Veritas Storage Foundation™ and High Availability Solutions Release Notes

HP-UX 11i v3

5.1 Service Pack 1 Rolling Patch 1

Veritas Storage Foundation™ and High Availability Solutions Release Notes

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

Product version: 5.1 SP1 RP1

Document version: 5.1SP1RP1.0

Legal Notice

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

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

About Symantec Connect

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

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

Contents

About this release

This chapter includes the following topics:

Introduction

This document provides information about the products in Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Rolling Patch 1 (5.1 SP1 RP1). Symantec strongly recommends installing the 5.1 SP1 Rolling Patch 1 immediately after installing Veritas Storage Foundation and High Availability Solutions 5.1 SP1.

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

http://www.symantec.com/docs/TECH144835

Review this entire document before installing and upgrading your Veritas Storage Foundation and High Availability product.

For further details, depending on the product for which you want to install this Rolling Patch, refer to one of the following release notes:

- *Veritas Storage Foundation Release Notes (Version 5.1 SP1)*
- *Veritas Cluster Server Release Notes (Version 5.1 SP1)*
- *Veritas Storage Foundation Cluster File System Release Notes (Version 5.1 SP1)*
- *Veritas Dynamic Multi-Pathing Release Notes (Version 5.1 SP1)*
- *Veritas Storage Foundation for Oracle RAC Release Notes (Version 5.1 SP1)*

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

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation High Availability (SFHA)
- Veritas Storage Foundation Cluster File System (SFCFS)
- Veritas Cluster Server (VCS)
- Veritas Dynamic Multi-Pathing (DMP)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)

Changes in this release

This section lists the changes introduced in Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1.

- Changes related to upgrade:
  See "Changes related to upgrade" on page 12.
- Changes in Veritas Storage Foundation High Availability:
  See "Changes in Veritas Storage Foundation High Availability" on page 12.
- Changes related to VCS agent for DB2:
  See "Changes related to VCS agent for DB2" on page 13.

Changes related to upgrade

This release supports rolling upgrade.

Changes in Veritas Storage Foundation High Availability

This release supports HP Integrity Virtual Machines (IVM) 4.3.

Changes related to VCS agent for DB2

This section contains the VCS agent for DB2 changes in this release.

IMF support for DB2 agent

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

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

- IMF notification module functions
- Administering the AMF kernel driver

Before enabling the agent for IMF

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

How the DB2 agent supports intelligent resource monitoring

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

Agent functions for the IMF functionality

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

imf_init

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

imf_getnotification

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

imf_register

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

Attributes that enable IMF

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

IMF

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

Mode

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

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

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

Type and dimension: integer-association

Default: 0

MonitorFreq

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

Default: 1

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

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

RegisterRetryLimit

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

Default: 3

IMFRegList

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

Type and dimension: string-vector

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

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

System requirements

For information on system requirements, refer to the product documentation for Veritas Storage Foundation and High Availability Solutions 5.1 SP1.

Note: This release requires that Version 5.1 SP1 is installed on your systems.

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.

For the latest information on supported hardware, visit:

https://sort.symantec.com/documents

Symantec recommends installing the latest HP-UX patches from HP.

List of patches

This section lists the patches included in this release.

- Veritas Storage Foundation:
  See "Veritas Storage Foundation patches in 5.1 SP1 RP1" on page 16.
- Veritas Cluster Server:
  See "Veritas Cluster Server patches in 5.1 SP1 RP1" on page 17.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System patches in 5.1 SP1 RP1" on page 18.
- Veritas Storage Foundation for Oracle RAC
  See "Veritas Storage Foundation for Oracle RAC patches in 5.1 SP1 RP1" on page 19.

Veritas Storage Foundation patches in 5.1 SP1 RP1

Table 1-1 lists the Veritas Storage Foundation patches included in this release.

Table 1-1 Veritas Storage Foundation patches

Patch	Version	Description
PHCO_42229	1.0	VRTS 5.1 SP1RP1 VRTSvxfs Command Patch (Veritas File System)
PHKL_42228	1.0	VRTS 5.1 SP1RP1 VRTSvxfs Kernel Patch (Veritas File System)
PHCO_42245	1.0	VRTS 5.1 SP1RP1 VRTSvxvm Command Patch (Veritas Volume Manager)

Table 1-1 Veritas Storage Foundation patches *(continued)*

Patch	Version	Description
PHKL_42246	1.0	VRTS 5.1 SP1RP1 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHCO_42093	1.0	VRTS 5.1 SP1RP1 VRTSdbed Command Patch (Veritas Storage Foundation for Databases Tools)
PHCO_42318	1.0	VRTS 5.1 SP1RP1 VRTSsfmh Command Patch (Veritas Operations Manager)
PHCO_42182	1.0	VRTS 5.1 SP1RP1 VRTSob Command Patch (Veritas Enterprise Administrator)
PHCO_42213	5.10.0.13	VRTS 5.1 SP1RP1 VRTSperl Command Patch (Perl Redistribution)

Veritas Cluster Server patches in 5.1 SP1 RP1

Table 1-2 lists the Veritas Cluster Server patches included in this release.

Table 1-2 Veritas Cluster Server patches

Patch	Version	Description
PHCO_42254	1.0	VRTS 5.1 SP1RP1 VRTSvxfen Command Patch
PHKL_42252	1.0	VRTS 5.1 SP1RP1 VRTSvxfen Kernel Patch
PVCO_03929	1.0	VRTS 5.1 SP1RP1 VRTSamf Command Patch
PVCO_03930	1.0	VRTS 5.1 SP1RP1 VRTScps Command Patch
PVCO_03931	1.0	VRTS 5.1 SP1RP1 VRTSvcs Command Patch
PVCO_03933	1.0	VRTS 5.1 SP1RP1 VRTSvcsag Command Patch
PVCO_03934	1.0	VRTS 5.1 SP1RP1 VRTSvcsea Command Patch

Veritas Storage Foundation Cluster File System patches in 5.1 SP1 RP1

Table 1-3 lists the Veritas Storage Foundation Cluster File System patches included in this release.

Table 1-3 Veritas Storage Foundation Cluster File System patches

Patch	Version	Description
PHCO_42229	1.0	VRTS 5.1 SP1RP1 VRTSvxfs Command Patch (Veritas File System)
PHKL_42228	1.0	VRTS 5.1 SP1RP1 VRTSvxfs Kernel Patch (Veritas File System)
PHCO_42245	1.0	VRTS 5.1 SP1RP1 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_42246	1.0	VRTS 5.1 SP1RP1 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHCO_42093	1.0	VRTS 5.1 SP1RP1 VRTSdbed Command Patch (Veritas Storage Foundation for Databases Tools)
PVCO_03931	1.0	VRTS 5.1 SP1RP1 VRTSvcs Command Patch (Veritas Cluster Server)
PVCO_03933	1.0	VRTS 5.1 SP1RP1 VRTSvcsag Command Patch (Veritas Cluster Server Bundled Agents)
PVCO_03934	1.0	VRTS 5.1 SP1RP1 VRTSvcsea Command Patch (Veritas Cluster Server Agent for Sybase)
PVCO_03932	1.0	VRTS 5.1 SP1RP1 VRTScavf Command Patch (Veritas Cluster Server Agents for Cluster File System)
PHKL_42342	1.0	VRTS 5.1 SP1RP1 VRTSglm Kernel Patch (Veritas Group Lock Manager)
PVCO_03929	1.0	VRTS 5.1 SP1RP1 VRTSamf Command Patch (Veritas Agent Monitoring Framework)
PVCO_03930	1.0	VRTS 5.1 SP1RP1 VRTScps Command Patch (Veritas Coordination Point Server)
PHCO_42254	1.0	VRTS 5.1 SP1RP1 VRTSvxfen Command Patch (Veritas I/O Fencing)

Table 1-3 Veritas Storage Foundation Cluster File System patches *(continued)*

Patch	Version	Description
PHKL_42252	1.0	VRTS 5.1 SP1RP1 VRTSvxfen Kernel Patch (Veritas I/O Fencing)
PHCO_42318	1.0	VRTS 5.1 SP1RP1 VRTSsfmh Command Patch (Veritas Operations Manager)
PHCO_42182	1.0	VRTS 5.1 SP1RP1 VRTSob Command Patch (Veritas Enterprise Administrator)
PHCO_42213	5.10.0.13	VRTS 5.1 SP1RP1 VRTSperl Command Patch (Perl Redistribution)

Veritas Storage Foundation for Oracle RAC patches in 5.1 SP1 RP1

The list of patches are the same as the patches mentioned in the following section of this guide:

See "Veritas Storage Foundation Cluster File System patches in 5.1 SP1 RP1" on page 18.

Fixed issues in this release

This section describes issues fixed in this release.

- Veritas Storage Foundation:
  See "Veritas Storage Foundation fixed issues" on page 19.
- Veritas Storage Foundation for Databases (SFDB) tools:
  See "Veritas Storage Foundation for Databases (SFDB) tools fixed issues" on page 34.
- Veritas Cluster Server:
  See "Veritas Cluster Server fixed issues" on page 34.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System fixed issues" on page 40.
- Veritas Storage Foundation for Oracle RAC:
  See "Veritas Storage Foundation for Oracle RAC fixed issues" on page 42.

Veritas Storage Foundation fixed issues

Table 1-4 lists the Veritas Volume Manager issues fixed in this release.

Table 1-4 Veritas Volume Manager fixed issues

Incident	Description

2492016	Multiple resize operations of Redundant Array of Inexpensive Disks (RAID5) or layered volumes may fail with the following message:
	`VxVM vxassist ERROR V-5-1-16092`
	`Volume TCv7-13263: There are other recovery activities.`
	`Cannot grow volume`

2491856	A Veritas Volume Replicator (VVR) primary node crashes while replicating in lossy and high latency network with multiple Transmission Control Protocol (TCP) connections.
2488042	A panic is triggered in the `vol_mv_commit_check()` function while accessing a Data Change Map(DCM) object.
2485288	The `vxpfto`(1M) command sets the Powerfail Timeout (PFTO) value on the wrong Veritas Volume Manager (VxVM) device.
2485278	In some cases, the error messages printed in the syslog file in the event of a master takeover failure are not enough to find out the root cause of the failure.
2485230	The `vxdisk`(1M) command displays the incorrect pubpath of an Extensible Firmware Interface (EFI) partitioned disk on the HP 11i v3 platform.

