Veritas Storage Foundation™ and High Availability Solutions Release Notes

HP-UX 11i v3

5.1 Service Pack 1 Rolling Patch 2

Symantec

Veritas Storage Foundation™ and High Availability Solutions Release Notes

Product version: 5.1 SP1 RP2

Document version: 5.1SP1RP2.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Contents

Chapter 1

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 2 (5.1 SP1 RP2). Symantec strongly recommends installing the 5.1 SP1 Rolling Patch 2 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 RP2.

- Changes in Veritas Storage Foundation High Availability:
  See "Changes in Veritas Storage Foundation High Availability" on page 12.

Changes in Veritas Storage Foundation High Availability

This release supports HP Integrity Virtual Machines (IVM) 6.1.

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 RP2" on page 13.
- Veritas Cluster Server:
  See "Veritas Cluster Server patches in 5.1 SP1 RP2" on page 14.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System patches in 5.1 SP1 RP2" on page 14.
- Veritas Storage Foundation for Oracle RAC
  See "Veritas Storage Foundation for Oracle RAC patches in 5.1 SP1 RP2" on page 15.

Veritas Storage Foundation patches in 5.1 SP1 RP2

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

Table 1-1 Veritas Storage Foundation patches

Patch	Version	Description
PHCO_43172	1.0	VRTS 5.1 SP1RP2 VRTSvxfs Command Patch (Veritas File System)
PHKL_43127	1.0	VRTS 5.1 SP1RP2 VRTSvxfs Kernel Patch (Veritas File System)
PHKL_42741	1.0	VRTS 5.1 SP1RP1P1 VRTSodm Kernel Patch (Veritas File System)
PHCO_43065	1.0	VRTS 5.1 SP1RP2 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_43064	1.0	VRTS 5.1 SP1RP2 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHCO_43070	1.0	VRTS 5.1 SP1RP2 VRTSdbed Command Patch (Veritas Storage Foundation for Databases Tools)

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

Patch	Version	Description
PHCO_42319	1.0	VRTS 5.1SP1 VRTSob Command Patch.
PHCO_42213	5.10.0.13	VRTS 5.1 SP1RP1 VRTSperl Command Patch (Perl Redistribution)

Veritas Cluster Server patches in 5.1 SP1 RP2

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

Table 1-2 Veritas Cluster Server patches

Patch	Version	Description
PHCO_43069	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Command Patch
PHKL_43068	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Kernel Patch
PHNE_43151	1.0	VRTS 5.1 SP1RP2 VRTSllt Patch
PHNE_43152	1.0	VRTS 5.1 SP1RP2 VRTSgab Patch
PVCO_03958	1.0	VRTS 5.1 SP1RP2 VRTSvcs Command Patch
PVCO_03959	1.0	VRTS 5.1 SP1RP2 VRTSvcsag Command Patch
PVCO_03960	1.0	VRTS 5.1 SP1RP2 VRTSvcsea Command Patch
PVCO_03962	1.0	VRTS 5.1 SP1RP2 VRTSamf Command Patch

Veritas Storage Foundation Cluster File System patches in 5.1 SP1 RP2

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_43172	1.0	VRTS 5.1 SP1RP2 VRTSvxfs Command Patch (Veritas File System)
PHKL_43127	1.0	VRTS 5.1 SP1RP2 VRTSvxfs Kernel Patch (Veritas File System)

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

Patch	Version	Description
PHCO_43065	1.0	VRTS 5.1 SP1RP2 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_43064	1.0	VRTS 5.1 SP1RP2 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHCO_43070	1.0	VRTS 5.1 SP1RP2 VRTSdbed Command Patch (Veritas Storage Foundation for Databases Tools)
PVCO_03956	1.0	VRTS 5.1 SP1RP2 VRTScavf Command Patch (Veritas Cluster Server Agents for Cluster File System)
PVCO_03962	1.0	VRTS 5.1 SP1RP2 VRTSamf Command Patch (Veritas Agent Monitoring Framework)
PHCO_43069	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Command Patch (Veritas I/O Fencing)
PHKL_43068	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Kernel Patch (Veritas I/O Fencing)
PHCO_42319	1.0	VRTS 5.1SP1 VRTSob Command Patch.
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 RP2

Table 1-4 lists the Veritas Storage Foundation for Oracle RAC patches included in this release.

