Solaris SPARC *(continued)*

Patch ID	Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
143287-09	VRTSvcs	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	116 MB	X	X
143288-14	VRTSvcsag	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	3.8 MB	X	X
143289-03	VRTScps	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	40 MB	X	X
143290-06	VRTSvcsea	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	16 MB	X	X
143680-04	VRTSglm	SFCFS, SFCFSHA, SFRAC, SVS	910 KB	X	
143681-04	VRTSglm	SFCFS, SFCFSHA, SFRAC, SVS	578 KB		X
143687-04	VRTSob	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	48 MB	X	X
144159-01	VRTSsfmh	DMP, SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	32 MB	X	X
145450-05	VRTSdbac	SFRAC	6.6 MB	X	
145451-05	VRTSdbac	SFRAC	5.2 MB		X
145454-05	VRTSvxfen	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	3.0 MB	X	
145455-05	VRTSvxfen	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	2.4 MB		X
145471-04	VRTSamf	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	1.7 MB	X	
145473-04	VRTSamf	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	2.2 MB		X
149806-01	VRTSperl	AT,DMP, SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	44 MB	X	X

Table 1-41   Patches and packages for Solaris x64

Patch ID	Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
142630-19	VRTSvxvm	DMP, SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	183 MB		X
142632-07	VRTSdbed	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	11 MB		X
142635-13	VRTSvxfs	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	29 MB		X
143272-10	VRTSodm	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	974 KB		X
143275-09	VRTScavf	SFCFS, SFCFSHA, SFRAC, SVS	838 KB		X
143291-05	VRTSllt	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	1.6 MB		X
143292-04	VRTSgab	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	1.6 MB		X
143294-08	VRTSvcs	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	204 MB		X
143295-14	VRTSvcsag	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	6.0 MB		X
143296-03	VRTScps	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	41 MB		X
143297-06	VRTSvcsea	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	21 MB		X
143682-04	VRTSglm	SFCFS, SFCFSHA, SFRAC, SVS	588 KB		X
143693-04	VRTSob	SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	46 MB		X
145452-05	VRTSdbac	SFRAC	4.7 MB		X
145456-05	VRTSvxfen	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	2.4 MB		X

Table 1-41   Patches and packages for Solaris x64 *(continued)*

Patch ID	Package Name	Products Affected	Patch Size	Solaris 9	Solaris 10
145458-01	VRTSsfmh	DMP, SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS	37 MB		X
145472-04	VRTSamf	SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	1.9 MB		X
149807-01	VRTSperl	AT,DMP, SF, SFCFS, SFCFSHA, SFHA, SFRAC, SVS, VCS	41 MB	X	X

Downloading the 5.1 SP1 RP4 archive

The patches that are included in the 5.1 SP1 RP4 release are available for download from the Symantec website. After downloading the 5.1 SP1 RP4 rolling patch, use gunzip and tar commands to uncompress and extract it.

For the 5.1 SP1 RP4 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

http://www.symantec.com/docs/TECH75362

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 RP4 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.

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 https://sort.symantec.com/patches to see if there are any patches available for the 5.1 SP1 Installer.

4 Change the directory to `/tmp/sfha51sp1`:

   # **cd /tmp/sfha51sp1**

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

     # **./installer**

6   Download SFHA 5.1 SP1 RP4 from https://sort.symantec.com/patches.

7   Extract it to a directory called /tmp/sfha51sp1rp4.

8   Check https://sort.symantec.com/patches to see if there are patches available
     for the 5.1 SP1 RP4 installer. Download applicable P-patches and extract them
     to the /tmp directory.

9   Change the directory to /tmp/sfha51sp1rp4:

     # **cd /tmp/sfha51sp1rp4**

10   Invoke the installrp script to install 5.1 SP1 RP4:

     # **./installrp**

11   If you did not configure the product after the 5.1 SP1 installation, the installer
     prompts you to configure the product during RP4 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
RP4 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.

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 RP4 with the Veritas Web-based installer

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

To install Veritas product

1   The 5.1 SP1 version of the Veritas product must be installed before upgrading to 5.1 SP1 RP4.

2   On the **Select a task and product** page, select **Install 5.1 SP1 RP4** from the **Task** drop-down list, and click **Next.**

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

4   You have the option to let the installer configure SSH or RSH communications between the systems. If you choose to allow this configuration, select the shell and provide the root passwords for each system.

5   After the validation completes successfully, it prompts you with a yes or no optional confirmation prompt, which contains warning messages and notes. If you want to ignore the warning messages, you can click **Yes** to the question **Do you want to continue?**.

6   Click **Next** to install 5.1 SP1 RP4 patches on the selected system.

7   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.

8   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**.

Installing the products using JumpStart

This chapter includes the following topics:

- Generating the finish scripts
- Overview of JumpStart installation tasks
- Preparing installation resources
- Adding language pack information to the finish file

Generating the finish scripts

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

To generate the finish script

1   Mount the 5.1 SP1 RP4 media and run the `installrp` program to generate the scripts.

   # **installrp -jumpstart** *directory_to_generate_scripts*

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

   For example:

   # **./installrp -jumpstart /js_scripts**

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

3 Specify a disk group name for the root disk.

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

4 Specify private region length.

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

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

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

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

```
The finish scripts for AT is generated at
/js_scripts/jumpstart_at.fin
The finish scripts for DMP is generated at
/js_scripts/jumpstart_dmp.fin
The finish scripts for FS is generated at
/js_scripts/jumpstart_fs.fin
The finish scripts for SF is generated at
/js_scripts/jumpstart_sf.fin
The finish scripts for SFCFS is generated at
/js_scripts/jumpstart_sfcfs.fin
The finish scripts for SFCFSHA is generated at
/js_scripts/jumpstart_sfcfsha.fin
The finish scripts for SFHA is generated at
/js_scripts/jumpstart_sfha.fin
The finish scripts for SFRAC is generated at
/js_scripts/jumpstart_sfrac.fin
The finish scripts for SVS is generated at
/js_scripts/jumpstart_svs.fin
The finish scripts for VCS is generated at
/js_scripts/jumpstart_vcs.fin
The finish scripts for VM is generated at
/js_scripts/jumpstart_vm.fin
The encapsulation boot disk script for VM is generated at
/js_scripts/encap_bootdisk_vm51102000.fin
```