2484695	In a Storage Foundation environment running Veritas Extension for Oracle Disk Manager (ODM), Veritas File System (VxFS) and Volume Manager (VxVM), a system panic may occur with the following stack trace:
	`vol_subdisksio_done()`
	`volkcontext_process()`
	`oldiskiodone()`
	`voldmp_iodone()`
	`gendmpiodone()`

2484466	I/O of large sizes like such as 512 K and 1024 K hang in Cluster Volume Replicator (CVR).

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

Incident	Description

2483476	The vxdisksetup(1M) command fails on disks which have stale Extensible Firmware Interface (EFI) information and the following error message is displayed:
	`VxVM vxdisksetup ERROR V-5-2-4686`
	`Disk <disk name> is currently an EFI formatted disk.`
	`Use -f option to force EFI removal.`

2480006	The `vxdmpadm listenclosure` command hangs because of duplicate enclosure entries in the `/etc/vx/array.info` file.
2479746	In case of I/Os on volumes having multiple subdisks (for example, striped volumes), the system panics.
2477291	Shared Disk Group (DG) import or node join fails with Hitachi Tagmastore storage.
2442850	When the `vxesd` daemon is invoked by the device attach and removal operations in a loop, it leaves open file descriptors with the `vxconfigd`(1M) daemon.
2440351	The grow operation on a Data Change Object (DCO) volume may grow it into any 'site' without following the allocation requirements.
2440031	In a Storage Foundation environment, running both Veritas File System (VxFS) and Veritas Volume Manager (VxVM), a system panic may occur when I/O hints are being used. One such scenario is when Veritas Extension for Oracle Disk Manager (ODM) is used.
2436288	I/O hangs occur in a Clustered Volume Replicator (CVR) environment.
2436287	In a Cluster Volume Replicator (CVR) configuration, I/Os are issued from both the master node and the slave node. Rebooting the slave node leads to a reconfiguration hang.
2436283	The Cluster Volume Manager (CVM) reconfiguration takes 1 minute for each Replicated Volume Group (RVG) configuration.
2435050	After Veritas Volume Replicator (VVR) is configured, the `vxconfigd`(1M) daemon hangs on the primary site when trying to recover Storage Replicator Log (SRL) after a system or storage failure.
2428179	The Veritas Volume Manager's (VxVM) subdisk operation – `vxsd mv <source_subdisk> <destination_subdisk>` - fails on subdisks with sizes greater than or equal to 2TB.

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

Incident	Description
2423086	Disabling a controller of an A/P-G type array can lead to an I/O hang even when there are paths available for I/O.
2421491	On Veritas Volume Manager (VxVM) rooted systems, during a machine bootup, the `vxconfigd`(1M) command dumps core and the machine fails to boot.

2421100	The system panics with the following stack trace:
	`dmp_get_path_state()`
	`do_passthru_ioctl()`
	`dmp_passthru_ioctl()`
	`dmpioctl()`
	`ioctl()`

2417205	The `vxassist`(1M) command dumps core if the `/etc/default/vxassist` file contains the line `wantmirror=<ctlr|target|...>`.
2417184	Application I/O hangs on Replicated Volume Group (RVG) volumes when RVG log owner is being set on the node which takes over the master's role either as part of the `vxclustadm setmaster` command or as part of the original master leave.
2415577	Enclosure attributes such as I/O policy and recovery option do not persist across reboots.
2415566	When disks of size greater than 2TB are used and the device responds to Small Computer System Interface (SCSI) inquiry but fails to service I/O, data corruption can occur as the write I/O is issued at an incorrect offset.
2413908	Performing Dynamic Logical Unit Number (LUN) reconfiguration operations (adding and removing LUNs) can cause corruption in the DMP database. This may lead the `vxconfigd`(1M) daemon to dump core or trigger a system panic.
2413077	In Veritas Volume Replicator (VVR) environment, the `vol_rv_async_childdone`() panic occurs because of a corrupted primary node pending queue, which is the queue that manages the remote I/O requests.
2411053	If a Disk Group (DG) is imported with the reservation key, then during DG deport, several reservation conflict messages are seen.

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

Incident	Description
2411052	1) On suppressing the underlying path of a PowerPath controlled device, the disk goes into an error state. 2) The `vxdmpadm exclude vxvm dmpnodename=<emcpower#>` command does not suppress the Third-party Driver (TPD) devices.

2409212	While doing a cold/ignite Ignite installation on Veritas Volume Manager (VxVM) 11i v3 5.1 SP1, the following warning messages are seen on a setup with an Asymmetric Logical Unit Access (ALUA) array:
	`VxVM vxconfigd WARNING V-5-1-0`
	`ddl_add_disk_instr: Turning off NMP Alua mode`
	`failed for dmpnode 0xffffffff with ret = 13`

2408864	Some Dynamic Multi-pathing (DMP) I/O statistics records are lost from the per-cpu I/O statistics queue. Hence, the DMP I/O statistics reporting command displays incorrect data.
2408209	Data corruption can be observed on a Cross-platform Data Sharing (CDS) disk, whose capacity is more than 1 TB.
2405446	Enhancements are made to customize the private region I/O size based on the maximum transfer size of underlying disk.
2397663	If the cloned copy of a Disk group (DG) and a destroyed DG exist on a system, an import operation imports the destroyed DG, instead of the cloned one.
2390822	On the Veritas Volume Replicator (VVR) Secondary cluster, if there is an error in a Storage Replicator Log (SRL) disk, `vxconfigd` may hang in the transaction code path.
2390815	In Veritas Volume Replicator (VVR) environment, a panic occurs in the `vol_rv_mdship_srv_done`() function.
2390804	Veritas Volume Replicator (VVR) volume recovery hang occurs at the `vol_ru_recover_primlog_done`() function in a dead loop.
2389095	In the presence of Not-Ready (NR) devices, `vxconfigd`(1M), the Veritas Volume Manager (VxVM) configuration daemon, goes into the DISABLED mode after it is restarted.

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

Incident	Description
2386763	The Dynamic Multi-Pathing Administration operations, such as `vxdmpadm exclude vxvm dmpnodename=<daname>` and `vxdmpadm include vxvm dmpnodename= <daname>` trigger memory leaks in the heap segment of the Veritas Volume Manager configuration daemon (`vxconfigd`).
2384844	When the `vxvm-recover` script is executed manually, the duplicate instances of the Veritas Volume Manager (VxVM) daemons, such as `vxattachd`, `vxcached`, `vxrelocd`, `vxvvrsecdgd` and `vxconfigbackupd` are invoked. When a user tries to kill any of the daemons manually, the other instances of the daemons remain on the system.
2384473	The `vxcdsconvert`(1M) utility fails if the disk capacity is greater than or equal to 1 TB.
2383705	The following message is displayed after a Disk Group (DG) creation: `VxVM ERROR V-5-3-12240: GPT entries checksum mismatch`.
2382717	The Veritas Volume Manager (VxVM) volume creation utility, `vxassist`(1M), does not function as expected while creating volumes with the `logtype=none` option.
2382714	In the presence of Not-Ready devices, when the Small Computer System Interface (SCSI) inquiry on the device succeeds, and the open, read or write operations fail, the status of the paths to such devices continuously alters between ENABLED and DISABLED for every Dynamic Multi-Pathing (DMP) restore task cycle.
2382710	A Disk Group (DG) import operation can fail with Serial Split Brain (SSB) though SSB does not exist.
2382705	The `vxconfigd` daemon leaks memory while excluding and including a Third party Driver-controlled Logical Unit Number (LUN) in a loop. As a part of this, vxconfigd loses its license information.
2379034	In Veritas Volume Manager (VxVM), changing the name of the enclosure does not work for all the devices present in /etc/vx/darecs.
2377317	Veritas Volume Manager (VxVM) does not show all the discovered devices. The number of devices shown by VxVM is lesser than those shown by the Operating System (OS).
2364700	If space-optimized snapshots exist at a secondary site, Veritas Volume Replicator (VVR) leaks kernel memory.

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

Incident	Description
2360719	1) The vxconfigbackup(1M) command fails with the error: `ERROR V-5-2-3720 dgid mismatch` 2) The `-f` option for the `vxconfigbackup`(1M) command is not documented in the man page. .
2360419	The `vxrecover`(1M) command fails to recover the data volumes with associated cache volume.

2360415	The system panics with the following stack trace:
	`voldrl_unlog+0001F0`
	`vol_mv_write_done+000AD0`
	`volkcontext_process+0000E4`
	`voldiskiodone+0009D8`
	`voldmp_iodone+000040`

2357820	Veritas Volume Replicator (VVR) leaks memory due to unfreed vol_ru_update structure. The memory leak is very small. However, it can be considerable, if VVR runs for many days.
2357579	While detecting unstable paths, the system panics.
2353922	The uninitialization of Veritas Volume Manager (VxVM) Cross Data platform Sharing (CDS) disks of size greater than 1 TB fails on HP-UX/IA-64 platform.
2353464	Duplicate device names are observed for Not Ready (NR) devices, when Veritas Volume Manager configuration daemon (`vxconfigd`) is restarted (`vxconfigd -k`).
2353427	The `vxdmpadm include`(1M) command includes all the excluded devices along with the device given in the command.
2353425	In Veritas Volume Manager (VxVM), the `vxconfigd`(1M) daemon dumps core and loses the Disk Group (DG) configuration.
2353421	In a Cluster Volume Manager (CVM) environment, a node join to the cluster gets stuck and leads to a hang unless the join operation is stopped on the joining node (SLAVE) using the command `/opt/VRTS/bin/vxclustadm stopnode`.

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

Incident	Description
2353410	The system panics in the Dynamic Multi-Pathing (DMP) kernel module due to a kernel heap corruption while the DMP path failover is in progress.
2353404	The vxconfigd(1M) daemon consumes a lot of memory when the Dynamic Multi-Pathing (DMP) tunable `dmp_probe_idle_lun` is set to on. The `pmap` command on vxconfigd process shows continuous growing heaps.
2353403	The `vxdisk -o thin list` command displays the size as zero for thin Logical Unit Numbers (LUNs) of capacity greater than 2 TB.
2353328	The `vxconfigd`(1M) daemon dumps core when array side ports are disabled/enabled in a loop for some iterations.
2353327	When using disks of size greater than 2TB, data corruption can occur in case of a write operation.
2353325	In a Veritas Volume Replicator (VVR) environment, replication doesn't start if the Replication Link (Rlink) detach and attach operations are performed just after a Storage Replicator Log (SRL) overflow.
2349653	Data corruption is observed on Dynamic Multi-Pathing (DMP) devices with single path during storage reconfiguration (Logical Unit Number (LUN) addition/removal).
2337091	If a CLARiiON array is configured in failover mode ‘x’ through one host controller and as failover mode ‘y’ through a different host controller, then the `vxconfigd`(1M)command dumps core.
2328286	Initialization of Veritas Volume Manager (VxVM) Cross Data platform Sharing (CDS) disk layout fails on a disk of size greater than or equal to 1 TB.
2328268	On Veritas Volume Manager (VxVM) rooted setup with boot devices connected through Magellan interface card, the system hangs at early boot time, due to transient I/O errors.
2328219	The `vxconfigd`(1M) command leaks memory while reading the default tunables related to SmartMove (a Veritas Volume Manager(VxVM) feature).
2323999	If the root disk is under the control of Veritas Volume Manager (VxVM) and the `/etc/vx/reconfig.d/state.d/install-db` file exists, the system becomes unbootable.