Table 1-4 Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 patches

Patch	Version	Description
PHCO_43172	1.0	VRTS 5.1 SP1RP2 VRTSvxfs Command Patch (Veritas File System)
PHKL_43127	1.0	VRTS 5.1 SP1RP2 VRTSvxfs Kernel Patch (Veritas File System)
PHCO_42319	1.0	VRTS 5.1SP1 VRTSob Command Patch
PHKL_42741	1.0	VRTS 5.1 SP1RP1P1 VRTSodm Kernel Patch

Table 1-4 Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 patches *(continued)*

Patch	Version	Description
PHCO_43065	1.0	VRTS 5.1 SP1RP2 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_43064	1.0	VRTS 5.1 SP1RP2 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PVCO_03957	1.0	VRTS5.1SP1RP2 VRTSdbac Command Patch (Storage Foundation for Oracle RAC)
PVCO_03958	1.0	VRTS5.1SP1RP2 VRTSvcs Command Patch
PVCO_03959	1.0	VRTS5.1SP1RP2 VRTSvcsag Command Patch
PVCO_03960	1.0	VRTS5.1SP1RP2 VRTSvcsea Command Patch
PHCO_43070	1.0	VRTS 5.1 SP1RP2 VRTSdbed Command Patch (Veritas Storage Foundation for Databases Tools)
PHKL_42342	1.0	VRTS 5.1 SP1RP2 VRTSglm Kernel Patch (Veritas Group Lock Manager)
PVCO_03962	1.0	VRTS 5.1 SP1RP2 VRTSamf Command Patch (Veritas Agent Monitoring Framework)
PVCO_03963	1.0	VRTS 5.1 SP1RP2 VRTScavf Command Patch
PVCO_03930	1.0	VRTS 5.1 SP1RP1 VRTScps Command Patch (Veritas Coordination Point Server)
PHCO_43069	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Command Patch (Veritas I/O Fencing)
PHNE_43151	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Kernel Patch (Veritas I/O Fencing)
PHNE_43152	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Kernel Patch (Veritas I/O Fencing)
PHKL_43068	1.0	VRTS 5.1 SP1RP2 VRTSvxfen Kernel Patch (Veritas I/O Fencing)
PHCO_42213	5.10.0.13	VRTS 5.1 SP1RP1 VRTSperl Command Patch (Perl Redistribution)

Fixed issues in this release

This section describes issues fixed in this release.

- Veritas Storage Foundation:
  See "Veritas Storage Foundation 5.1 SP1 RP2 fixed issues" on page 17.
- Veritas Storage Foundation for Databases (SFDB) tools:
  See "Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP2" on page 50.
- Veritas Cluster Server:
  See "Veritas Cluster Server 5.1 SP1 RP2 fixed issues" on page 51.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 fixed issues" on page 66.
- Veritas Storage Foundation for Oracle RAC:
  See "Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 fixed issues" on page 69.

Veritas Storage Foundation 5.1 SP1 RP2 fixed issues

Table 1-5 lists the Veritas Volume Manager issues fixed in this release.