List the js_scripts directory.

```
# ls /js_scripts
```

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

For SF:

```
encap_bootdisk_vm51102000.fin jumpstart_sf.fin
```

For SFHA:

```
encap_bootdisk_vm51102000.fin jumpstart_sfha.fin
```

For SFCFS:

```
encap_bootdisk_vm51102000.fin jumpstart_sfcfs.fin
```

For SF Oracle RAC:

```
encap_bootdisk_vm51102000.fin jumpstart_sfrac.fin
```

For VCS:

```
encap_bootdisk_vm51102000.fin jumpstart_vcs.fin
```

For DMP:

```
encap_bootdisk_vm51102000.fin jumpstart_dmp.fin
```

For SVS:

```
encap_bootdisk_vm51102000.fin jumpstart_svs.fin
```

For FS:

```
encap_bootdisk_vm51102000.fin jumpstart_fs.fin
```

For VM:

```
encap_bootdisk_vm51102000.fin jumpstart_vm.fin
```

For AT:

```
encap_bootdisk_vm51102000.fin jumpstart_at.fin
```

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

    For example:

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

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

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

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

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

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

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

   ■ `minpkgs`

   ■ `recpkgs`

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

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

     # **./installrp -listpatches**

     See "The installrp script options" on page 12.

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

Overview of JumpStart installation tasks

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

Summary of tasks

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

2 Read the JumpStart installation instructions.

3 Generate the finish scripts.

See "Generating the finish scripts" on page 168.

4 Modify the rules file for JumpStart.

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

5 Prepare installation resources.

See "Preparing installation resources" on page 173.

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

For Solaris SPARC:

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

For Solaris x64:

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

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

---

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

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

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

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1   Copy the contents of 5.1 SP1 disk and 5.1 SP1 RP4 disk both to buildsrc.

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

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

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

See "Generating the finish scripts" on page 168.

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

```
# ./installrp -listpatches

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

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

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

4   Copy the install and uninstall scripts that you created when you generated the finish script to *BUILDSRC* if you want to configure or uninstall from /opt/VRTS/install. Otherwise, you need to configure and uninstall from disc.

See "Generating the finish scripts" on page 168.

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

See "Generating the finish scripts" on page 168.

6 Modify the rules file as required.

For example:

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

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

Adding language pack information to the finish file

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

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

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

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

do
...
done
```

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

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

do
.
.
.
done
```

Upgrading to 5.1 SP1 RP4

This chapter includes the following topics:

- Prerequisites for upgrading to 5.1 SP1 RP4
- Downloading required software to upgrade to 5.1 SP1 RP4
- Supported upgrade paths
- Upgrading to 5.1 SP1 RP4
- Verifying software versions

Prerequisites for upgrading to 5.1 SP1 RP4

The following list describes prerequisites for upgrading to the 5.1 SP1 RP4 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 RP4 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.

Downloading required software to upgrade to 5.1 SP1 RP4

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

To download required software to upgrade to 5.1 SP1 RP4

1 Download SFHA 5.1 SP1 RP4 from https://sort.symantec.com/patches.

2 Extract it to a directory such as `/tmp/sfha51sp1rp4`.

3 Check https://sort.symantec.com/patches to see if there are patches available for the 5.1 SP1 RP4 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 5.1 SP1 RP4 installer patches.

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 5.1 SP1 to 5.1 SP1 RP4

- 5.1 SP1 PR3 to 5.1 SP1 RP4

- 5.1 SP1 RP1 to 5.1 SP1 RP4

- 5.1 SP1 RP2 to 5.1 SP1 RP4

- 5.1 SP1 RP3 to 5.1 SP1 RP4

- 5.1 SP1 P-patch to 5.1 SP1 RP4

- 5.1 SP1 RP1 P-patch to 5.1 SP1 RP4

- 5.1 SP1 RP2 P-patch to 5.1 SP1 RP4

- 5.1 SP1 RP3 P-patch to 5.1 SP1 RP4

Upgrading to 5.1 SP1 RP4

This section describes how to upgrade from 5.1 SP1 or later RPs, or 5.1 SP1 PR1 or later RPs to 5.1 SP1 RP4 on a cluster or a standalone system.

- Performing a full upgrade to 5.1 SP1 RP4 on a cluster
  Use the procedures to perform a full upgrade to 5.1 SP1 RP4 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System (SFCFS), Veritas Storage Foundation for Oracle RAC (SFRAC), or Symantec VirtualStore (SVS) installed and configured.

- Perfoming a full upgrade to 5.1 SP1 RP4 on a standalone system
  Use the procedure to upgrade to 5.1 SP1 RP4 on a system that has SF installed.

- Performing a full upgrade to Veritas products using Live Upgrade
  Use the procedure to upgrade your Veritas product with a Live Upgrade.

- Performing a rolling upgrade using the installer
  Use the procedure to upgrade your Veritas product with a rolling upgrade.

Performing a full upgrade to 5.1 SP1 RP4 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 RP4:

- Performing a full upgrade to 5.1 SP1 RP4 on a Veritas Cluster Server

- Performing a full upgrade to 5.1 SP1 RP4 on an SFHA cluster

- Performing a full upgrade to 5.1 SP1 RP4 on an SFCFSHA cluster

- Performing a full upgrade to 5.1 SP1 RP4 on an SF Oracle RAC cluster
  See "Downloading required software to upgrade to 5.1 SP1 RP4 " on page 176.

Performing a full upgrade to 5.1 SP1 RP4 on a Veritas Cluster Server

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

To upgrade VCS

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

2   Log in as superuser.

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

3   If you install VCS on Solaris 10 systems that run non-global zones, make sure that all non-global zones are booted and in the running state on each node before you upgrade the VCS stack in the global zone.

4   On any nodes in the cluster, make the VCS configuration read-only:

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

5   Check the readiness of the nodes where you plan to upgrade. From the directory
     that contains the extracted and untarred 5.1 SP1 RP4 rolling patch binaries,
     change to the directory that contains the installrp script. Start the pre-upgrade
     check:

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

6   Resolve any issues that the precheck finds.

7   Start the upgrade:

     # **./installrp** *node1 node2 ... nodeN*

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

Performing a full upgrade to 5.1 SP1 RP4 on an SFHA cluster

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

To perform a full upgrade to 5.1 SP1 RP4 on an SFHA cluster

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

2   Log in as superuser.

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

4   If you install SFHA on Solaris 10 systems that run non-global zones, make
     sure that all non-global zones are booted and in the running state on each
     node before you upgrade the SFHA stack in the global zone.

5   On each node in the cluster, make the VCS configuration read only:

     # **haconf -dump -makero**

6   Stop VCS.

     To stop applications, unmount VxFS file systems and stop VxVM volumes
     managed by VCS.

     # **hastop -all**

7   Stop all the applications that are using VxFS files systems and VxVM volumes
     which are not managed by VCS.

     Use application's native commands to stop applications.

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

   # **df -F vxfs**

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

   # **umount /*checkpoint_name*

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

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

   # **df -F vxfs**

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

   # **umount /*filesystem*

10 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.

11 Stop activity to all VxVM volumes that are not managed by VCS.

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

12 On each node, stop all VxVM volumes by entering the following command for each disk group, which are not managed by VCS:

    # **vxvol -g *diskgroup* stopall**

    Verify that no volumes remain open:

    # **vxprint -Aht -e v_open**

13 Deport all the disk groups which are not managed under VCS.

    # **vxdg deport *diskgroup***

14 If required, apply the OS kernel patches.

    See Oracle's documentation for the procedures.

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

    # **/opt/VRTS/bin/vxsvcctrl status**

    If the VEA service is running, stop it:

    # **/opt/VRTS/bin/vxsvcctrl stop**

16 On each node, stop the VCS command server:

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

    pid_of_CmdServer is the process ID of CmdServer.

17 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP4 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

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

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

    Resolve any issue that the precheck finds.

18 Start the upgrade.

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

    Review the output.

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

    # **hagrp -online** *service_group* **-sys** *nodename*

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

21 Import all the diskgroups that are not managed by VCS:

    # **vxdg import** *diskgroup*

22 Restart all the volumes by entering the following command for each disk group
    that are not managed by VCS:

    # **vxvol -g** *diskgroup* **startall**

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

    # **vxrvg -g** *diskgroup* **start** *rvg_name*

24 Remount all VxFS file systems on all nodes, which are not managed by VCS:

    # **mount -F vxfs** *blockdevice mountpoint*

25 Remount all Storage Checkpoints on all nodes:

    # **mount -F vxfs -o ckpt=***name blockdevice checkpoint_name*

26 Start all applications which are using VxFS files systems that are not managed
    by VCS.

    Use application native commands to start the applications.

Performing a full upgrade to 5.1 SP1 RP4 on an SFCFSHA cluster

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

To perform a full upgrade to 5.1 SP1 RP4 on an SFCFSHA cluster

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

2 Log in as superuser.

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

4   If you install SFCFSHA on Solaris 10 systems that run non-global zones, make sure that all non-global zones are booted and in the running state on each node.

Run the following command to get the state of non-global zones:

```
#zoneadm list -cv
```

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

```
# haconf -makerw
```

6   Enter the following command to freeze HA service group operations on any node:

```
# hagrp -freeze groupname -persistent
```

7   Make the configuration read-only:

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

8   Stop VCS. To stop applications, unmount VxFS/CFS file systems and stop VxVM or CVM volumes managed under VCS.

```
# hastop -all
```

9   Stop all applications which are using CFS file systems and VxVM volumes not managed by VCS. Use application native commands to stop applications.

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

```
# df -F vxfs
```

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

```
# cfsumount /checkpoint_name
```

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

```
# df -F vxfs
```

■   If any VxFS/CFS file systems are present, on each node in the cluster, stop IOs on the file systems, unmount all of the VxFS/CFS file systems:

```
# cfsumount /filesystem
```

Or

```
# umount /filesystem
```

12 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.

13 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.

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

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

Verify that no volumes remain open:

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

15 Deport all the disk groups which are not managed under VCS:

```
# vxdg deport diskgroup
```

16 On each node, stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

17 If required, apply the OS kernel patches.

See "System requirements" on page 20.

See Oracle's documentation for the procedures.

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

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

If the VEA service is running, stop it:

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

19 From the directory that contains the extracted and untarred 5.1 SP1 RP4 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

Review the output.

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

21 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

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

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

23 Make the configuration read-only:

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

24 Bring the CVM service group online on each node:

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

25 Import all the VxVM or CVM diskgroups that are not managed by VCS:

```
# vxdg import diskgroup
```

or

```
$vxdg import -s disgroup
```

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

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

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

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

28 Remount all VxFS/CFS file systems on all nodes:

```
# cfsmount /mountpoint
```

Or

```
# mount -Ffstype blockdevice mountpoint
```

29 Remount all Storage Checkpoints on all nodes:

```
# cfsmount /checkpoint_name
```

Or

```
# mount vxfs -o ckpt=name blockdevicemountpoint
```

30 Start all applications which are using VxFS/CFS file systems that are not managed by VCS. Use the application native commands to start the applications.

Performing a full upgrade to 5.1 SP1 RP4 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 RP4 on an SF Oracle RAC cluster

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

2 Log in as superuser.

3 Verify that `/opt/VRTS/bin` and /opt/VRTSvcs/bin are 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 Database is not controlled by VCS, enter the following command on any node of the cluster to stop Oracle Database:

```
# $ORACLE_HOME/bin/srvctl stop database -d $DB_NAME
```

8 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
```

9 Stop VCS.

```
# hastop -all
```

10 If required, apply the Solaris OS patches.

See System Requirements.

See *Oracle's* documentation for the procedures.

11 If required, upgrade the operating system and reboot the systems.

12 From the directory that contains the extracted and untarred 5.1 SP1 RP4 rolling patch binaries, change to the directory that contains the installrp script. If ssh key authentication is configured then enter:

   # **./installrp *node1 node2 ... nodeN***

   If ssh is not configured then enter:

   # **./installrp -rsh *node1 node2 ... nodeN***

   Review the output.

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

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

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

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

17 Relink the SF Oracle RAC libraries with Oracle.

   Refer to *Veritas Storage Foundation for OracleRAC 5.1 SP1 or later Installation and Configuration Guide* for more information.

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

   # **haconf -makerw**

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

   # **hasys -unfreeze -persistent *nodename***

20 Make the configuration read-only:

   # **haconf -dump -makero**

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

   # **hagrp -online *service_group* -sys *nodename***

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

   # **vxvol -g *diskgroup* startall**

23 Remount all VxFS file systems on all nodes:

> # **mount /filesystem**

24 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**

25 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**

Perfoming a full upgrade to 5.1 SP1 RP4 on a standalone system

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

To upgrade to 5.1 SP1 RP4 on a standalone system

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

2 Log in as superuser.

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

4 If required, apply the OS kernel patches.

See "System requirements" on page 20.

See Oracle's documentation for the procedures.

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

> # **df -F vxfs**

6 Unmount all Storage Checkpoints and file systems:

> # **cfsumount /checkpoint_name**
> # **cfsumount /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 Copy the patch archive downloaded from the patch central to temporary location, untar the archive and browse to the directory containing the installrp installer script. Enter the `installrp` script:

  # **./installrp *nodename***

12 If necessary, reinstate any missing mount points in the /etc/vfstab 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 7, 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**

Performing a full upgrade to Veritas products using Live Upgrade

This section describes how to upgrade 5.1 SP1 to 5.1 SP1 RP4 using Live Upgrade.