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

Incident	Description

2316309	The following error messages are printed on the console during system boot up:
	`VxVM vxdisk ERROR V-5-1-534`
	`Device [DEVICE NAME]: Device is in use`

2256728	The `vxdg`(1M) import command hangs if called from a script with STDERR redirected.
2253269	The `vxdg`(1M) man page does not clearly describe the Disk Group (DG) import and destroy operations for the case in which the original DG is destroyed and cloned disks are present.
2248354	When Veritas Volume Replicator (VVR) is replicating over a Network Address Translation (NAT) based firewall, Replication Links (Rlinks) fail to connect for NAT configurations, resulting in replication failure.
2247645	Initialization of Veritas Volume Manager (VxVM) Cross Data platform Sharing (CDS) disk layout on a disk with size greater than or equal to 1 TB fails on HP-UX/IA-64 platform.
2241149	The `vxdg`(1M) move/split/join command may fail during high I/O load.
2234292	A diskgroup (DG) import fails with a non-descriptive error message when multiple copies (clones) of the same device exist and the original devices are either offline or not available.
2232829	With NetApp MetroCluster disk arrays, takeover operations (toggling of Logical Unit Number (LUN) ownership within NetApp filer) can lead to I/O failures on Veritas Volume Manager (VxVM) volumes.
2220064	The Volume Replicator Administrative Services (VRAS) `vradmind` daemon hangs on the Veritas Volume Replicator (VVR) secondary site.
2214184	In the Veritas Volume Replicator (VVR) environment, transactions on Replication link (Rlink) are not allowed during the Storage Replicator Log (SRL) to Data Change Map (DCM) flush.
2211971	On a system with heavy I/O load, the `dmpdaemon` requests 1 MB of continuous memory paging which slows down the system.
2204146	In a Campus Cluster environment, some disks are left detached and not recovered by the vxattachd(1M) daemon.

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

Incident	Description
2198041	When creating a space-optimized snapshot by specifying the cache object size either as a percentage of the volume size or the absolute size, the snapshot creation can fail with the error that the cache size does not align with Disk Group (DG) alignment.
2169348	During node reconfiguration in a Cluster Volume Manager (CVM) environment, the master node hangs with a lot of I/Os in the queue due to a node leave.

2163809	The internal testing utility, volassert, prints the following message:
	`Volume TCv1-548914:`
	`recover_offset=0, expected 1024`

Table 1-5 lists the Veritas File System issues fixed in this release.

Table 1-5 Veritas File System fixed issues

Incident	Description
2559801	The memory used by the Veritas File System (VxFS) internal buffer cache may grow significantly after 497 days of uptime, when LBOLT that is the global system variable that gives the current system time, wraps over.
2559601	A full fsck operation displays corrupted Inode Allocation Unit (IAU) headers.
2529356	An `f:vx_iget:1a` assert is seen in Veritas File System (VxFS) during an internal stress test.
2508164	Access to a file system may hang if the customer creates large number of shares with numerous user quotas.
2496959	Using the `vxtunefs`(1M) command, the `pdir_enable` tunable can be set to invalid values.
2494464	The `vx_ino_update:2` assert is hit during internal testing.
2486597	On a machine with severe inode pressure, multiple threads may wait on a Mutex in the `vx_ireuse_steal`() function.
2482337	A `kernel null pointer dereference` panic may occur in Veritas File System (VxFS).

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

Incident	Description

2480949	The system log file may contain the following error message on a multi-threaded environment with SmartTier.
	`UX:vxfs fsppadm: ERROR: V-3-26626:`
	`File Change Log IOTEMP and ACCESSTEMP`
	`index creation failure for /vx/fsvm with message`
	`Argument list too long`

2478325	The `fsck`(1M) command takes a long time to complete the intent log replay.

2478237	The following asserts are seen during internal stress and regression runs:
	`f:vx_do_filesnap:1b`
	`f:vx_inactive:2a`
	`f:xted_check_rwdata:31`
	`f:vx_do_unshare:1`

2427281	The `vxfs_fcl_seektime`() Application Program Interface (API) seeks to the first record in the File Change Log (FCL) file after a specified time. This API can incorrectly return an EINVAL (FCL record not found) error while reading the first block of the FCL file.
2427269	In Veritas File System (VxFS), truncating-up of new files using the file control command, `fcntl`(2), followed by a small write operation of 512 bytes results in an incorrect file size of 512 bytes.
2426039	A time limit is established for each file system which determines how long a user is allowed to exceed the soft limit. But currently, the user is allowed to exceed the soft limit on Veritas File System (VxFS) file system, even after the time limit is exceeded.
2413015	In Veritas File System (VxFS) with partitioned directory enabled (disk layout 8) and accessed through a read-only mount, in some cases, the directory listing lists less number of entries.
2413010	In Veritas File System (VxFS) with partitioned directory enabled (disk layout 8) and accessed through Network File System (NFS), directory listing lists less number of entries.
2412179	The quota usage gets set to ZERO when umount/mount is performed on the file system, though files owned by users exist. This issue may occur after some file creations and deletions.

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

Incident	Description
2412177	A user quota file corruption occurs when the DELICACHE feature in Veritas File System (VxFS) is enabled. The current inode usage of the user becomes negative after frequent file creations and deletions.
2412173	During an internal testing of write operations using direct I/O, the system panics with the following panic string: `pfd_unlock: bad lock state!`
2412029	When named streams are used in Veritas File System (VxFS), the system may panic.
2409792	In a system under severe memory pressure, the sequential read performance reduces to up to 20% of the original.
2403663	The `vxrestore`(1m) man page does not mention that the `vxrestore`(1m) command fails to restore dumps with block sizes greater than 63 when the `-b` option is not used.
2402643	The full `fsck`(1M) command with '-o full' option on Veritas File System (VxFS) performs a large directory index validation during pass2c. However, if the number of large directories is more, then this pass takes a lot of time.
2386483	Access to a file system hangs when creating a named attribute, due to a read/write lock being held exclusively and indefinitely. This causes a thread to loop in the `vx_tran_nattr_dircreate`() function.
2373565	The system may panic when the `fsadm`(1M) command with the `-e` option is run on a file system containing file level snapshots.
2371923	In Veritas File System (VxFS), the performance of the delete operation is affected.
2371921	The mkfs(1M) command fails to create a VxFS file system with Disk Layout Version 4 (DLV4).

2368788	When the vx_ninode variable is being tuned with a value less than (250*vx_nfreelists), the following message is displayed:
	`vmunix: ERROR: mesg 112: V-2-112:`
	`The new value requires changes to Inode table`
	`which can be made only after a reboot`

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

Incident	Description
2368738	If a file which has shared extents, has corrupt indirect blocks, then in certain cases the reference count tracking system can try to interpret this block and panic the system. Since this is an asynchronous background operation, this process is retried repeatedly on every file system mount and hence, a panic occurs every time the file system is mounted.
2360821	When retrieving information about the checkpoints using the command `fsckptadm-C blockinfo <pathname> <ckpt-name> <mountpoint>`, the command fails with error 6 (ENXIO) and the file system is disabled.
2360820	Users may sometimes get access denial message while accessing files in directories with Access Control List (ACL).
2341007	When a file is newly created, issuing `fsppadm query -a /mount_point` could show the incorrect IOTemp information.
2340839	Shortly after removing files in a file system, commands such as `df`(1M), which use the `statfs`() function, can take about 10 seconds to complete.
2340825	When the `fsdb_vxfs`(1M) command is used to look at the bmap of an ILIST file ("mapall" command), a large hole at the end of the ILIST file is wrongly reported.
2340817	The system may panic when performing File Change Log (FCL) commands like `getacl`(1), and `setacl`(1) on Veritas File System (VxFS).
2340813	The Veritas File System (VxFS) mmap I/O operation on HP-UX 11i v3 is slower than the same operation on HP-UX 11i v2.
2340802	The `top`(1m) command shows that after some directories are deleted, the file system daemon process (vxfsd) consumes a significant amount of CPU time.
2340799	In Veritas File System (VxFS), a panic occurs because of a NULL pointer in the `vx_unlockmap`() function.

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

Incident	Description

2340755	When an IO-BOX cell without any CPUs is brought online, the following message is logged in the syslog file:
	`vmunix: ERROR: mesg 112: V-2-112:`
	`The new value requires changes to Inode table`
	`which can be made only after a reboot`

2340741	The vxdump(1M) command may dump core while backing up layout 7 VxFS file system, if the files in the file system are getting accessed.
2329893	The vxfsstat(1M) command's vxi_bcache_maxkbyte variable counter shows the maximum memory available for buffer allocation. The maximum memory available for buffer allocation depends on the total memory available for buffer cache (buffers + buffer headers), which is "vx_bc_bufhwm" global. Therefore, the vxi_bcache_maxkbyte variable should never be greater than the vx_bc_bufhwm variable.
2320049	There is a requirement for a new option to specify fileset-inode pairs. Currently, it is not possible to specify an inode that is unique to the file system since inode numbers are reused in multiple filesets.
2320044	In Veritas File System (VxFS), the ncheck(1M) command with the -i option does not limit the output to the specified inodes.
2311490	When a hole is created in the file using Data Management Application Programming Interface (DMAPI), the dm_punch_hole() function can leave the file in a corrupted state.
2296277	While querying a mount point, the fsppadm(1M) command displays the message, Operation not applicable in the output.
2289610	The vxfsstat(1M) command does not reflect the change in the vx_ninode(5) tunable after the tunable is changed using the kctune(1M) command.
2289528	The fsppadm(1M) command which is used to query a file, returns invalid file access time and update time. The fsppadm(1M) command used to enforce the log can display invalid file size.
2280386	After upgrading from disk layout version 6 to 7, the fsadm(1M) command for defragmentation may show the bad file number' error on a VxFS file system.

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

Incident	Description
2275543	On a VxFS filesystem, `write`() system call hangs for more than 10 seconds causing critical applications to timeout.
2257904	The `df`(1M) command with the `-h` option takes 10 seconds to execute and reports an inaccurate free block count, shortly after a large number of files are removed.
2243063	When a file is created in a large directory, the system hangs.
2222244	A backup operation using Symantec NetBackup (NBU) may seem to progress slowly.
2169326	When a clone is mounted on a locally mounted file system, a size limit is assigned to the clone. If the clone exceeds this limit, then it is removed. If the files from the clone are being accessed at the time of the removal of the clone, then an assert may be triggered in the function `vx_idelxwri_off`() through the function `vx_trunc_tran`().

Table 1-6 lists the Veritas Perl Redistribution issue fixed in this release.

Table 1-6 Veritas Perl Redistribution fixed issue

Incident	Description
2255106	VRTSperl package swverify warning messages are logged in the swverify/swagent logs after SFHA 5.0 HP-UX 11i v2 is upgraded to SFHA 5.1SP1 HP-UX 11i v3 on the Itanium platform.

Table 1-7 lists the Veritas Operations Manager issue fixed in this release.

Table 1-7 Veritas Operations Manager fixed issue

Incident	Description
2182417	The memory used by the `vxdclid` process in Veritas Operations Manager (VOM) managed hosts increases over time.