Table 1-5 Veritas Volume Manager fixed issues

Incident	Description
2560843	In VVR (Veritas Volume Replicator) setup I/Os can hang in slave nodes after one of the slave node is restarted.
2556781	In cluster environment, import attempt of imported disk group may return wrong error.
2216951	`vxconfigd` dumps core because `chosen_rlist_delete()` function hits NULL pointer in linked list of clone disks.
1431223	`vradmin syncvol` and `vradmin syncrvg` commands do not work if the remote disk group and vset names are specified while synchronizing vsets.
2739709	Disk Group rebuild fails as the links between volume and vset missing from `vxprint -D -` output.
1675482	`vxdg list <dgname>` command shows configuration copy in new failed state.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2441937	`vxconfigrestore precommit` fails with awk errors.
2425259	`vxdg` join operation fails with VE_DDL_PROPERTY: Property not found in the list.
2774406	System may panic while accessing Data Change Map (DCM) volume.
2566174	Null pointer dereference in `volcvm_msg_rel_gslock()` results in panic.
2688308	When re-import of disk group fails during master takeover, other shared disk groups should not be disabled.
2715129	`Vxconfigd` hangs during Master takeover in a CVM (Clustered Volume Manager) environment.
2354046	Man page for dgcfgrestore is incorrect.
2513101	User data corrupted with disk label information.
2626741	Using `vxassist` with `-o ordered` and `mediatype:hdd` options together does not work as expected.
2348199	`vxconfig` dumps core while importing a disk group .
2627126	I/O hang seen due to I/Os' stuck at DMP level.
2760181	Panic hit on secondary slave during log owner operation.
2648176	Performance difference on Master versus Slave during recovery by Document Change Object (DCO).
2620556	I/O hung after Storage Replicator Log (SRL) overflow.
2763206	`vxdisk rm` command dumps core when disk name of very large length is given.
2838059	VVR Secondary panic in `vol_rv_update_expected_pos`.
2836798	In VxVM, resizing simple EFI disk fails and causes system to panic or hang.
1291519	After multiple VVR migrate operations, `vrstat` fails to output statistics.
2000585	`vxrecover` does not start the remaining volumes if one of the volumes is removed during `vxrecover` command run.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2104887	`vxdg` import error message needs improvement for cloned disk group import failure.
2756059	System may panic when large `cross-dg` mirrored volume is started at boot.
2492451	The start up script `vxvm-startup2` launches `vxesd` without checking the presence of `install-db` file.
2578336	Failed to online the Cross Data platform Sharing (CDS) disk.
2647975	Shared Disk Group (DG) had a split brain condition, when customer ran `hastop -local` command.
2826125	VxVM script daemon is terminated abnormally on its invocation.
2739601	`vradmin repstatus` output occasionally reports abnormal timestamp.
2792748	Node join fails because of closing of wrong file descriptor.
2515369	`vxconfigd`(1M) can hang in the presence of EMC BCV devices.
2801962	Growing a volume takes significantly large time when the volume has version 20 Document Change Object (DCO) attached to it.
2606695	Panic in CVR (Clustered Volume Replicator) environment when performs I/O Operations.
2516584	Startup scripts use `quit` instead of `exit`, causing empty directories in `/tmp`
1903700	Removing mirror using vxassist does not work.
2585239	VxVM commands run very slow on setup with tape devices.
2626199	`vxdmpadm list dmpnode` command shows incorrect path-type.
2818840	Enhance the `vxdmpasm` utility to support various permissions and `root:non- system` ownership can be set persistently.
2677016	Veritas Event Manager (vxesd(1M)) daemon dumps core when main thread tries to close one of its thread (which hold connection with HP Event Manager).
2413763	Uninitialized memory read results in a `vxconfigd` coredump.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2753954	When a cable is disconnected from one port of a dual-port FC HBA, the paths by another port are marked as SUSPECT PATH.
2637217	Document new storage allocation attribute support in vradmin man page for resizevol and resizesrl.
2617277	Man pages for the `vxautoanalysis` and `vxautoconvert` commands missing from the base package.
2526623	Memory leak detected in Cluster Volume Manager (CVM) code.
2556467	Disabling all paths and restarting of the host causes losing of `/etc/vx/.vxdmprawdev` records.
2656803	VVR (Veritas Volume Replicator) panics when `vxnetd` start and stop operations are invoked in parallel.
2227678	Second rlink goes into DETACHED STALE state in multiple secondaries environment when Storage Replicator Log (SRL) has overflowed for multiple rlinks.
2277558	`vxassist` outputs a misleading error message during snap-shot related operations.
2389554	`vxdg listssbinfo` output is incorrect.
2815517	`vxdg adddisk` allows mixing of clone and non-clone disks in a disk group.
2775960	In secondary Cluster Volume Replicator (CVR) case, I/O hang seen on a disk group during Storage Replicator Log (SRL) disable activity on other disk group.
2252680	`vxtask abort` does not appropriately clean up the tasks.
2664825	Disk Group (DG) import fails when disk contains no valid UDID tag on config copy and config copy is disabled.
2560835	I/Os and `vxconfigd` hung on master node after slave is restarted under heavy I/O load.
2423608	Panic in vol_dev_strategy() following FC problems.
2495332	`vxcdsconvert` fails if the private region of the disk to be converted is less than 1 MB.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2620555	I/O hangs due to Storage Replicator Log (SRL) overflow and Cluster Volume Manager (CVM) reconfiguration.
925653	Node join fails for higher Cluster Volume Manager (CVM)Timeout value.
2657797	Starting a RAID5 volume fails, when one of the subdisks in the RAID5 column starts at an offset greater than 1TB.
2576602	`vxdg listtag` should give error message and display correct usage when executed with wrong syntax.
2599526	I/O hang seen when DCM is zero.
2689845	Data disk can go in error state when data at the end of the first sector of the disk is same as MBR signature.
2606709	Storage Replicator Log (SRL) overflow and Cluster Volume Replicator (CVR) reconfiguration lead to the reconfiguration hang.
2575172	I/Os hung on master node after restarting the slave node.
2149922	Record the disk group import and deport events in system log.
2088426	Re-onlining of disks in disk groups during disk group deport/destroy.
2729911	During a controller or port failure the UDEV removes the associated path information from Dynamic Multi-pathing (DMP). When the paths are being removed the I/O occurring to this disk could still get re-directed to this path, after it has been deleted, leading to an I/O failure.
2495351	VxVMconvert utility was incompatible to migrate data across platforms from native LVM configurations.
2754819	Disk Group (DG) rebuild through 'vxmake –d' loops infinitely if the disk group configuration has multiple objects on a single cache object.
2257850	`vxdiskadm` leaks memory while performing operations related to enclosures.
2735951	Uncorrectable write error is seen on subdisk when SCSI device/bus reset occurs.
2567618	VRTSexplorer coredumps in `checkhbaapi/print_target_map_entry`.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2535716	LVM Volume Group (VG) to VxVM Disk Group conversion fails requesting user to reduce the number of configuration records.