Supported live upgrade paths:

- Upgrading Veritas Products without Solaris OS upgrade:
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 9 Update x 5.1 SP1 RP4
  - Upgrading Solaris 10 Update x 5.1 SP1 to Solaris 10 Update x 5.1 SP1 RP4
- Upgrading Veritas Products with Solaris OS upgrade
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 9 Update y 5.1 SP1 RP4
  - Upgrading Solaris 9 Update x 5.1 SP1 to Solaris 10 Update y 5.1 SP1 RP4
  - Upgrading Solaris 10 Update x 5.1 SP1 to Solaris 10 Update y 5.1 SP1 RP4

Prerequisites to upgrade to 5.1 SP1 RP4 using Live Upgrade:

- The node should have an alternate boot disk that is identical to the primary boot disk.
- Installation disc for 5.1 SP1 and 5.1 SP1 RP4 to be installed on the ABE.
- Installation disc for target OS to be installed on ABE.
- The latest list of required patches is available in the Oracle Solaris Live Upgrade Software:
  Patch Requirements (Doc ID 1004881.1) document in My Oracle Support (https://support.oracle.com/).

- If OS upgrade is involved, then remove the currently installed SUNWluu, SUNWlur and SUNWlucfg packages and install SUNWluu, SUNWlur, SUNWlucfg packages from target OS. Also replace SUNWluzone if zones are involved.

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

Upgrading Veritas products using Live Upgrade from 5.1 SP1 to 5.1 SP1 RP4 without OS upgrade

This section describes how to upgrade SF, SFHA, SFCFSHA, VCS, or SF for Oracle RAC from 5.1 SP1 to 5.1 SP1 RP4 using Live Upgrade where OS upgrade is not involved.

To upgrade your Veritas product using Live Upgrade

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

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

2   Run the `vxlustart` command with -v and other required options to ensure there are no problems before beginning the Live Upgrade process.

    If the `vxlustart` command with -v and other required options reports success, proceed with running the `vxlustart` command.

    If the `vxlustart` command with -v and other required options reports errors, correct the problem, and run the command again.

    Note: This `vxlustart` command with -v and other required options does not catch failures that are reported by Solaris Live Upgrade commands.

    ---

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

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

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

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

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

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

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

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

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

6 Run the `lustatus` command to verify that the alternate boot environment is active. For example:

```
# lustatus
Boot Environment         Is        Active Active    Can    Copy
Name                     Complete  Now    On Reboot Delete Status
------------------------ -------- ------ --------- ------ ----------
source.4605              yes       no     no        yes    -
dest.4605                yes       yes    yes       no     -
```

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

```
# gabconfig -a
```

Upgrading Veritas products using Live Upgrade from 5.1 SP1 to 5.1 SP1 RP4 with OS upgrade

This section describes how to upgrade SF, SFHA, SFCFSHA, VCS, or SF for Oracle RAC from 5.1 SP1 to 5.1 SP1 RP4 using Live Upgrade where OS upgrade is involved.

To upgrade your Veritas product using Live Upgrade

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

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

2 Run the `vxlustart` command with -v and other required options to ensure there are no problems before beginning the Live Upgrade process.

If the `vxlustart` command with -v and other required options reports success, proceed with running the `vxlustart` command.

If the `vxlustart` command with -v and other required options reports errors, correct the problem, and run the command again.

Note: This `vxlustart` command with -v and other required options does not catch failures that are reported by Solaris Live Upgrade commands.

---

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

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

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

> `# ./uninstallprod -rootpath /altroot_path`

5 Install Veritas products again on ABE.

> `# ./installprod -rootpath /altroot_path`

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

> `# ./installrp -rootpath  /altroot_path`

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

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

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

  > `# ./vxlufinish -u target_os_version`

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

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

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

> `# shutdown -g0 -y -i6`

10 Verify that the alternate boot environment is active.

> `# lustatus`

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

```
# gabconfig -a
```

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

Performing a rolling upgrade using the installer

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

- About rolling upgrades
- Prerequisites for a rolling upgrade
- Performing a rolling upgrade using the installer

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 for Oracle RAC
- Symantec VirtualStore

You can perform a rolling upgrade from 5.1 SP1 or later, or 51 SP1 PR1.

Prerequisites for a rolling upgrade

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 SF Oracle RAC, stop Oracle Clusterware before upgrading Kernel packages.

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

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

Performing a rolling upgrade using the installer

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

Performing a rolling upgrade on kernel packages for VCS, SFHA, SVS, SFCFS and SFCFSHA: phase 1

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

To perform the rolling upgrade on kernel packages: phase 1

1   Stop all applications that access volumes.

2   Unmount all the file systems that are managed by SF.

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

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

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

5   The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel packages.

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

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

8   Before you proceed to phase 2, complete step 1 to step 6 on the second subcluster.

Performing a rolling upgrade on non-kernel packages for VCS, SFHA, SVS, SFCFS and SFCFSHA : phase 2

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

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

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

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

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

3   Installer will start HA daemon (had) on all nodes, HA will be available once HA daemon is up.

4   Verify the cluster's status:

```
# hastatus -sum
```

5   If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Preparing to perform a rolling upgrade for SFRAC

Perform the preparatory steps in this section if you are performing a rolling upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

Note: Perform the steps on the first subcluster.

To prepare to upgrade SF Oracle RAC:

1 Log in as superuser to one of the nodes in the cluster.

2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/lltab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

   For example:

   ```
   # cp /etc/VRTSvcs/conf/config/main.cf \
   /etc/VRTSvcs/conf/config/main.cf.save
   # cp /etc/VRTSvcs/conf/config/types.cf \
   /etc/VRTSvcs/conf/config/types.cf.save
   # cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
   /etc/VRTSvcs/conf/config/OracleTypes.cf.save
   # cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
   /var/VRTSvcs/conf/config/PrivNIC.cf.save
   # cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
   /var/VRTSvcs/conf/config/MultiPrivNIC.cf.save
   ```

3 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

   If the applications are under VCS control:

   ```
   # hagrp -offline grp_name -sys node_name
   ```

   If the applications are not under VCS control, use native application commands to stop the application.

4 For Oracle RAC 10g and Oracle RAC 11g

   Stop the Oracle RAC resources on each node.