Table 1-8 lists the Veritas Enterprise Administrator issue fixed in this release.

Table 1-8 Veritas Enterprise Administrator fixed issue

Incident	Description
2394915	The Veritas Enterprise Administrator (VEA) service (vxsvc) crashes and dumps core.

Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Table 1-9 lists the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-9 Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Incident	Description
2395194	The `vxdbd` daemon consumes excessive CPU resources.
2395173	The `vxdbd` daemon allows the use of ciphers but provides no way to configure the strength of the ciphers.
2361363	Running the `qio_convertdbfiles`(1m) command results in the following error: `/opt/VRTSdbed/bin/qio_convertdbfiles: Command not found.`

Veritas Cluster Server fixed issues

Table 1-10 lists the issues fixed 5.1 SP1 Rolling Patch 1

Table 1-10 Veritas Cluster Server fixed issues

Incident	Description
2406748	When AMF processes the offline registration request, stale data from earlier cycles causes an online process to be reported as offline.
2403851	An error in the code prevented the unloading of AMF module though the module was not being used.
2301731	There was a race condition between a system call and the AMF driver unconfiguration which causes the kernel to panic.
2330045	The RemoteGroup resource tries to connect to the remote cluster even after the agent has invoked its offline entry point. If the connection to the remote cluster is not available, the resource goes into the UNKNOWN state and prevents the service group from going offline

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

Incident	Description
2426663	The kernel driver of the Veritas fencing module (VxFEN) starts the vxfend process only when VxFEN is configured in a customized mode. However, when you change the fencing mode to ‘scsi3’ using vxfenswap utility, the VxFEN kernel driver fails to terminate the vxfend process.
2382559	Online Migration fails with the message `I/O fencing does not appear to be configured on node`
2382559	Oracle agents use /usr/lib in LD_PRELOAD before Oracle libraries and Oracle CC asked the customer to use Oracle Library path before /usr/lib.
2354932	When HAD is running in onenode mode, `hacli` command tries to send unicast messages to other systems(which are not a part of the cluster since HAD is running in onenode mode). This attempt to send unicast message to other systems causes HAD to coredump .
2407653	The module reference count on the filesystem registered with AMF for mount offline monitoring or mount online monitoring is not released when you forcefully unconfigure AMF using the `amfconfig -Uof` command. This causes extra reference counts on these modules to remain even after AMF driver is unloaded.
2394176	If you run the vxfenswap utility on a multinode VCS cluster, then after some time, the vxfenswap operation stalls and no output appears on the console. However, the console does not freeze (the system does not hang). If you run the `ps -ef | grep vxfen` command on every node, the output indicates that the ‘vxfenconfig -o modify’ process is running on some nodes, but it is not running at least on one node.
2372072	If the `hacf` command cannot get its current working directory, it logs an error. At this point the log object is not initialized.
2386326	The fencing module runs a SCSI3 query on disk to determine its serial number. The buffer size for the query is 96 bytes whereas the size of the output is much larger. Therefore, the serial number of the disk is truncated, and appears to be the same for all disks.
2417846	Cluster Manager (Java Console) does not encrypt Oracle agent attribute "DBAPword" strings.
2398807	In /opt/VRTSvcs/bin/vcsenv, Soft and Hard limit of file descriptors is set To 2048. If Hard limit is set to a higher value, then the Hard limit is overridden with the lower value(2048).

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

Incident	Description
2438261	Failed to perform online migration from scsi raw to scsi dmp policy
2253349	The IP agent makes use of IP address and Netmask value pair to perform online and offline operations. When the Netmask value on the interface is changed outside of VCS control, the VCS expected value of Netmask mismatches with the netmask value present on the device and hence offline operation fails.
2382493	Parent service group does not failover if it is dependent on an Online-Local-Firm parallel child service group.
2517333	The AMF driver tries to release the memory that it has not allocated. This causes the node to panic.
2423990	When you configure the User attribute, if you specify a user name that does not exist on a system, then the `getpwnam` command fails during online/offline entry points. As a result, the agent logs the above messages in the engine log.
2382592	In `hares -display`, there is a limit of 20 characters. Any attribute value greater than 20 characters is truncated. Hence 'Status' keys is not displayed as limit of 20 characters is exhausted by other keys like State, Msg, TS of ResourceInfo.
2276622	I/O fencing (vxfen) fails to start using coordinator disks from certain disk arrays. Even if you configure multiple coordinator disks, the component displays the following error message: `V-11-2-1003 At least three coordinator disks must be defined` If you run the SCSI-extended inquiry command `vxfenadm -i <disk_name>` on the disks, it reports same serial number for all the coordinator disks.
2399898	When you run the 'hagrp -switch' command, the VCS engine checks the state of a parent service group before switching a child service group. If more than one parent service group is ONLINE, the VCS engine rejects the command irrespective of the rigidity (soft or firm) or location (local, global, or remote) of the group dependency.
2366201	If VxFEN fails to get the UUID or serial number of one of the specified CPs, then VxFEN treats it as a fatal failure. The node cannot then join a cluster or start a cluster. As a result, every coordination point becomes a potential single point of failure, and compromises high availability (HA).

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

Incident	Description
2366701	The application agent checks for the existence of the Monitor program on a node. On nodes that cannot access the shared disk, this check fails. As a result, the agent marks the status of the application as UNKNOWN on such nodes. Further, if an application goes down on an active node, then the application cannot fail over to nodes where the status of the application is UNKNOWN. The Application agent can handle only the following set of values returned by the monitor program: 100 --> OFFLINE 101 to 110 --> ONLINE Any other value --> UNKNOWN If the monitor program returns "0" as success and "1" as failure, the Application agent reports the state of application as UNKNOWN.
2382452	The configure_cps.pl utility contains an irregular apostrophe character that causes the syntax error.
2439772	In case of network interruption, wide-area connector becomes unresponsive and is not able to receive or send any updates to the remote cluster. Even `wacstop` is not able to stop wac.
2382463	When Preferred Fencing is enabled, the least node weight that can be assigned is 1. Hence, VCS adds 1 to the value specified in FencingWeight attribute. If the value of FencingWeight is set to 10000 then VCS tries to set the node weight 10001 which fails since the maximum node weight that can be assigned is 10000.
2382583	When a CP server becomes inaccessible, the engine log does not provide sufficient information for a root-cause analysis (RCA).
2400485	When the Veritas fencing module (VxFEN) starts, it may encounter issues in reading coordination point information because the variable that tracks the fencing mode is incorrectly passed from the user land configuration files to the VxFEN kernel driver.
2426572	VCS assumes that the state of a resource is OFFLINE before it is probed. If the first monitor cycle of the agent returns OFFLINE, then the persistent resource is not reported as FAULTED.

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

Incident	Description
2330980	When you add a node to the SystemList of a group, the related agent must start monitoring resources from that group on the new node. So, the High Availability Daemon (HAD) module sends a snapshot (information related to the monitored resources in the group, including attribute values) to the agent on the new node. However, HAD also sends the snapshot to the existing nodes. Therefore, the agent framework may incorrectly modify certain attributes, and the agent may report an incorrect state of a resource on an existing node.
2367721	Virtual fire drill matches the output of the id command with the output of same command on the system where the resource is in online state. Some fields in this output differ when SELinux is enabled and the Virtual fire drill fails.
2330041	When a child service group is auto-started, the parallel parent service groups that have a online-global dependency are not auto-started.
2175599	VxFEN's user mode process, vxfend, uses a limited size buffer to store the snapshot of cluster membership. This buffer can only accommodate the snapshot of up to 33 nodes.
2382335	In a shared diskgroup that contains more than one disk, the `vxfentsthdw -g <diskgroup>` command fails to map a shared disk correctly to the nodes that share it.
2400330	When service group is manually switched, whyonlining parameter of PreOnline script is shown as "Manual".

2403782	Sybase agent scripts use an incorrect path for the `cat` command on the Linux platform. As a result, the operation to bring a Sybase resource online fails, and the following message appears as part of the command output:
	`Can't exec "/usr/bin/cat": No such file or directory at /opt/VRTSagents/ha/bin/Sybase/online line 244.`
	`Use of uninitialized value $encrypted_passwd in substitution (s///) at /opt/VRTSagents/ha/bin/Sybase/online line 245`

2438621	The MultiNICB agent compares the subnets of all the interfaces on a system, and in the above case, reports an incorrect resource state.

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

Incident	Description
2382460	If you configure the Veritas fencing module (VxFEN) in one of the following two ways, then you may not be able to distinguish certain important messages in the log file: ■ /etc/vxfenmode file contains 3 or more coordination points with single_cp=1 ■ etc/vxfenmode file contains 1 disk as a coordination point with single_cp=1
2416842	The "had" process can accept a limited number of connections from clients. This limit (FD_SETSIZE) is determined by the operating system. However, the accept system call can return a file descriptor greater than the limit. In such a case "had" cannot process this file descriptor using the select system call. As a result "had" goes into a unrecoverable loop.
2417843	The action entry point master.vfd queries the DNS servers without specifying the Domain name in the `dig` command. Therefore, it failes to query the SOA record for the configured domain.
2491635	VxVM introduced autostartvolumes feature in 5.1SP1 release. The DiskGroup agent online entry point in VCS 5.1SP1 fails to correctly verify the VxVM version to check the availability of autostartvolumes feature.
2569036	The MonitorTimeStats attribute may intermittently reflect incorrect values of the average monitor time for a resource.
2271882	The default value of Netlsnr attribute was set for traditional monitoring but was not set for AMF.
2318334	Oracle agents use /usr/lib in LD_PRELOAD before Oracle libraries.
2417840	The offline entry points must delete all configured resource records for multi-home ResRecord attribute values if OffDelRR attribute is set to 1. This functionality failed to work properly and removed only some of the resource records during offline operation.
2417841	The clean entry points must delete the resource records configured in ResRecord attribute of DNS type resource if OffDelRR attribute is set to 1. This functionality failed to work. It is required at the time of execution of clean entry point when some configured resource records are present at the DNS server and are required to be deleted.
2424812	When the Oracle owner uses CSH shell, Oracle agent fails after an upgrade to version 5.0MP4RP1(for Linux) or 5.0MP3RP5(for other unices).

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

Incident	Description
2438621	The MultiNICB agent compares the subnets of all the interfaces on a system, and in the above case, reports an incorrect resource state.
2511385	The Sybase agent declares a resource online before the Sybase database can complete its recovery. The following message also appears in the engine logs: `recovery state is UNKNOWN- online script terminating`
2296172	Failover Policy does not take into consideration AutoFailover = 2 and SystemZones when the nodes within a SystemZone are brought down/rebooted.
2477280	Application resource does not failover when system reboots after Concurrency Violation
2477296	When a service group is in the process of failover, if a flush operation is performed on the target node when the service group is not active on the target node, then the value of TargetCount attribute is inadvertently set to 0. Hence the service group does not failover when the node panics in future.
2483044	When a group is in the middle of a transition (i.e. with failover set and resources in path going offline) and one resource in path faults and other resource in path offlines then is engine received one more probe for offlined resource it result in assertion and HAD core dumps.
2439695	VXFEN module gets loaded even though the user chooses not to enable VXFEN

Veritas Storage Foundation Cluster File System fixed issues

Table 1-11 lists the Veritas Storage Foundation Cluster File System (SFCFS) issues fixed in this release.