2627056	`vxmake(1M)` command when run with a very large description file fails with following error:
	`VxVM vxmake ERROR V-5-1-10127 creating subdisk <subdisk_name>:`
	`Operation requires transaction`

2680343	Manual disable/enable of paths to an enclosure leads to system panic.
2858853	After master switch, vxconfigd dumps core on old master.
2929300	When `kctune(1M)` or `kcmodule(1M)` commands are run, the following harmless warning message is displayed: `Warning: The file '/usr/conf/mod/dg.o' does not contain valid kernel code. It will be ignored.`
2365486	In 2-nodes SFRAC configuration, after enabling ports systems panics due to improper order of acquire and release of locks.
2495186	With TCP protocol used for replication, I/O throttling happens due to memory flow control.
2527289	In a Campus Cluster setup, storage fault may lead to DETACH of all the configured site. This also results in I/O failure on all the nodes in the Campus Cluster.
2621465	When detached disk after connectivity restoration is tried to reattach gives 'Tagid conflict' error.
2608849	On VVR primary logowner, local I/O starved with heavy I/O load from logclient.
2417546	Raw devices are lost after reboot and cause permissions problem.
2061082	`vxddladm -c assign names` command should work for devices with native support not enabled (VxVM labeled or Third Party Devices).
2635476	DMP (Dynamic Multi Pathing) driver does not automatically enable the failed paths of Logical Units (LUNs) that are restored.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2643634	Message enhancement for a mixed (non-cloned and cloned) diskgroup import.
2666163	A small portion of possible memory leak incase of mix (clone and non-cloned) diskgroup import.
2675538	`vxdisk` resize may cause data corruption.
2674465	Adding/removing new LUNs causes data corruption.
2688747	Under a heavy I/O load on logclient node, the writes on VVR Primary logowner takes a very long time to complete. Writes appear to be hung.
2700486	`vradmind` can core dumps when Primary and Secondary nodes have the same hostname and an active Stats session exists on Primary node.
2700792	The VxVM volume configuration daemon may dump a core during the Cluster Volume Manager(CVM) startup.
2700086	EMC BCV (Not Ready) established devices are resulting in multiple DMP events messages (paths being disabled/enabled).
2698860	`vxassist` mirror failed for thin LUN because `statvfs` failed.
2710579	Do not write backup labels for CDS (Cross-platform Data Sharing) disk - irrespective of disk size
2390998	System panicked during SAN reconfiguration because of the inconsistency in dmp device open count.
2722850	DMP fail over hangs when the primary controller is disabled while I/O activity is ongoing.
2729501	`vxdmpadm` exclude `vxvm path=<>` results in excluding unexpected set of paths.
2423701	Upgrade of VxVM caused change in permissions.
2741240	Invoking "vxdg join" operation during heavy I/O load results in a transaction failure and leaves disks in an intermediate state.
2771452	I/O hung because of hung port deletion.
2428170	I/O hung on Mirror volume and return error on DMP disk, but phydisk (/dev/sdbw) is OK.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2169726	Importing Disk Group using a Non-cloned and cloned disks can lead to data corruption.
2419803	Secondary Site panics in VVR (Veritas Volume Replicator).
2489350	Memory leak in VVR (Veritas Volume Replicator).
2235382	I/O hung in DMP while restoring a path in presence of pending I/Os on local A/P class LUN.
2438426	VxVM fails to correctly discover ZFS LUNs presented via PP after excluding/including libvxpp.so
2484334	System panic occurs in DMP with `dmp_stats_is_matching_group()` stack.
2419486	Data corruption occurs on changing the naming scheme.
2390431	In a Disaster Recovery environment, when DCM (Data Change Map) is active and during SRL(Storage Replicator Log)/DCM flush, the system panics due to missing parent on one of the DCM in an RVG (Replicated Volume Group).
2420386	Corrupted data is seen near the end of a sub-disk, on thin-reclaimable disks with either CDS EFI or sliced disk formats.
2513101	User data corrupted with disk label information.
2530279	Vxesd has been built without any thread locking mechanism.
2148851	Vxdisk resize operation failed to resize the disk which is expanded physically from array console.
2528133	`vxprint -l` command gives following error (along with the output), when multiple Disk Groups' have same `DM_NAME`.
2483053	Primary slave node runs out of memory, system hangs on VRTSvxvm.
2510523	In CVM-VVR configuration, I/Os on "master" and "slave" nodes hang when "master" role is switched to the other node using `vxclustadm setmaster` command.
2524936	Disk Group is disabled after rescanning disks with `vxdctl enable` command.
2432006	Pending read count with `kio cache` is not decremented when read object is locked in transaction.

Table 1-5 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2536667	Slave node panics when private region I/O and `dg deport` operation are executed simultaneously.
2431448	In Cluster Volume Replicator (CVR), I/O hangs while transitioning to DCM mode.
2344186	Volume recovery is not clearing the need sync flag from volumes with DCO in BADLOG state. Thus, nodes are unable to join the cluster.
2185069	In a CVR setup, while the application I/Os are going on all nodes of the primary site, bringing down a slave node results in a panic on the master node.
2419948	Race between the SRL flush due to SRL overflow and the kernel logging code, leads to a panic.
2553729	Disk groups do not get imported and 'clone_disk' flag is seen on non-clone disks after uprade of VxVM.

Table 1-6 lists the Veritas File System issues fixed in this release.