   ■ If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

   ```
   # hagrp -offline grp_name -sys node_name
   ```

   ■ If the database instances are not managed by VCS, then run the following on one node:

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

5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS

starts. Failing to perform this step results in the database attempting to come online after the upgrade; the attempt fails due to the presence of old libraries on the system.

```
# haconf -makerw
# hagrp -modify oracle_group AutoStart 0
# haconf -dump -makero
```

■ If the Oracle database is not managed by VCS, change the management policy for the database to manual:

```
$ srvctl modify database -d db_name -y manual
```

6   Take all the VCS service groups offline:

```
# hagrp -offline grp_name -sys sys_name
```

7   Unmount all the VxFS file system which is not under VCS control.

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

```
# fuser -c /mount_point
```

```
# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

Performing a rolling upgrade on kernel packages for SFRAC: phase 1

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

To perform the rolling upgrade on kernel packages: phase 1

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

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

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

3   The installer checks system communications, package versions, product versions, and completes prechecks.

   It then upgrades applicable product kernel packages.

4   The installer loads new kernel modules.

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

   In case of failure in the startup of some of the processes, you may need to reboot the nodes and manually check the cluster's status.

   Note: The Oracle service group is offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically. The service group is started later in the process.

   ---

6   Relink the SF Oracle RAC libraries with Oracle by choosing the option **Relink Oracle Database Binary** from the program menu.

7   Bring the Oracle database service group online.

   ■ If VCS manages the Oracle database:

      # **hagrp -online *oracle_group* -sys *node_name***

   ■ If VCS does not manage the Oracle database:

      # **srvctl start database -d *db_name***

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

9   Start all applications that VCS does not manage. Use native application commands to start the applications.

10 If the boot disk is encapsulated, reboot the first sub-cluster's system.

11 Before you proceed to phase 2, complete step 1 to 10 on the second subcluster.

12 ■ If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online when VCS starts:

      # **haconf -makerw**
      # **hagrp -modify oracle_group AutoStart 1**
      # **haconf -dump -makero**

   ■ If VCS does not manage the Oracle database, change the management policy for the database to automatic:

```
$ srvctl modify database -d db-name -y AUTOMATIC
```

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

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

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

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

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

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

3   The installer loads the new kernel modules.

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

5   Verify the cluster's status:

    ```
    # hastatus -sum
    ```

6   If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1,
    make sure that you upgraded all application clusters to version 5.1 SP1. Then,
    upgrade VCS or SFHA on the CP server systems.

    For instructions to upgrade VCS or SFHA on the CP server systems, see the
    VCS or SFHA installation guide.

Verifying software versions

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

```
# pkginfo -l pkgname
```

Uninstalling version 5.1 SP1 RP4

This chapter includes the following topics:

-

-

-

About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP4

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

Note: The version of VRTSaslapm doesn't change after rolling back.

Roll back of version 5.1 SP1 RP4 to the 5.1 SP1, 5.1SP1RP1, 5.1SP1RP2, 5.1SP1RP3 release is supported for the following products:

- Storage Foundation and High Availability (SFHA)

- Veritas Storage Foundation Cluster File System (SFCFS)

- Veritas Storage Foundation for Oracle RAC (SF for Oracle RAC)

- Veritas Cluster Server (VCS)

- Symantec VirtualStore (SVS)

- Dynamic Multi-Pathing (DMP)

Rolling back using the uninstallrp script

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

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

Rolling back 5.1 SP1 RP4 on an SFRAC cluster

Use the following procedure to roll back on an SFRAC cluster to the previous version.

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

    # **hagrp -offline oracle_group -sys *node_name***

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

    $ **srvctl stop database -d *db_name***

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

    ■   For 10gR2 or 11gR1:

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

    ■   For 11gR2:

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

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

    ■   Using native application commands, stop the applications that use CVM or CFS on all nodes.

    ■   Verify that no processes use the CFS mount point:

    # **fuser -c mount_point**

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

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

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

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

  ```
  # umount mount_point
  ```

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

On any node execute following command to stop VCS:

```
# hastop -all
```

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

- Using native application commands, stop the applications that use VxVM or VxFS.

- Verify that no processes use the VxFS mount point:

  ```
  # fuser -c mount_point
  ```

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

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

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

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

  ```
  # umount mount_point
  ```

8   Run the uninstallrp command, type:

```
# ./uninstallrp nodeA nodeB nodeC...
```

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

Rolling back 5.1 SP1 RP4 on a Veritas Cluster Server

Use the following procedure to roll back on a VCS cluster to the previous version.

1   Log in as superuser.

2   To rollback VCS on Solaris 10 systems that run non-global zones, make sure
that all non-global zones are booted into the running state on each node.

      Run the following command to get the state of non-global zones:

```
# zoneadm list -cv
```

3   Start the rollback. Enter:

```
# /opt/VRTS/install/uninstallrp node1 node2 ... nodeN
```

4   After the rollback, review the log files for any issues.

Rolling back 5.1 SP1 RP4 on an SFHA cluster

Use the following procedure to roll back on a SFHA cluster to the previous version.

1   Log in as superuser.

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

3   To rollback SFHA on Solaris 10 systems that run non-global zones, make sure
that all non-global zones are booted into the running state on each node.

      Run the following command to get the state of non-global zones:

```
# zoneadm list -cv
```

4   On any nodes in the cluster, make the VCS configuration read only. Enter:

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

5   Stop VCS.

      To stop applications, unmount the VxFS file systems and stop VxVM volumes
that are managed by VCS. Enter:

```
# hastop -all
```

6   Stop all the applications that are using VxFS files systems and VxVM volumes
that are not managed by VCS.

      Use the applications' native commands to stop applications.

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

```
# df -F vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# df -F vxfs
```

10 If any VxFS file systems are in use, on each node in the cluster, stop I/Os on the file systems and unmount all the VxFS file systems. Enter:

```
# umount /filesystem
```

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

- Stop all the applications that are involved in the 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
  ```

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

12 Stop activity to all VxVM volumes that are not managed by VCS. For example, stop any applications such as databases that access the volumes, and unmount any file systems that you create on the volumes.

Use applications' native commands to stop the applications.

13 On each node, stop all VxVM volumes.

For each disk group that is not managed by VCS, enter:

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

To verify that no volumes remain open, enter:

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

14 Deport all the disk groups that are not managed under VCS. Enter:

```
# vxdg deport diskgroup
```

15 If required, apply the OS kernel patches.

See Oracle's documentation for the procedures.

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

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

If the VEA service is running, stop it:

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

17 Start the rollback.

```
# /opt/VRTS/install/uninstallrp [-rsh] node1 node2 ... nodeN
```

Review the output.