Table 1-11 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
2425429	On Cluster File System (CFS), if any node is rebooted and it rejoins the cluster after quota is turned ON, it fails to mount the file system.
2420060	In a Cluster File System (CFS) setup, a hang occurs in the cluster when one of the nodes in the cluster leaves or is rebooted.

Table 1-11 Veritas Storage Foundation Cluster File System fixed issues *(continued)*

Incident	Description
2413004	In a Cluster File System (CFS) environment with partitioned directory enabled (disk layout 8), the system hangs when there is an attempt to create a large number of subdirectories within a directory.
2360819	When creating a new file, Cluster File System (CFS) unexpectedly and prematurely displays a 'file system out of inodes' error.
2340834	When multiple `vxassist mirror` commands are running on different nodes of a cluster, the nodes may panic.
2340831	When the Veritas Cluster Server (VCS) engine, High Availability Daemon (HAD) does not respond, the Group Atomic Broadcast (GAB) facility causes the system to panic.
2329887, 2412181, 2418819	In a Cluster File System (CFS) environment , the file read performances gradually degrade up to 10% of the original read performance and `fsadm -F vxfs -D -E` *`mount_point`* shows a large number (> 70%) of free blocks in extents smaller than 64k.
2247299	In a Cluster File System (CFS) setup, one of the nodes may hang repeatedly during the execution of the `vx_event_wait()` function.
2243061	Performing a nested mount on a CFS file system triggers a data page fault if a forced unmount is also taking place on the CFS file system.

Table 1-12 lists the Veritas Group Lock Manager issues fixed in this release.

Table 1-12 Veritas Group Lock Manager fixed issues

Incident	Description
2406572	The System Activity Reporter (SAR) utility on HP-UX shows some processes for Group Lock Manager (GLM) in Primary Rate Interface (PRI) state.
2241125	During an internal tetsing, some specific tests related to the `glmdump` command fail on a 4-node cluster.

Table 1-13 lists the Veritas Cluster Server Agents for Cluster File System issues fixed in this release.

Table 1-13 Veritas Cluster Server Agents for Cluster File System fixed issues

Incident	Description
2422830	Errors in `main.cf` are observed, after installing SFCFS 5.1 SP1RP1 by using OS command `swinstall`, and starting the cluster after a system reboot.
2241317	During an internal testing, some Veritas Cluster Server Agents for Cluster File System related tests fail on a four-node cluster.
2235677	CFS mount point deletion by using `cfsmntadm delete` *`mount_point`* fails due to improper mount point search.

Veritas Storage Foundation for Oracle RAC fixed issues

There are no fixes specific to this release.

For issues fixed in version 5.1 SP1, see the *Veritas Storage Foundation for Oracle RAC Release Notes (5.1 SP1)*.

Software limitations in this release

This section describes the software limitations in this release.

- Veritas Storage Foundation:
  See "Veritas Storage Foundation software limitations" on page 42.
- Veritas Storage Foundation for Databases (SFDB) Tools:
  See "Veritas Storage Foundation for Databases tools software limitations" on page 44.
- Veritas Cluster Server:
  See "Veritas Cluster Server software limitations" on page 45.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System software limitations" on page 45.
- Veritas Storage Foundation for Oracle RAC:
  See "Veritas Storage Foundation for Oracle RAC software limitations" on page 46.

Veritas Storage Foundation software limitations

This section describes the software limitations in this release of Veritas Storage Foundation.

Veritas Dynamic Multi-Pathing software limitations

The following are software limitations in this release of Veritas Dynamic Multi-Pathing (DMP).

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, set the following Dynamic Multi-Pathing (DMP) tunables:

Table 1-14

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

The change is persistent across reboots.

To change the tunable parameters

1 Issue the following commands:

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

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

```
# vxdmpadm gettune dmp_restore_interval
# vxdmpadm gettune dmp_path_age
```

LVM volume group is in unusable state if last path is excluded from DMP (1976620)

When a DMP device is used by a native LVM volume group, do not exclude the last path to the device. This can put the LVM volume group in an unusable state.

Veritas Volume Replicator software limitations

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

Replication in a shared environment

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

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

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

VVR support for replicating across Storage Foundation versions

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

Veritas Storage Foundation for Databases tools software limitations

The following are software limitations of Storage Foundation for Databases (SFDB) tools in this release.

Oracle Data Guard in an Oracle RAC environment

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

Upgrading if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 5.1SP1 RP1, upgrade the Oracle binaries and database to version 11.1.0.7, and move to SP1 and then upgrade to SP1 RP1.

Veritas Cluster Server software limitations

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

Limitations related to I/O fencing

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

Stopping systems in clusters with I/O fencing configured

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

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

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

Workaround:

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

Veritas Storage Foundation Cluster File System software limitations

The following are software limitations in this release of Veritas Storage Foundation Cluster File System.

cfsmntadm command does not verify the mount options (2078634)

You must confirm if the mount options are correct which are then passed to the`cfsmntadm` command. If the mount options are incorrect, the mount fails and the CFSMount resource will not come online. You can check the VCS engine log file for any mount failure messages.

Veritas Storage Foundation for Oracle RAC software limitations

There are no software limitations in this release.

For software limitations in version 5.1 SP1, see the *Veritas Storage Foundation for Oracle RAC Release Notes (5.1 SP1)*.

Known issues in this release

This section describes the known issues in this release.

- Issues related to installation:
  See "Issues related to installation" on page 46.
- Veritas Storage Foundation:
  See "Veritas Storage Foundation 5.1 SP1 RP1 known issues" on page 50.
- Veritas Storage Foundation for Databases (SFDB) Tools:
  See "Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP1 known issues" on page 63.
- Veritas Cluster Server:
  See "Veritas Cluster Server 5.1 SP1 RP1 known issues" on page 66.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System 5.1 SP1 RP1 known issues" on page 72.
- Veritas Storage Foundation for Oracle RAC:
  See "Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP1 known issues" on page 74.

Issues related to installation

This section describes the known issues during installation and upgrade.

The Web-based installer does not work from the disc (2321818)

The Web-based installer fails to run.

Workaround:

For this workaround, you need to have about 1.7 GB of local storage available. Copy the disc to a local system and start the Web-based installer from the local copy. Symantec recommends that you use cpio for these operations.

To start the Web-based installer workaround

1 Create a mount point.

```
# mkdir /mnt/dvd
```

2 Optionally to find the specific device path (/dev/dsk/cxtxdx), run this command:

```
# /usr/sbin/ioscan -fnkC disk
```

3 Mount the disc to the mount point.

```
# mount /dev/dsk/cxtxdx /mnt/dvd
```

4 Create a temporary installation directory.

```
# mkdir /tmp/HXRT51SP1
```

5 Create a symbolic link from the disc to the temporary installation directory.

```
# ln -s /mnt/dvd/* /tmp/HXRT51SP1/
```

6 Remove the installer link from the temporary installation directory.

```
# rm -rf /tmp/HXRT51SP1/scripts
```

7 Copy the installer scripts from the disc to the temporary installation directory.

```
# cp -rf /mnt/dvd/scripts/ /tmp/HXRT51SP1/
```

8 Start the Web-based installer from the temporary installation directory.

```
# /tmp/HXRT51SP1/webinstaller start
```

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 (\").

Incorrect error messages: error: failed to stat, etc. (2120567)

During installation, you may receive errors such as, "error: failed to stat /net: No such file or directory."

Ignore this message. You are most likely to see this message on a node that has a mount record of /net/x.x.x.x. The /net directory, however, is unavailable at the time of installation.

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*

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

NetBackup 6.5 or older version is installed on a VxFS file system. Before upgrading to Veritas Storage Foundation (SF) 5.1 SP1, the user umounts all VxFS file systems including the one which hosts NetBackup binaries (/usr/openv). While upgrading SF 5.1 SP1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure packages VRTSpbx, VRTSat, and VRTSicsco, which causes NetBackup to stop working.

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

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

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

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

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

The VRTSacclib depot is deprecated (2032052)

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

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

The `-help` option for certain commands prints an erroneous argument list (2138046)

For installsf, installat, and the installdmp scripts , although the `-help` option prints the `-security`, `-fencing`, `-addnode` options as supported, they are in fact not supported. These options are only applicable for high availability products.

Web installation looks hung when -tmppath option is used (2160878)

If you select the -tmppath option on the first page of the webinstaller after installing or uninstalling is finished on the last page of webinstaller, when you click the **Finish** button, the webpage hangs. Despite the hang, the installation or the uninstallation finishes properly and you can safely close the page.

Veritas Storage Foundation 5.1 SP1 RP1 known issues

The Veritas Storage Foundation known issues in the 5.1 SP1 release are listed in *Veritas Storage Foundation Release Notes (Version 5.1 SP1)*.

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

Veritas Volume Manager known issues

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

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: Specify explicitly the length of privoffset, puboffset, publen, and privlen while initializing the disk.

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.

vxdg split or join operations can fail for disks with a disk media name greater than or equal to 27 characters (2063387)

If a disk's media name is greater than or equal to 27 characters, certain operations, such as diskgroup split or join, can fail with the following error:

```
VxVM vxdg ERROR : vxdg move/join dg1 dg2 failed subdisk_name : Record
already exists in disk group
```

VxVM uses disk media names to create subdisk names. If multiple subdisks are under the same disk, then the serial number, starting from 1, is generated and appended to the subdisk name so as to identify the given subdisk under the physical disk. The maximum length of the subdisk name is 31 characters. If the disk media name is long, then the name is truncated to make room for serial numbers. Therefore, two diskgroups can end up having same subdisk names due to this truncation logic, despite having unique disk media names across diskgroups. In such scenarios, the diskgroup split or join operation fails.

Workaround: To avoid such problems, Symantec recommends that disk media name length should be less than 27 characters.

Shared disk group creation on slave fails if the naming scheme on slave is operating system native scheme with the mode as the new name (2148981)

While creating shared disk groups on slaves using the command shipping feature, the disk group creation may fail if the naming scheme on the slave where the command was issued is the operating system's native scheme with the mode as the new name.

Workaround: You can create the shared disk group from the slave by changing the naming scheme to the operating system's native scheme while in the "Legacy" mode.

After initializing a disk for native LVM, the first instance of vxdisk list fails with a 'get_contents' error and errant flags are displayed (2074640)

After you initialize a disk that is under the operating system's native LVM control and not under Veritas Volume Manager (VxVM) control by using the `pvcreate path_to_physical_disk` command, the first time that you run the vxdisk list disk_name command, the command displays the following error:

```
VxVM vxdisk ERROR V-5-1-539 Device disk_name: get_contents failed:
Disk device is offline
```

In addition, the `flags` field is incorrectly populated. However, in the next instantiation of the same command, VxVM does not produce an error and the flags are correctly populated with the LVM tag.