Table 1-6 Veritas File System fixed issues

Incident	Description
2845175	In the `vx_do_getacl()` function, a local variable is accessed without being initializing as a result leading to a panic.
2373266	A thread doing a lookup may get into a loop which can lead to a system hang.
2839871	On systems using delicache and thin-provisioning features, some synchronus I/O may hang.
2841059	VxFS marked for full `fsck` fails to clear the corruption in attribute inode.
2709869	The system panics with a redzone violation while releasing inodes File Input/Output (FIO) statistics structure.
2138025	Internal test noise with replay functionality is failing on local mount file system.
2925918	Checkpoint promotion may lead to deadlock.

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

Incident	Description
2597347	The `fsck` command when run on a 29TB file system crashes with segmentation fault.
2086902	System crashed when spinlock was held too long.
2693010	VxFS patches to not remove formatted/cached man pages.
2715028	`fsadm -d` hang during `vx_dircompact`.
2566875	A write(2) operation exceeding the quota limit fails with an EDQUOT error.
2651922	Performance degradation of 'll' and high SYS% CPU in `vx_ireuse()` .
2670022	Duplicate file names can be seen in a directory.
2696067	When a `getaccess()` command is issued on a file which inherits the default Access Control List (ACL) entries from the parent, it shows incorrrect group object permissions.
2715186	System panic "spinlock: locker forgot to unlock".
2730894	All binary patches.
2745357	Performance enhancements are made for the read/write operation on Veritas File System (VxFS) structural files.
2730759	Poor sequential read performance.
2753944	VxFS hang in `vx_pd_create`.
2779609	The creation of directory hangs during an internal conformance test.
2426648	The `fsck(1M)` operation on a Veritas File Systems (vxfs) can fail with an EINVAL error while validating the inode attributes.
2376382	Vxrestore man page to add `-b option` details.
2246127	Mount should perform read ahead on IAUs.
2515380	The `ff(1M)` command hangs and exits if the program exceeds the memory limit.
2510903	When Veritas File System (VxFS) tries to write to clones loops permanently on HP-UX 11.31, a few commands hang, for example, `bdf(1M)`.

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

Incident	Description
2428964	In postinstall script for VRTSvxfs package, invoke `increase_tunable` without `-i option .`
2526174	Wrong offset calculation affects replication functionality.
2561334	Using `flockfile()` instead of adding new code to take lock on `__fsppadm_enforcesq` file descriptor before writing into it.
2515459	Local mount hangs in `vx_bc_binval_cookie`.
2527578	Panic in `vx_bhash_rele`.
2588593	`df(1M)` shows wrong usage value for volume when large file is deleted.
2528819	VxFS thread create warning messages.
2561739	Class perm changed to "rwx" after adding user ACL entry with null perm.
2492304	`Find` command displays duplicate directory entries.
2599590	Expanding or shrinking a DLV5 file system using the fsadm(1M)command causes a system panic.
2534693	A man page for `vx_dexh_sz(5)` tunable is not available.
2631276	Lookup fails for the file which is in partitioned directory and is being accessed using its VxFS namespace extension name.
2271797	Internal Noise Testing with locally mounted VxFS filesystem hit an assert `f:vx_getblk:1a .`
2350956	Internal noise test on locally mounted filesystem exited with error message `run_fsck : Failed to full fsck cleanly` on SLES10_SP4.
2326037	Internal Stress Test on cluster file system with clones failed while writing to file with error ENOENT.
2555198	The sendfile() interface does not create the DMAPI events for Hierarchical Storage Management (HSM) on Veritas File System (VxFS).

Table 1-7 lists the Veritas Enterprise Administrator issue fixed in this release.

Table 1-7 Veritas Enterprise Administrator fixed issue

Incident	Description
2535298	With PHCO_42182 enclosure, controllers are not visible with Veritas Enterprise Administrator.

Veritas Storage Foundation 5.1 SP1 RP1 fixed issues

This section lists the fixed issues in 5.1 SP1 RP1 release.

Veritas Storage Foundation fixed issues

Table 1-8 lists the Veritas Volume Manager issues fixed in this release.

Table 1-8 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.

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

Incident	Description

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).

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.

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

Incident	Description
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.
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.

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

Incident	Description
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.
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.

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

Incident	Description
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.
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.

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

Incident	Description
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.
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`).

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

Incident	Description
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`.
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.

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

Incident	Description
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.

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.

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

Incident	Description
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.
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 1-9 lists the Veritas File System issues fixed in this release.

Table 1-9 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.

Table 1-9 Veritas File System fixed issues *(continued)*

Incident	Description
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).

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.

Table 1-9 Veritas File System fixed issues *(continued)*

Incident	Description
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.
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.

Table 1-9 Veritas File System fixed issues *(continued)*

Incident	Description
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`

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).

Table 1-9 Veritas File System fixed issues *(continued)*

Incident	Description
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.
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.

Table 1-9 Veritas File System fixed issues *(continued)*

Incident	Description
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.
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 1-10 lists the Veritas Perl Redistribution issue fixed in this release.

Table 1-10 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 1-11 lists the Veritas Operations Manager issue fixed in this release.