18 Bring service groups online on each node. Enter:

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

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

20 Import all the disk groups that are not managed by VCS. Enter:

```
# vxdg import diskgroup
```

21 Restart all the volumes for each disk group that is not managed by VCS. Enter:

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

22 If you stopped any RVGs in step 11, restart each RVG. Enter:

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

23 Remount all VxFS file systems on all nodes that are not managed by VCS. Enter:

```
# mount -F vxfs blockdevice mountpoint
```

24 Remount all Storage Checkpoints on all nodes:

```
# mount -F vxfs -o ckpt=ckpt_name blockdevice checkpoint_name
```

25 Start all applications which are using VxFS files systems that are not managed by VCS.

Use applications' native commands to start the applications.

Rolling back 5.1 SP1 RP4 on an SFCFS cluster

Use the following procedure to roll back on a SFCFS cluster to the previous version.

1 Log in as superuser.

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

3 To roll back VCS on Solaris 10 systems that run non-global zones, make sure that all non-global zones are booted into the running state on each node.

Run the following command to get the state of non-global zones:

```
# zoneadm list -cv
```

4 On any nodes in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

5 Freeze HA service group operations on any node:

```
# hagrp -freeze groupname -persistent
```

6 Make the VCS configuration read only:

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

7 Stop VCS.

To stop applications, unmount the VxFS or CFS file systems and stop VxVM
or CVM volumes that are managed under VCS:

```
# hastop -all
```

8 Stop all the applications that are using CFS files systems and VxVM volumes
that are not managed by VCS.

Use the applications' native commands to stop applications.

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

```
# df -F vxfs
```

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

```
# cfsumount /checkpoint_name
```

11 On each node, check if any VxFS file systems are mounted:

```
# df -F vxfs
```

12 If any VxFS or CFS file systems are in use, on each node in the cluster, stop
I/Os on the file systems and unmount all the VxFS or CFS file systems:

```
# cfsumount /filesystem
```

or

```
# umount /filesystem
```

13 If you have created any Veritas Volume Replicator (VVR) replicated volume
groups (RVG) on your system, perform the following steps in order:

- Stop all the applications that are involved in the 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
  ```

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

  ---

14 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that you create on the volumes.

Use applications' native commands to stop the applications.

15 On each node, stop all the VxVM volumes for each disk group.

For each disk group that is not managed by VCS, enter:

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

Verify that no volumes remain open:

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

16 Deport all the disk groups which are not managed under VCS:

```
# vxdg deport diskgroup
```

17 On each node, stop the VCS command server:

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

`pid_of_CmdServer` is the process ID of CmdServer.

18 If required, apply the OS kernel patches.

See Oracle's documentation for the procedures.

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

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

If the VEA service is running, stop it:

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

20 Start the rollback.

```
# /opt/VRTS/install/uninstallrp [-rsh] node1 node2 ... nodeN
```

Review the output.

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

22 Make the VCS configuration writable again on any nodes

```
# haconf -makerw
```

23 Unfreeze HA service group operations on any nodes:

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

24 Make the configuration read only:

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

25 Bring the CVM service group online on each node:

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

26 Restart all the volumes for each disk group:

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

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

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

28 Remount all VxFS or CFS file systems on all nodes. Enter

```
# cfsmount /mountpoint
```

or

```
# mount -F fstype blockdevice mountpoint
```

29 Remount all Storage Checkpoints on all nodes. Enter:

```
# cfsmount /checkpoint_name
```

or

```
# mount -F vxfs -o ckpt=ckpt_name blockdevice mountpoint
```

30 Start all applications that are using VxFS or CFS files systems that are not managed by VCS.

Use applications' native commands to start the applications.

Rolling back manually

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

- Rolling back Storage Foundation or Storage Foundation and High Availability manually

- Rolling back Storage Foundation Cluster File System manually

- Rolling back Storage Foundation for Oracle RAC manually

- Rolling back Veritas Cluster Server manually

- Rolling back Symantec VirtualStore manually

- Rolling back Dynamic Multi-Pathing manually

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

Rolling back Storage Foundation or Storage Foundation and High Availability manually

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

To roll back SF or SFHA

1   Log in as superuser.

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

3   Unmount all Storage Checkpoints and file systems:

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

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

    ```
    # df -v /
    ```

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

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

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

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

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

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

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

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

df -F vxfs

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

umount */filesystem*

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the vxrvg stop command to stop each RVG individually:

  # **vxrvg -g *diskgroup* stop *rvg_name***

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

  # **vxrlink -g *diskgroup* status *rlink_name***

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

  ---

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

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

vxvol -g *diskgroup* stopall

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

vxprint -Aht -e v_open

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

- For Solaris 9:

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

- For Solaris 10:

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

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

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

11 Unmount /dev/odm:

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

12 Unload the ODM module:

- For Solaris 9:

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

- For Solaris 10:

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

13 Unload the cluster fencing (vxfen) module:

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

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

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SF 5.1 SP1 RP4 patches.

■ Get the list of 5.1 SP1 RP4 patches, type:

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

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

```
# patchrm 143287-07
```

Rolling back Storage Foundation Cluster File System manually

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

To roll back SFCFS or SFCFS HA manually

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# df -v /
```

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

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

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

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

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

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

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

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

```
# df -F vxfs
```

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

```
# umount /filesystem
```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the vxrvg stop command to stop each RVG individually:

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

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

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

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

---

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

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

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

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

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

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

    ■ For Solaris 9:

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

    ■ For Solaris 10:

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

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

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

11 Unmount /dev/odm:

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

12 Unload the ODM module:

    ■ For Solaris 9:

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

- For Solaris 10:

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

13 Unload the cluster fencing (vxfen) module:

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

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

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the SFCFS 5.1 SP1 RP4 patches.

- Get the list of 5.1 SP1 RP4 patches, type:

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

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

Rolling back manually

```
# patchrm 143287-07
```

Rolling back Storage Foundation for Oracle RAC manually

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

To roll back SF for Oracle RAC manually

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

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

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

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

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

    ■ For 10gR2 or 11gR1:

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

    ■ For 11gR2:

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

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

    ■ Using native application commands, stop the applications that use CVM or CFS on all nodes.

    ■ Verify that no processes use the CFS mount point:

      ```
      # fuser -c mount_point
      ```

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

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

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

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

      ```
      # umount mount_point
      ```

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

On any node execute following command to stop VCS:

```
# hastop -all
```

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

- Using native application commands, stop the applications that use VxVM or VxFS.