Workaround: Issue the vxdisk list disk_name command a second time.

vxconfigd fails to allocate memory until the daemon is restarted (2112448)

Veritas Volume Manager (VxVM) utilities may fail with the following error message:

```
Memory allocation failure
```

This error implies that there is insufficient memory for the vxconfigd daemon. A program's data segment size is enforced by the operating system tunablemaxdsiz. The default value of maxdsiz is 1 GB. With this default maxdsiz value, the vxconfigd daemon can allocate a maximum of 1 GB of memory.

Workaround: You might need to increase the operating system maxdsiz tunable's value appropriately to increase the data storage segment for the programs.

See the maxdsiz(5) manual page for more information.

After increasing the value, you must stop and restart the vxconfigd daemon. Depending on the maxdsiz tunable value, vxconfigd can allocate a maximum up to 2 GB of memory on PA machines, and 4 GB of memory on IA machines.

vxdisksetup fails to give a LUN the cdsdisk format if the LUN is larger than 1 TB and the system is using Tachyon HBAs (2146340)

The vxdisksetup command fails to initialize a LUN to have the cdsdisk format if the LUN is larger than 1 TB and the system is using Tachyon HBAs. The vxdisksetup command displays the following error:

```
VxVM vxdisk ERROR V-5-1-5433 Device disk_name: init failed:
Disk is not useable, bad format
```

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.

Workaround:

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

Shared disk group creation on slave fails if the naming scheme on slave is operating system native scheme with the mode as the new name (2148981)

While creating shared disk groups on slaves using the command shipping feature, the disk group creation may fail if the naming scheme on the slave where the command was issued is the operating system's native scheme with the mode as the new name.

Workaround:

You can create the shared disk group from the slave by changing the naming scheme to the operating system's native scheme while in the "Legacy" mode.

Messages observed for VxVM rooted systems during boot up (2423036)

- For VxVM rooted systems, you may observe the following message during boot up:

  ```
  WARNING: VxVM vxdmp V-5-3-0 APM for array type
   Array_type is not available
  ```

  For VXVM rooted disks, the BOOT volume `/stand` is under VxVM control. During very early boot time, only `/` is mounted and available. Since the BOOT volume `/stand` is not available at this time and Dynamic Multi-Pathing (DMP) Array Policy Module (APM) modules are `AUTO DLKM` modules, APMs are not loaded. DMP makes use of procedures of generic array type during this stage if it does not find any loaded APMs. Further during the boot process, the BOOT volume `/stand` is mounted and available. APMs are loaded successfully at this time.
  This message is harmless.

- For VxVM rooted systems, you may observe the following message during boot up when using a boot disk under DMP control:

  ```
  NOTICE: VxVM vxio V-5-0-0Could not find VxVM entry in Kernel Registry
  Service,so root disk is being claimed by DMP
  ```

  VxVM uses Kernel Registry Service to maintain persistently whether the boot disk belongs to DMP or native MultiPathing (nMP). VxVM reads kernel registry and takes this decision. In case VxVM does not find an entry in the kernel registry, this message is printed.
  This message is informative and harmless.

Veritas Dynamic Multi-Pathing known issues

This section describes the Veritas Dynamic Multi-Pathing (DMP) known issues for this release.

Issues with ALUA arrays that support standby Asymmetric Access State (AAS) when using EFI disks on HP 11i v3 (2057649)

This issue was seen with ALUA arrays that support standby Asymmetric Access State (AAS) when Extensible Firmware Interface (EFI) disks are present on the system. The HP-UX native multipath plugin (NMP) driver does not recognize the hardware path that DMP has selected and selects the standby path for internal I/Os.

This issue causes delays with Veritas Volume Manager (VxVM) device discovery and other VxVM commands. Veritas Volume Manager does not support SAN booting with these arrays on HP-UX 11i v3.

DMP not supported with LVM 2.2 (2071691)

In this release, Veritas Dynamic Multi-Pathing (DMP) is supported with Logical Volume Manager (LVM) versions 1.0, 2.0, and 2.1. DMP devices cannot be used with LVM 2.2. The vgchange command hangs or causes a panic.

DMP path discovery behavior when a device is removed from PowerPath control (2144891)

To remove a device from PowerPath control, you use the powermt unmanagecommand. When you remove a device from PowerPath control, DMP requires two device discovery cycles to discover the attributes of the paths of the device correctly.

Workaround:

Issue the following command to start the device discovery:

```
# vxdisk scandisks
```

After the discovery completes, issue the command again to start a second device discovery cycle.

Path name character limit when converting LVM volumes over DMP to VxVM volumes over DMP (2035399)

The HP-UX lvdisplay utility truncates physical volume path names to 22 characters. If a path name is truncated, utilities such as vxvmconvert orvxautoconvert that depend on the lvdisplay output may not function properly. If you intend to use the vxvmconvert utility or the vxautoconvert utility to convert LVM over DMP to VxVM over DMP, Symantec recommends that you reduce the length of the enclosure name to at most 8 characters before enabling native stack support.

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 theNODE_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.

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

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

Table 1-15shows the Hitachi arrays that have new array names.

Table 1-15 Hitachi arrays with new names

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

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

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

Enclosure name limitation when using HP-UX LVM pvcreate command on DMP device

For HP-UX LVM on a DMP device, you cannot use the pvcreate command if the enclosure-based name of the DMP device contains the 's' character. This is a limitation of the pvcreate utility on HP-UX LVM.

Workaround: Rename the enclosure to replace the 's' with some other character in the name of the enclosure before you run the pvcreate command. To rename the enclosure, use the following command:

```
# vxdmpadm setattr enclosure enclr_name name=new_enclr_name
```

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.

Veritas File System known issues

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

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.

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 the `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: Use the `vxtunefs` command and set`write_pref_io` and `write_nstream` to high values, such that `write_pref_io`multiplied by `write_nstream` is around 8 MB.

Storage Checkpoints can exceed the quota limit (2102201)

Under some circumstances, Storage Checkpoints can exceed the quota limit set by 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.

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.

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.

Storage Foundation 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 Storage Foundation 5.0 MP3 and Secondary sites running Storage Foundation 5.1 SP1, or vice versa, you must install the Storage Foundation 5.0MP3 Rolling Patch 2 on the nodes using 5.0MP3. This patch resolves several outstanding issues for replicating between versions.

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

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

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:

```
# /sbin/init.d/vras-vradmind.sh stop
# /sbin/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.

There is no workaround for this issue.

vxassist relayout removes the DCM (2162522)

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

To replace the DCM, enter the following:

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

1 Pause or stop the applications.

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

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

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.

vradmin functionality may not work after a master switch operation (2163712)

In certain situations, if you switch the master role, vradmin functionality may not work. The following message displays:

```
VxVM VVR vxrlink ERROR V-5-1-15861 Command is not supported for command
shipping. Operation must be executed on master
```

Workaround:

To restore vradmin functionality after a master switch operation

1 Restart vradmind on all cluster nodes.

2 Re-enter the command that failed.

Veritas Enterprise Administrator known issues

The following are the Veritas Enterprise Administrator (VEA) known issues for this release.

The system properties wizard in the Veritas Enterprise Administrator GUI displays the VxFS and OSFS provider versions as 6.0.000.0 instead of 5.1.100.0 (2325730)

The system properties wizard in the Veritas Enterprise Administrator (VEA) GUI displays the VxFS and OSFS provider versions as 6.0.000.0 instead of 5.1.100.0.

This is a cosmetic issue that has no impact on functionality.

Dynamic Multi-Pathing objects are not visible with Veritas Enterprise Administrator (2535298)

After you install the PHCO_42182 patch, the Dynamic Multi-Pathing (DMP) objects are not visible with Veritas Enterprise Administrator (VEA).

Resolution: Install the PHCO_42319 patch to fix this issue.

Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP1 known issues

This section lists the SFDB tools known issues in this release.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0MP2 to 5.1SP1 (2003131)

While upgrading from 5.0MP2 to 5.1SP1 the following error message could be seen when running `sfua_rept_migrate`:

```
# /opt/VRTSdbed/migrate/sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
SFORA sfua_rept_migrate ERROR V-81-8903 Could not start repository
database.
/usr/lib/dld.sl: Can't find path for shared library: libcur_colr.1
/usr/lib/dld.sl: No such file or directory
sh: 3845 Abort(coredump)
Symantec DBMS 3.0.85.0 vxdbms_start_db utility
ASA failed. Sybase ASA error code: [134].
Sybase ASA Error text: {{{}}}
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround:

To upgrade without an existing SFDB repository set up

1. Verify X/Open curses is installed on the system.
2. Create the following link:

   ```
   ln -s /usr/lib/libxcurses.1
   /usr/lib/libcur_colr.1
   ```

3. Run the following command:

   ```
   # sfua_rept_migrate
   ```

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0.x to 5.1SP1 (2184482)

When upgrading from Storage Foundation version 5.0 or 5.0.1 to Storage Foundation 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 whensfua_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.

Relinking ODM after upgrading from 5.0.x

The VRTSodm library path has changed from /opt/VRTSodm/lib/libodm.sl to/opt/VRTSodm/lib/libodm.so.

After upgrading to from 5.0.x you must update the ODM link for your database to the new VRTSodm library path /opt/VRTSodm/lib/libodm.so.

Upgrading in an HP Serviceguard environment (2116455)

When upgrading SFDB to 5.1SP1 from the previous release in an HP Serviceguard environment, first verify that the cmviewcl command can be executed by a non-root user. This permission change must be done before executing SFDB upgrade commands.

Using SFDB tools after upgrading Oracle to 11.2.0.2 (2203228)

The procedure which Oracle recommends for upgrading to Oracle 11.2.0.2 results in the database home changing. After you upgrade to Oracle 11.2.0.2, you must run the dbed_update command with the new Oracle home provided as an argument to the -H option before using any SFDB tools. After this step, the SFDB tools can be used normally.

Database fails over during Flashsnap operations (1469310)

In a Storage Foundation 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.

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`.

There is no workaround for this issue.

VCS agent for Oracle: Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 (1985055)

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2.

Workaround: Set MonitorOption attribute for Oracle resource to 0.

Veritas Cluster Server 5.1 SP1 RP1 known issues

This section describes the known issues for VCS in this release.

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.

Verification for VRTSat package or patch returns errors (1244204)

If you run the `swverify` command on VRTSat package or patch, the command returns errors for missing files on VRTSat.CLIENT-PA32.

Workaround:

This message may be safely ignored.

Issues related to LLT

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

LLT port stats sometimes shows recvcnt larger than recvbytes (1788315)

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.

Workaround:

This does not impact the LLT functionality.

LLT can incorrectly declare port-level connection for nodes in large cluster configurations (1809827)

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.

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 (1965954)

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.