Table 1-11 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 1-12 lists the Veritas Enterprise Administrator issue fixed in this release.

Table 1-12 Veritas Enterprise Administrator fixed issue

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

Veritas Storage Foundation 5.1 SP1 fixed issues

This section lists the Veritas Storage Foundation fixed issues in 5.1 SP1 release.

Fixed issues related to installation and upgrades

This section describes the incidents that are fixed related to installation and upgrades in this release.

Table 1-13 Fixed issues related to installation and upgrades

Incident	Description
1851632	The software media can now be ejected after installing, configuring, and starting any of the products that contain Veritas Volume Manager.
1852746	Installing the VRTSvxfs, VRTSdbed, and VRTSodm packages no longer fails on the second node of a cluster when only the required packages are installed.

Veritas Storage Foundation fixed issues

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

Veritas File System fixed issues

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

Table 1-14 Veritas File System fixed issues

Incident	Description
1283531	Fixed the cause of a hang when mounting a frozen file system.
1176062	Fixed the cause of an assertion panic.
1539171	Fixed the cause of a panic in the `volkio_to_kio_copy`() call.
1805046	Fixed an incorrect alert generation from VxFS when the file system usage threshold is set.
1798708	Fixed the cause of a system panic in `inctext`, in which `VTEXT` was not set and `tcount` was greater than 0.
1903977	Fixed an issue in which a panic could happen because of a mutex getting destroyed while the protected structure was still in use.
1922948	You can now mount deprecated disk layout versions so that they can be upgraded.
1926141	Direct I/O can now be enabled using the `mincache` and `convosync` mount options with the base VxFS license.
1935374	Storage Checkpoint creation now properly fails to assign a metadata allocation policy on a data-only volume.
1936959	Setting vxfs_bc_bufhwm higher than the available physical memory now reports an error.
1946063	Fixed a performance issue in `fsadm` where `fsadm` kept relocating and copying already reorganized regions of a file in subsequent passes.
1630971	Interchanged the order in which `fcache_vn_destroy`() and `VFS_TEARDOWN_STACK`() are called to avoid a panic.
2017776	The virtual memory area is no longer destroyed if there are active mappings on the vnode Fix to during a force unmount.
2018439	The `fsppadm` command no longer dumps core if a volume does not have placement tags.
1874185	Added quota support for the user "nobody".

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

Incident	Description
1975547	The `sar -v` command now properly reports VxFS inode table overflows.
2014708	Fixed the cause of ENOTBLK being returned via the async driver due to fdd not being a block device.
2028782	Fixed an issue in which `Q_SETQUOTA` was not setting the current usage using the `quotactl`() API.
2017776	Fixed an issue in which the volume manager area was destroyed when spinlock was held.
2068824	Fixed the cause of a hang that was due to a disowned beta semaphore.
2098371	Fixed a performance issue in which write I/O performance degraded rapidly when the size of a file reached 64 MB and I/O size was not a multiple of 64 bytes.
2106668	Fixed an issue in which in some cases EFBIG was returned soon after resizing a file system through the `fsadm` command.
2111614	Fixed an issue in which file modification time was not updated when the `O_SYNC` and `nodatainlog` mount options were used.
2149407	Fixed an issue in which modification and access times were not getting updated through mmap and msync.
2163013	Fixed an issue in which the `odmstat` command was showing very high average IO time.

Veritas Volume Manager fixed issues

This section describes the incidents that are fixed in Veritas Volume Manager in this release. This list includes Veritas Volume Replicator and Cluster Volume Manager fixed issues.

Table 1-15 Veritas Volume Manager fixed issues

Incident	Description
150476	Add T for terabyte as a suffix for volume manager numbers
248925	If vxdg import returns error, parse it

Table 1-15 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
311664	vxconfigd/dmp hang due to a problem in the `dmp_reconfig_update_cur_pri`() function's logic
321733	Need test case to deport a disabled dg.
339282	Failed to create more than 256 config copies in one DG.
597517	Tunable to initialize EFI labeled >1tb PP devices.
1089875	Increasing vol_maxspecialio to 1 MB on HP-UX.
1097258	vxconfigd hung when an array is disconnected.
1239188	Enhance vxprivutil to enable, disable, and display config+log copies state.
1301991	When vxconfigd is restarted with -k option, all log messages are sent to stdout. syslog should be the default location.
1321475	Join Failure Panic Loop on axe76 cluster.
1405756	CVM: Add support to set PFTO values cluster-wide.
1441406	'vxdisk -x list' displays wrong DGID.
1458792	After upgrade from SF5.0mp1 to SF5.0mp3, *unit_io and *pref_io was set to 32m.
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.
1485075	DMP sending I/O on an unopened path causing I/O to hang
1504466	VxVM: All partitions aren't created after failing original root disk and restoring from mirror.
1513385	VVR:Primary panic during autosync or dcm replay.
1528121	FMR: wrong volpagemod_max_memsz tunable value cause buffer overrun
1528160	An ioctl interrupted with EINTR causes frequent vxconfigd exits.
1586207	"vxsnap refresh" operations fail occasionally while data is replicating to secondary.
1589022	Infinite looping in DMP error handling code path because of CLARIION APM, leading to I/O hang.
1594928	Avoid unnecessary retries on error buffers when disk partition is nullified.