- Verify that no processes use the VxFS mount point:

```
# fuser -c mount_point
```

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

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

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

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

```
# umount mount_point
```

8 To stop the process, type:

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

9 Remove the SF for Oracle RAC 5.1 SP1 RP4 patches.

- Get the list of 5.1 SP1 RP4 patches, type:

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

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

```
# patchrm 143287-07
```

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

11 Reboot the nodes point.

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

Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 5.1 SP1 RP4 to VCS 5.1 SP1 RP1, VCS 5.1 SP1 RP2 or VCS 5.1 SP1 PR1 on your cluster manually. To uninstall VCS, see the *Veritas Cluster Server Installation Guide*.

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

To roll back VCS manually

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

    # **hagrp -state**

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

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

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

    # **haconf -makerw**

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

    # **hagrp -freeze** *service_group* **-persistent**

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

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

    # **haconf -dump -makero**

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

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

7 Shut down VCS. On any node, type:

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

8 Shut down CmdServer. On each node, type:

/opt/VRTSvcs/bin/CmdServer -stop

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

/sbin/gabconfig -a

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

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

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

  # **zoneadm list -icv**

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

  # **zoneadm -z** *zone* **boot**

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

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

---

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

/sbin/vxfenconfig -U

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

- Identify the vxfen kernel module, for example:

  # **modinfo|grep vxfen**

- Unload vxfen using the module number.

  # **modunload -i 210**

13 Unconfigure GAB. On each node, type:

    # **/sbin/gabconfig -U**

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

- Identify the GAB kernel module. For example:

    # **modinfo | grep gab**

- Unload GAB using the module number:

    # **modunload -i 149**

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

- Type:

    # **/sbin/lltconfig -U**

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

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

- Identify the LLT kernel module. For example:

    # **modinfo | grep llt**

- Unload LLT using the module number:

    # **modunload -i 147**

17 Remove the VCS 5.1 SP1 RP4 patches. On each node, perform the following steps:

- Get the list of 5.1 SP1 RP4 patches, type:

    # **./installrp -listpatches**

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

    # **patchrm 143287-07**

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

    # **showrev -p | grep VRTS**

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

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

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

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

21 After VCS has started, perform the following steps:

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

```
# hastatus -summary
```

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

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

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

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

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

where *system* is the node name.

Rolling back Symantec VirtualStore manually

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

To roll back SVS manually

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# df -v /
```

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

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

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

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

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

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

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

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

```
# df -F vxfs
```

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

```
# umount /filesystem
```

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

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

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

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

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

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

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

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

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

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

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

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

■ For Solaris 9:

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

■ For Solaris 10:

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

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

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

11   Unmount /dev/odm:

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

12   Unload the ODM module:

■ For Solaris 9:

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

- For Solaris 10:

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

13 Unload the cluster fencing (vxfen) module:

- For Solaris 9:

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

- For Solaris 10:

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

14 Stop GAB and LLT in the following order:

  For Solaris 9:

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

  For Solaris 10:

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

15 Check if the VEA service is running:

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

  If the VEA service is running, stop it:

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

16 Remove the SVS 5.1 SP1 RP4 patches.

- Get the list of 5.1 SP1 RP4 patches, type:

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

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

```
# patchrm 143287-07
```

Rolling back Dynamic Multi-Pathing manually

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

To roll back DMP manually

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# df -v /
```

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

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

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

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

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

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

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

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

       # **df -F vxfs**

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

       # **umount** */filesystem*

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

       ■   Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

       ■   Use the vxrvg stop command to stop each RVG individually:

           # **vxrvg -g** *diskgroup* **stop** *rvg_name*

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

           # **vxrlink -g** *diskgroup* **status** *rlink_name*

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

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

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

       # **vxvol -g** *diskgroup* **stopall**

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

       # **vxprint -Aht -e v_open**

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

       ■   For Solaris 9:

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

■ For Solaris 10:

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

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

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

11 Unmount /dev/odm:

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

12 Unload the ODM module:

■ For Solaris 9:

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

■ For Solaris 10:

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

13 Unload the cluster fencing (vxfen) module:

■ For Solaris 9:

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

■ For Solaris 10:

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

14 Stop GAB and LLT in the following order:

For Solaris 9:

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

For Solaris 10:

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

15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

16 Remove the DMP 5.1 SP1 RP4 patches.

■ Get the list of 5.1 SP1 RP4 patches, type:

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

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

```
# patchrm 143287-07
```

Post-upgrade tasks for Oracle RAC High Availability IP (HAIP) support on Oracle RAC 11.2.0.2 and later versions

This appendix includes the following topics:

- Configuring private IP addresses for Oracle RAC 11.2.0.2 and later versions
- Configuring the CSSD resource using the Veritas product script-based installer

Configuring private IP addresses for Oracle RAC 11.2.0.2 and later versions

Private IP addresses are required by Oracle RAC to provide communication between the cluster nodes. Depending on your private network configuration, you may need one or more IP addresses. For Oracle RAC 11.2.0.2 and later versions, Symantec recommends the use of link aggregation, NIC bonding, or Oracle High Availability IP (HAIP) to provide redundant, dedicated private communication.

Note: IPv6 addresses are not supported in this release.

Table A-1 lists the available options for configuring the private network for Oracle RAC. Use one of the following options to configure the private network.

Table A-1   Options for configuring the private network for Oracle RAC 11.2.0.2 and later

Option	Description
Link aggregation or NIC bonding	See "Using link aggregation or NIC bonding" on page 234.
Oracle High Availability IP (HAIP)	See "Using Oracle High Availability IP (HAIP)" on page 234.

Using link aggregation or NIC bonding

To configure link aggregation or NIC bonding, see the operating system documentation.

Note: Ensure that a link configured under an aggregated link or NIC bond is not configured as a separate LLT link.

When LLT is configured over a bonded interface, perform one of the following steps to prevent GAB from reporting jeopardy membership:

- Configure an additional network interface under LLT in addition to the bonded NIC.
- Add the following line in the /etc/llttab file:

  `set-dbg-minlinks <n>`

  Where <n> is the number of links configured under an aggregated link. Restart LLT.

  # `svcadm enable llt`

Using Oracle High Availability IP (HAIP)

Starting with Oracle RAC 11.2.0.2, Oracle introduced the High Availability IP (HAIP) feature for supporting IP address failover. The purpose of HAIP is to perform load balancing across all active interconnect interfaces and fail over existing non-responsive interfaces to available interfaces. HAIP has the ability to activate a maximum of four private interconnect connections. These private network adapters

can be configured during the installation of Oracle Grid Infrastructure or after the installation using the `oifcfg` utility.