Preferred fencing does not work as expected for large clusters in certain cases (2161816)

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

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

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

See the *Veritas Storage Foundation Cluster File System Administrator's Guide* for more information on preferred fencing.

Server-based I/O fencing fails to start after configuration on nodes with different locale settings (2112742)

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

Workaround:

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

Reconfiguring Storage Foundation Cluster File System HA with I/O fencing fails if you use the same CP servers (2076240)

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.

Workaround:

Manually remove the application cluster information from the CP servers after you reconfigure Storage Foundation Cluster File System HA 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.

Workaround:

No known resolution for this issue.

Coordination Point agent does not provide detailed log message for inaccessible CP servers (1907648).

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

Issues related to IMF

This section describes the known issues realted to IMF.

If you forcefully unconfigure AMF when a script based agent is registered with AMF, the getnotification thread continuously polls and floods the engine log. (2521893).

Workaround:

Restarting the agent will resolve the issue.

Forcefully unconfiguring AMF does not change the monitor method of agent to TRADITIONAL (2521881).

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

Workaround:

Restarting the agent will resolve the issue.

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

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

Workaround:

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

Issues related to Bundled agents

This section describes the known issues related to Bundled agents in this release.

The vip service group added by `cfsshare addvip` command comes in FAULTED state (2556356).

When PingOptimize is set to 1 and no NetworkHosts is specified, NIC agent depends on packet count to report the health of the interface. If there is not enough traffic on the interface, NIC agent can report incorrect state of the interface.

Workaround:

Any of the following workarounds should resolve the issue:

- Setting PingOptimize = 0. This will make NIC agent ping the broadcast address whenever there is no traffic on the interface.
- Setting valid NetworkHosts value. This will make NIC agent to ping NetworkHosts to check health of status.

An error message is displayed when the Options attribute is not specified for IPMultiNICB agent (2557189).

When the Options attribute for IPMultiNICB is not specified, the following error message is logged by the online entry point of IPMultiNICB agent:

```
V-16-10021-14446 IPMultiNICB:ipmnicb:online:Error in configuring IP
address
```

Workaround:

The functionality is not affected by this error message.

Application Agent does not handle a case when user is root, envfile is set, and shell is csh (2513774).

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 as per the csh syntax.

Workaround:

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

The preonline_ipc trigger functionality of VCS, that performs certain checks before bringing a group online, does not work for resources other than IP resources (2528475).

This is a known limitation. There is an enhancement requirement to extend preonline_ipc trigger support to other resources types.

Issues related to VCS Engine

This section describes the known issues related to VCS Engine in this release.

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

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

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

Workaround:

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

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

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

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

When resources r2 and r3, online of resource r1 is initiated. Before OnlineTimeout, resources r2 and r3 fault. The sequence of fault detection is important i.e. first r2 and then r3. When fault of both resources is detected the group is in a OFFLINE|FAULTED state and resource r1 is waiting to go online. If `hagrp -clear` command is executed to clear the fault then HAD dumps core on all nodes due to assertion.

Workaround:

Before clearing the fault, user should flush the pending online operation using`hagrp -flush`.

In a VCS cluster that is deployed in a secure environment, VCS fails to authenticate users with an authentication broker that resides outside the VCS cluster (2272352).

For example, in LDAP-based authentication, if you install the LDAP client on a system that is not a VCS node, then you cannot use that system as an authentication broker to authenticate users on VCS nodes.

Workaround:

Symantec has introduced the VCS_REMOTE_BROKER environment variable, which you can use to authenticate users on VCS nodes, with a remote broker. VCS_REMOTE_BROKER works only with non-root users, as the root user does not require authentication to run ha commands in a local cluster

In a GCO (Global Cluster Option) setup, you may be unable to bring an IPMultiNIC resource online (2358600).

In a GCO setup, the IPMultiNIC resource may be unable to successfully use certain commands to detect the state of the corresponding MultiNICA resource. As a result, the IPMultiNIC resource does not come online.

Workaround:

Symantec has modified the IPMultiNIC agent code to fix this issue.

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

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

Workaround:

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

Issues related to installation

This section describes issues related to installation.

installrp fails to install 5.1SP1 RP1 when the root user shell is set to csh (2523643)

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

Workaround:

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

Veritas Storage Foundation Cluster File System 5.1 SP1 RP1 known issues

This section describes the Veritas Storage Foundation Cluster File System (SFCFS) known issues in this release.

Installer assigns duplicate node ID during -addnode procedure

While performing an -addnode operation using a CPI installer to a cluster where a node has failed, VCS appends the new node with a duplicate node ID of its last node. This happens only to the cluster in which any but the last node has failed. In this case,/etc/llthost displays two nodes with same node IDs. This is because VCS assigns the node ID by simply counting the number of node entries without checking the assigned node IDs.

Workaround: Instead of running the CPI command, add the new node manually as described in the *Veritas Cluster Server Installation Guide (Version 5.1 SP1)*.

SFCFSHA upgrade shows partial upgrade warning

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

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

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
```

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.

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

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 5.1 SP1 RP1 known issues

There are no known issues in this release.

For known issues in version 5.1 SP1, see the *Veritas Storage Foundation for Oracle RAC Release Notes (5.1 SP1).*

Documentation errata

The following section provides documentation updates.

Veritas Cluster Server Administrator's Guide (2444653)

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

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

Downloading the patches

The patches included in Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1 are available for download from the Symantec website. After downloading the file, use gunzip and tar to uncompress and extract.

For the 5.1 SP1 RP1 download archive and instructions, visit:

http://sort.symantec.com/patch/matrix

Chapter 2

Upgrading to version 5.1 SP1 RP1

This chapter includes the following topics:

About the installrp script

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1 provides an installation script. To install the patches that are included in this release, the recommended method is to use the `installrp` script. The `installrp` script lets you install all the patches that are associated with the packages installed. After using the `installrp` script, you may need to restart systems.

Table 2-1 lists the command line options for the installrp script.

Table 2-1 Command line options for the installrp script

Command Line Option	Function
`[ <system1> <system2>... ]`	Specifies the systems on which to run the installation options. If not specified, the command prompts for a system name.

Table 2-1 Command line options for the installrp script *(continued)*

Command Line Option	Function
`[ -precheck ]`	The `-precheck` option is used to confirm that systems meet the products install requirements before installing.
`[ -postcheck ]`	The `-postcheck` option is used to check for any issues after installation or upgrading.
`[ -logpath <`*`log_path`*`> ]`	The `-logpath` option is used to select a directory other than `/opt/VRTS/install/logs` as the location where `installrp` log files, summary file, and response file are saved.
`[ -responsefile <`*`response_file`*`> ]`	The `-responsefile` option is used to perform automated installations or uninstallations using information stored in a file rather than prompting for information. `<`*`response_file`*`>` is the full path of the file that contains configuration definitions.
`[ -tmppath <`*`tmp_path`*`> ]`	The `-tmppath` option is used to select a directory other than `/var/tmp` as the working directory for `installrp`. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
`[ -hostfile <`*`hostfile_path`*`> ]`	The `-hostfile` option specifies the location of a file containing the system names for installer.
`[ -keyfile <`*`ssh_key_file`*`> ]`	The `-keyfile` option specifies a key file for SSH. When this option is used, `-i <`*`ssh_key_file`*`>` is passed to every SSH invocation.
`[ -patchpath <`*`patch_path`*`> ]`	The `-patchpath` option is used to define the complete path of a directory available to all install systems (usually NFS mounted) that contains all patches to be installed by `installrp`.

Table 2-1 Command line options for the installrp script *(continued)*

Command Line Option	Function
`[ -rsh ]`	The `-rsh` option is used when `rsh` and `rcp` are to be used for communication between systems instead of `ssh` and `scp`. When the `-rsh` option is not used, systems must be pre-configured such that `ssh` commands between systems execute without prompting for passwords or confirmations.
`[ -redirect ]`	The `-redirect` option is used to display progress details without showing advanced display functionality so output can be redirected to a file.
`[ -listpatches ]`	The `-listpatches` option is used to display product patches in the correct installation order.
`[ -makeresponsefile ]`	The `-makeresponsefile` option generates a response file without doing an actual installation. The text displaying install, uninstall, start, and stop actions are a part of a simulation. These actions are not actually performed on the system.
`[ -pkginfo ]`	The `-pkginfo` option is used to display the correct install order of packages and patches. This option is available with or without one of following options: `-allpkgs`, `-minpkgs`, and `-recpkgs`.
`[ -serial ]`	The `-serial` option is used to perform install, uninstall, start, and stop operations, typically performed simultaneously on all systems, in a serial fashion.
`[ -upgrade_kernelpkgs ]`	The `-upgrade_kernelpkgs` option is used to perform rolling upgrade Phase-I. In this phase, the product kernel packages are upgraded to the latest version.

Table 2-1 Command line options for the installrp script *(continued)*

Command Line Option	Function
`[ -upgrade_nonkernelpkgs ]`	The `-upgrade_nonkernelpkgs` option is used to perform rolling upgrade Phase-II. In this phase, VCS and other agent packages are upgraded to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
`[ -version ]`	The `-version` option is used to check the status of installed products on the system.

Special upgrade instructions

The following special upgrade instructions apply to the respective patches in this release.

VRTS 5.1 SP1RP1 VRTSdbed Command Patch (PHCO_42093)

- By default, configurable ciphers are not enabled with the `vxdbd` daemon. To use configurable ciphers with the vxdbd daemon, ensure the following after upgrading to 5.1 SP1RP1:
  - Set the SSLCipherSuite parameter to the appropriate cipher string in the `/opt/VRTSdbed/eat/root/.VRTSat/profile/VRTSatlocal.conf` file.

    Note: By default, LOW strength ciphers are not supported if you are using configurable ciphers. The default SSLCipherSuite string is `SSLCipherSuite"="HIGH:MEDIUM:!eNULL:!aNULL:!SSLv2"`.

  - Restart the vxdbd daemon after setting the VXDBD_USE_ENCRYPT environment variable to 1.
  - All the client side scripts/binaries run with the VXDBD_USE_ENCRYPT environment variable set to 1.

VRTS 5.1 SP1RP1 VRTSperl Command Patch (PHCO_42213)

- This patch applies only to 11i v3 IPF/Integrity systems. It is not applicable to PA-RISC-based systems.

Performing a full upgrade to 5.1 SP1 RP1 on a cluster

Performing a full upgrade on a cluster requires stopping cluster failover functionality during the entire procedure. However, if you use Veritas Storage Foundation Cluster File System (SFCFS) and Cluster Volume Manager (CVM), the SFCFS and CVM services remain available.

The following are the stages of performing a full upgrade on a cluster:

- Freeze service group operations and stop Veritas Cluster Server (VCS) on the cluster.
- Take the nodes offline and install the software patches.
- Bring the nodes online to restart cluster failover services.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.1 SP1 RP1:

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

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

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

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

To upgrade VCS

1. Log in as superuser.
2. Upgrade the Operating System and reboot the systems if required.
3. Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

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

4 Resolve any issues that the precheck finds.

5 Start the upgrade:

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

6 Restart the nodes:

```
# shutdown -r now
```

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

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

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

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

1 Log in as superuser.

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

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

```
# mount | grep vxfs
```

4 Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

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

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

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

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

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

Verify that no volumes remain open:

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

8 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
```

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

```
# ./installrp -precheck [-rsh] node1node2 ... nodeN
```

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

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

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

12 Start the upgrade.

```
# ./installrp [-rsh] node1node2 ... nodeN
```

Review the output.