Table 1-15 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
1662744	RVG offline hung due to I/Os pending in TCP layer
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1665094	Snapshot refresh causing the snapshot plex to be detached.
1713670	'vxassist -g <dg-name> maxsize' doesn't report no free space when applicable
1715204	Failure of vxsnap operations leads to orphan snap object which cannot be removed.
1766452	vradmind dumps core during collection of memory stats.
1792795	Supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
1825270	I/O failure causes VCS resources to fault, as dmpnode get disabled when storage processors of array are rebooted in succession
1825516	Unable to initialize and use ramdisk for VxVM use.
1826088	After pulling out the Fibre Channel cables of a local site array, plex becomes DETACHED/ACTIVE.
1829337	Array firmware reversal led to disk failure and offlined all VCS resources
1831634	CVR: Sending incorrect sibling count causes replication hang, which can result in I/O hang.
1831969	VxVM: ddl log files are created with world write permission
1835139	I/Os hung after giveback of NetApp array filer
1840673	After adding new LUNs, one of the nodes in 3 node CFS cluster hangs
1846165	Data corruption seen on cdsdisks on Solaris-x86 in several customer cases
1857558	Need to ignore jeopardy notification from GAB for SFCFS/RAC, since oracle CRS takes care of fencing in this stack
1857729	CVM master in the VVR Primary cluster panicked when rebooting the slave during VVR testing
1860892	Cache Object corruption when replaying the CRECs during recovery
1869995	VVR: Improve Replication performance in presence of SO snapshots on secondary.

Table 1-15 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1874034	Race between modunload and an incoming IO leading to panic
1880279	Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment.
1881336	VVR: Primary node panicked due to race condition during replication
1884070	When running iotest on a volume, the primary node runs out of memory
1897007	vxesd coredumps on startup when the system is connected to a switch which has more than 64 ports
1899688	VVR: Every I/O on smartsync enabled volume under VVR leaks memory
1899943	CPS based fencing disks used along with CPS servers does not have coordinator flag set
1901827	vxdg move fails silently and drops disks.
1907796	Corrupted Blocks in Oracle after Dynamic LUN expansion and vxconfigd core dump
1915356	I/O stuck in vxvm causes a cluster node panic.
1933375	Tunable value of 'voliomem_chunk_size' is not aligned to page-size granularity
1933528	During Dynamic reconfiguration vxvm disk ends up in error state after replacing physical LUN.
1936611	vxconfigd core dump while splitting a diskgroup
1938907	WWN information is not displayed due to incorrect device information returned by HBA APIs
1946941	vxsnap print shows incorrect year
1954062	vxrecover results in os crash
1956777	CVR: Cluster reconfiguration in primary site caused master node to panic due to queue corruption
1969526	Panic in voldiodone when a hung priv region I/O comes back
1972848	vxconfigd dumps core during upgradation of VxVM

Table 1-15 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
1974393	Cluster hangs when the transaction client times out
1982178	vxdiskadm option "6" should not list available devices outside of source diskgroup
1982715	vxclustadm dumps core during memory re-allocation.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout
1992872	vxresize fails after DLE.
1993953	CVM Node unable to join in Sun Cluster environment due to wrong coordinator selection
1998447	Vxconfigd dumps core due to incorrect handling of signal
1999004	I/Os hang in VxVM on linked-based snapshot
2002703	Misleading message while opening the write protected device.
2009439	CVR: Primary cluster node panicked due to queue corruption
2010426	Tag setting and removal do not handle wrong enclosure name
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
2016129	Tunable to disable OS event monitoring by vxesd
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1
2021737	vxdisk list shows HDS TrueCopy S-VOL read only devices in error state.
2025593	vxdg join hang/failure due to presence of non-allocator inforecords and when tagmeta=on
2027831	vxdg free not reporting free space correctly on CVM master. vxprint not printing DEVICE column for subdisks.
2029480	Diskgroup join failure renders source diskgroup into inconsistent state
2029735	System panic while trying to create snapshot
2034564	I/Os hung in serialization after one of the disks which formed the raid5 volume was pulled out
2036929	Renaming a volume with link object attached causes inconsistencies in the disk group configuration