Oracle Grid Infrastructure currently creates an alias IP (known as virtual private IP) on a private network interface using the 169.254.*.* subnet for HAIP. However, if that subnet range is already in use, Oracle Grid Infrastructure will not attempt to use it. Therefore, to enable HAIP using the Oracle Grid Infrastructure installer, configure all private interfaces with static private IP addresses before starting the Grid installation. During the installation of Oracle Grid Infrastructure, select these interfaces as 'private' interfaces. It is recommended to use all the LLT links as cluster interconnects.

You can configure HAIP in one of the following ways:

SF Oracle RAC script-based installer   See "Configuring static IP addresses for HAIP using the Veritas product script-based installer" on page 235.

Manual   See "Configuring static IP addresses for HAIP manually" on page 237.

Configuring static IP addresses for HAIP using the Veritas product script-based installer

The high availability IP address is managed by Oracle Grid Infrastructure. SF Oracle RAC supports the configuration of multiple network interfaces for HAIP. One static IP address is required for each network interface. All nodes must use a common network interface for the private interconnects in the cluster.

Review the following information before you configure HAIP:

- Do not use the public network interface for Oracle HAIP configuration.

- All LLT links are used for configuring HAIP.

- The same network interfaces must be used for private interconnect on each node of the cluster.

- The private IP addresses must be added to /etc/hosts on all nodes before installing Oracle Clusterware. If you choose the installer to add IP addresses in /etc/hosts then

To configure static IP addresses for HAIP using the Veritas product script-based installer

1   From the SF Oracle RAC menu, enter **3** to select the option **Oracle Network Configuration**.

```
1) Create Oracle Users and Groups
2) Create Storage for OCR and Voting disk
3) Oracle Network Configuration
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-4,b,q] (1) 3
```

2   Enter **1** to select the option **Configure private IP addresses (For HAIP Configuration) - For Oracle RAC 11.2.0.2 and later**.

```
1) Configure private IP addresses (For HAIP Configuration)
- For Oracle RAC 11.2.0.2 and later
2) Configure private IP addresses (PrivNIC Configuration)
- For Oracle RAC prior to 11.2.0.2
3) Configure private IP addresses (MultiPrivNIC Configuration)
- For Oracle RAC prior to 11.2.0.2
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: 1-3,b,q (1) 1
```

The installer discovers available LLT links.

3   Enter **y** to add the IP addresses to the `/etc/hosts` file.

```
Do you want the Installer to add IP addresses in /etc/hosts
file? [y,n,q] (y)
```

Enter **n** if you choose to add the IP addresses to the file manually. Go to step 5.

4 Perform this step only if you enabled the installer to add the IP address to the `/etc/hosts` file in the previous step. Provide the private IP address and the private node name for the IP address that must be added to the file.

Note: All IP addresses must be in the same subnet, failing which Oracle Clusterware/Grid Infrastructure will not be able to communicate properly across the nodes. If the private IP address entries are already present in the `/etc/hosts` file on one of nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster. Only those IP addresses that are not already present in the file on any node are added.

```
Enter the IP for bge1 on sys1: [b] 192.168.12.1
Enter Hostname alias for the above IP address: [b] sys1-haip1
Enter the private IP for bge2 on sys1: [b] 192.168.13.1
Enter Hostname alias for the above IP address: [b] sys1-haip2
Enter the IP for bge1 on sys2: [b] 192.168.12.2
Enter Hostname alias for the above IP address: [b] sys2-haip1
Enter the private IP for bge2 on sys2: [b] 192.168.13.2
Enter Hostname alias for the above IP address: [b] sys2-haip2
```

5 Enter the netmask information for the private network:

```
Enter the Netmask for private network: [b] (255.255.255.0)
```

The SF Oracle RAC installer now displays the configuration information.

6 Enter **y** to review and confirm the configuration information. The installer updates the `/etc/hosts` file (if you chose an installer-based update).

7 If you chose to add the IP address information to the `/etc/hosts` file manually, proceed to update the file as described in the following procedure.

Configuring static IP addresses for HAIP manually

You can configure a static IP addresses for HAIP manually using native operating system commands. See the operating system documentation for the appropriate commands.

For example, to configure a static IP address:

Edit the file `/etc/hostname.interface_name ip-address`

Configuring the CSSD resource using the Veritas product script-based installer

Configure the CSSD resource using the Veritas product installer if the OCR and voting disk storage is configured on CVM raw volumes or CFS.

Note: If the OCR and voting disk storage is configured on ASM disk groups, configure the CSSD resource manually.

The installer performs the following configuration tasks:

■ Adds the CSSD resource to the VCS configuration in the cvm group.

Note: If the CSSD resource already exists, the installer enables reconfiguration of the resource by deleting the existing resource.

■ For Oracle RAC 10g Release 2 and 11.2.0.1 versions: Sets the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resource that manages the private IP address for Oracle Clusterware/Grid Infrastructure.

■ Sets the dependency of the CSSD resource on the CFSMount or CVMVolDg resources that manage OCR and voting disk.

■ Disables automatic startup of Oracle Clusterware/Grid Infrastructure.

■ Enables the CSSD resource and saves the new configuration.

To configure the CSSD resource using the Veritas product script-based installer

1   Start the Veritas product installer, if it is not already running. Select the option **Post Oracle Installation Tasks**.

```
    1)   Configure SFRAC sub-components
    2)   SFRAC Installation and Configuration Checks
    3)   Prepare to Install Oracle
    4)   Install Oracle Clusterware/Grid Infrastructure and Database
    5)   Post Oracle Installation Tasks
    6)   Exit SFRAC Configuration
    Choose option: [1-6,q] (1) 5
```

2   Select the option **Configure CSSD agent**.

3   Enter the full path to the Oracle Clusterware/Grid Infrastructure home directory.

```
Enter absolute path of Oracle Clusterware/Grid Infrastructure
Home directory: [b]
```

The agent uses this information to locate the clusterware process binaries. If the information is not provided, the agent reads the information from the Oracle configuration file.

The installer verifies that Oracle Clusterware/Grid Infrastructure is running on all the nodes. It also checks that the PrivNIC/MultiPrivNIC (for supported versions), OCR, and voting disk resources are configured in the same service group.

4   Press **Return** to continue with the configuration. The installer reads the resource and group mappings for the CSSD agent from the VCS configuration file and displays the information.

5   Enter **y** to continue with the configuration. Review the messages as the installer configures the CSSD agent and sets the appropriate dependencies.

6   Press **Return** to return to the installer menu.

At the end of the configuration, the installer disables automatic startup of Oracle Clusterware/Grid Infrastructure.