13 Restart the nodes:

```
# shutdown -r now
```

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

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

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

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

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

```
# mount /filesystem
```

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

```
# mount /checkpoint_name
```

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

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

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

1 Log in as superuser.

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

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

```
# mount | grep vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# mount | grep vxfs
```

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

  ```
  # umount /filesystem
  ```

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

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

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

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

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

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

6 Stop activity to all VxVM volumes.

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

7 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
```

8 If required, apply the OS kernel patches.

9 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
```

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

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

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

11 Restart the nodes:

```
# shutdown -r now
```

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

13 Bring the CVM service group online on each node:

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

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

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

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

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

16 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

17 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

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

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

To upgrade to 5.1 SP1 RP1 on a SF Oracle RAC cluster

1 Log in as superuser.

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

3 Stop VCS by running the following command:

```
# hastop -all
```

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

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

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

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

Unmount the VxFS or CFS file system:

```
# umount /mount_point
```

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

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

Verify that no volumes remain open:

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

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

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

8 Restart the nodes:

```
# shutdown -r now
```

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

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

Performing a full upgrade to 5.1 SP1 RP1 on a standalone system

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

To upgrade to 5.1 SP1 RP1 on a standalone system

1 Log in as superuser.

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

3 If required, apply the OS kernel patches.

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

```
# mount | grep vxfs
```

5 Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

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

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

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

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

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

Verify that no volumes remain open:

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

9 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

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

```
# ./installrp system
```

11 Restart the system.

```
# shutdown -r now
```

Performing a rolling upgrade to 5.1 SP1 RP1on a cluster

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

- About rolling upgrades
- Prerequisites for rolling upgrades
- Performing a rolling upgrade using the script-based installer
- Performing a rolling upgrade of the product using the Web-based 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 are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

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

You can perform a rolling upgrade for the following products:

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

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

Prerequisites for rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

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

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

Performing a rolling upgrade using the script-based installer

Navigate to the installer program to start the rolling upgrade.

To perform the rolling upgrade on kernel packages: phase 1

1 Log in as superuser to one of the nodes in the first sub-cluster.

2 Back up the configuration files on your system.

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 Unmount all the VxFS file systems 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
```

5 Start the installer.

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

6 The installer checks system communications, depot versions, product versions, and completes prechecks. Press **y** to continue.

7 The installer performs a pre-check on the nodes in the cluster. You can address the findings of the precheck, or select **y** to continue.

8 The installer lists the patches to upgrade on the selected node or nodes.

9 The installer prompts you to stop the applicable processes. Select **y** to continue.

Failover service groups now fail over to the other node or nodes. Applications in failover service groups now experience normal downtime during the failover.

10 The installer stops relevant processes and installs the new patches. It performs the configuration for the upgrade and re-starts processes.

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.

11 For SF/SFHA/SFCFS/SF Oracle RAC: Restart the nodes in the first sub-cluster:

```
# shutdown -r now
```

12 Perform the following steps on the nodes in the first sub-cluster:

- Manually mount the VxFS and CFS file systems that VCS does not manage.
- Start all applications that VCS does not manage. Use native application commands to start the applications.

13 Complete step 1 to step 4 on the nodes in the second sub-cluster.

14 Start the installer on the nodes in the second sub-cluster.

```
# ./installrp -upgrade_kernelpkgs node3 node4
```

15 For VCS: Repeat step 6 through step 9 and step 12.

For SF/SFHA/SFCFS/SF Oracle RAC: Repeat step 6 through step 12.

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

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

1 Start the installer:

```
# ./installrp -upgrade_nonkernelpkgs node1 node2 node3 node4
```

2 The installer checks system communications, depot versions, product versions, and completes prechecks. Press **y** to continue.

3 The installer performs a pre-check on the nodes in the cluster. You can address thefindings of the precheck, or select **y** to continue.

4 The installer lists the patches to upgrade on the selected node or nodes.

5 The installer prompts you to stop the applicable processes. Select **y** to continue.

6 The installer stops relevant processes and installs the new patches. It performs the configuration for the upgrade and re-starts processes.

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.

7 Verify the cluster's status:

```
# hastatus -sum
```

Performing a rolling upgrade of the product using the Web-based installer

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

Note: This release does not support a rolling upgrade of SF Oracle RAC using the Web-based installer.

The rolling upgrade is divided into two phases. In the first phase, the installer upgrade kernel packages. In the second phase, it upgrades non-kernel packages. The second phase is required for upgrades that have high-availability components. When you perform a rolling upgrade, you need to divide the number of systems that you plan to upgrade roughly in half. Half of the systems' available capacity is needed to take over processes during the rolling upgrade.

To start the rolling upgrade—phase 1

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

2 Start the Web-based installer.

3 In the Task pull-down menu, select **Rolling Upgrade**.

In the Product pull-down menu, select the product that you want to upgrade using a rolling upgrade.

Note that the Upgrade Kernel packages for Rolling Upgrade Phase-1 radio button is selected.

Click the **Next** button to proceed.

4 In the Systems Names field, enter the sub-cluster's system names. Separate system names with a space.

The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.

5 The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted.

6 When the upgrade completes, perform step 3 through step 6 on the second subcluster.

To upgrade the non-kernel components—phase 2

1. In the Task pull-down menu, make sure that **Rolling Upgrade** and the product are selected.

   Note that the Upgrade Non-Kernel packages for Rolling Upgrade Phase-2 radio button is selected.

   Click the **Next** button to proceed.

2. In the Systems Names field, enter the names of all the systems that you want to upgrade. Separate system names with a space.

   The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.

3. The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted.

Chapter 3

Uninstalling version 5.1 SP1 RP1

This chapter includes the following topics:

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

This section describes how to roll back either by using the `uninstallrp` script or the Web-based installer.

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

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

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

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

About the uninstallrp script

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1 provides a script that you can use to roll back to the 5.1 SP1 release. To uninstall the patches that are included in this release, the recommended method is to use the `uninstallrp` script.

Table 3-1 lists the command line options for the uninstallrp script.

Table 3-1 Command line options for the uninstallrp script

Command Line Option	Function
`[ <system1> <system2>... ]`	Specifies the systems on which to run the `uninstallrp` script. If not specified, the command prompts for a system name.
`[ -logpath <log_path> ]`	The `-logpath` option is used to select a directory other than `/opt/VRTS/install/logs` as the location where `uninstallrp` log files, summary file, and response file are saved.
`[ -responsefile <response_file> ]`	The `-responsefile` option is used to perform automated installations or uninstallations using information stored in a file rather than prompting for information. *`<response_file>`* is the full path of the file that contains configuration definitions.
`[ -tmppath <tmp_path> ]`	The `-tmppath` option is used to select a directory other than `/var/tmp` as the working directory for `uninstallrp`. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
`[ -hostfile <hostfile_path> ]`	The `-hostfile` option specifies the location of a file containing the system names for `uninstallrp`.

Table 3-1 Command line options for the uninstallrp script *(continued)*

Command Line Option	Function
`[ -keyfile <ssh_key_file> ]`	The `-keyfile` option specifies a key file for SSH. When this option is used, `-i <ssh_key_file>` is passed to every SSH invocation.
`[ -rsh ]`	The `-rsh` option is used when `rsh` and `rcp` are to be used for communication between systems instead of `ssh` and `scp`.
`[ -redirect ]`	The `-redirect` option is used to display progress details without showing advanced display functionality so that output can be redirected to a file.
`[ -makeresponsefile ]`	The `-makeresponsefile` option generates a response file without doing an actual installation. The text displaying install, uninstall, start, and stop actions are a part of a simulation. These actions are not actually performed on the system.
`[ -serial ]`	The `-serial` option is used to perform install, uninstall, start, and stop operations, typically performed simultaneously on all systems, in a serial fashion.
`[ -version ]`	The `-version` option is used to check the status of installed products on the system.

Rolling back using the uninstallrp script

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

To roll back on a standalone system

1 Browse to the directory that contains the uninstallrp script.

2 Stop all VxVM volumes. For each disk group enter:

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

Verify that no volumes remain open.

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

3 Unmount all the Storage Checkpoints and the file systems.

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

Verify that you unmounted the the Storage Checkpoints and the file systems.

```
# mount | grep vxfs
```

4 Run the uninstallrp script to rollback patches, type:

```
# ./uninstallrp
```

5 Restart the system:

```
# shutdown -r now
```

The uninstallrp script removes 5.1 SP1 RP1 patches. After patch rollback completes, modules are loaded and processes are restarted. `uninstallrp` will also report any warning happened during uninstallation.

To roll back in a cluster setup

1 Stop VCS:

```
# hastop -all
```

2 Use native application commands to stop the applications that use VxFS or VxVM disk groups on each node and that are not under VCS control, whether local or CFS.

3 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
```

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

```
# ./uninstallrp
```

To roll back on all the cluster nodes in one go, type:

```
# ./uninstallrp system1 system2 systemn
```

5 Restart the nodes:

```
# shutdown -r now
```

6 Manually mount the VxFS and CFS file systems that VCS does not manage.

7 Start all applications that VCS does not manage. Use native application commands to start the applications.

Uninstalling 5.1 SP1RP1 with the Web-based installer

This section describes how to uninstall this release with the Web-based installer.

To uninstall 5.1 SP1RP1

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

2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

3. Start the Web-based installer.

4. On the **Select a task and a product** page, select **Uninstall a Product** from the **Task** drop-down list.

5. Select **Storage Foundation or Storage Foundation High Availability** from the **Product** drop-down list, and click **Next**.

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

7. After the validation completes successfully, click **Next** to uninstall SFHA on the selected system.

8. If there are any processes running on the target system, the installer stops the processes. Click **Next**.

9. After the installer stops the processes, the installer removes the products from the specified system. Click **Next**.

10. After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.

11. Click **Finish**.

    The Web-based installer prompts you for another task.

Chapter 4

Reference documentation

This chapter includes the following topics:

- Reference documentation

Reference documentation

The documentation for this release is available on the software disc in the PDF format. Symantec recommends copying documentation from the disc to your system directory. This release includes the following document.

Table 4-1 lists the document included in this release.

Table 4-1 Documentation in 5.1 SP1 RP1

Title	File Name
Veritas Storage Foundation and High Availability Solutions Release Notes	`sfha_notes_51sp1rp1_hpux11iv3.pdf`

To refer to the product documentation for the Veritas Storage Foundation and High Availability Solutions 5.1 SP1 release, go to the following URL:

https://sort.symantec.com/documents