Table 1-15 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2038137	System panics if volrdmirbreakup() is called recursively.
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.
2040150	Existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration
2052203	Master vold restart can lead to DG disabled and abort of pending transactions.
2052459	CFS mount failed on slave node due to registration failure on one of the paths
2055609	Allocation specifications not being propagated for DCO during a grow operation
2060785	Primary panics while creating primary rvg
2061066	vxisforeign command fails on internal cciss devices
2061758	Need documentation on list of test suites available to evaluate CDS code path and verification of the code path.
2063348	Improve/modify error message to indicate its thin_reclaim specific
2070531	Campus cluster: Couldn't enable site consistency on a dcl volume, when trying to make the disk group and its volumes siteconsistent.
2075801	VVR: "vxnetd stop/start" panicked the system due to bad free memory
2076700	VVR: Primary panic due to NULL pointer dereference
2094685	Diskgroup corruption following an import of a cloned BCV image of a SRDF-R2 device
2097320	Events generated by dmp_update_status() are not notified to vxconfigd in all places.
2105547	Enabling tagmeta=on on a disk group no longer causes a delay in disk group split/join operations.
2105722	VVR: I/O hang on Primary with link-breakoff snapshot
2112568	System panics while attaching back two Campus Cluster sites due to incorrect DCO offset calculation
2122009	vxddladm list shows incorrect hba information after running vxconfigd -k

Table 1-15 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2126731	vxdisk -p list output is not consistent with previous versions
2131814	VVR: System panic due to corrupt sio in _VOLRPQ_REMOVE

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

Table 1-16 lists the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in 5.1 SP1 RP2.

Table 1-16 Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Incident	Description
2664050	vxdbd core dump after running for about five hours.
2816482	Memory consumed by vxdbd continues to increase with time, even when the daemon is totally idle.
2664794	Sometimes after startup vxdbd process dumps core due to segmentation fault.
1957142	Sometimes `reverse_resync_abort` and `reverse_resync_commit` operation fail with an error while shutting down the database instance.

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

This section describes fixed issues in 5.1 SP1 RP1 release.

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

Table 1-17 lists the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in 5.1 SP1 RP1.

Table 1-17 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.

Table 1-17 Veritas Storage Foundation for Databases (SFDB) tools fixed issues *(continued)*

Incident	Description
2361363	Running the `qio_convertdbfiles`(1m) command results in the following error: `/opt/VRTSdbed/bin/qio_convertdbfiles: Command not found.`

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

This section lists the Veritas Storage Foundation for Databases (SFDB) tools fixed issues for 5.1 SP release.

Storage Foundation for Databases (SFDB) tools fixed issues for 5.1 SP1

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

Table 1-18 Veritas Storage Foundation for Databases tools fixed issues

Incident	Description
1873738	The dbed_vmchecksnap command may fail
1399393	Clone command fails on an Oracle RAC database
1736516	Clone command fails for instant checkpoint on Logical Standby database
1789290	dbed_vmclonedb -o recoverdb for offhost fails for Oracle 10gr2 and prior versions
1810711	Flashsnap reverse resync command fails on offhost flashsnap cloning

Veritas Cluster Server 5.1 SP1 RP2 fixed issues

Table 1-19 lists the issues fixed in 5.1 SP1 RP2.

Table 1-19 Veritas Cluster Server 5.1 SP1 RP2 fixed issues

Incident	Description
1919382	The mount agent fails to detect the mounted file system if either the BlockDevice or MountPoint attribute has a trailing forward slash.

Table 1-19 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Incident	Description
2199924	Veritas Cluster Server (VCS) is unable to load configuration if the size of an attribute is greater than 4 KB.
2292294	When you run a process-usage tracing command on the gablogd daemon, the command may freeze. Following are some of the commands where you may experience this behavior: ■ `pfiles` ■ `lsof` ■ `procfiles` ■ `truss`
2528470	The `preonline_ipc` script does not support any of the resources other than IP resources.
2531558	Graceful shutdown of a node triggers I/O fencing race condition on peer nodes.
2558988	When a system rejoins a cluster after fault recovery, the `CurrentLimits` attribute value of that system is not correctly reflected in the other member nodes of the cluster.
2561722	The agent does not try to re-register the resource with IMF to the value specified in the RegisterRetryLimit key of the IMF attribute.
2593173	The online entry point of the DiskGroup agent does not detect the serial split brain situation and fails to log a warning.
2636874	The system panics while monitoring Veritas File System (VxFS) resources.
2660011	The resource moves to FAULTED state even if the `ManageFaults` attribute value is set to NONE at the service group level and the service group faults if the resource is critical.
2684818	VCS may fail to implement the configured values of the following attributes: ■ `PreOnline` ■ `ContainerInfo` ■ `TriggersEnabled` ■ `AutoStartList`
2692173	The service groups come online on the same node even if you set an online remote dependency between two service groups.
2710892	One or more nodes may not be able to join an already running fencing cluster.

Table 1-19 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Incident	Description
2724829	Sybase and SybaseBk agents does not support PRON IMF.
2728802	If the directory or filename specified as part of the `httpdDir` attribute does not exist on the cluster node, the Apache agent cannot perform detail monitoring or start the HTTP server. The Apache agent logs an error with the following message ID: `V-16-10061-20495`
2729816	VCS does not fail over a service group even if the value of the `OnlineRetryLimit` attribute is exhausted.
2729867	When High Availability Daemon (HAD) goes down and a node at the primary site crashes, VCS is unable to fail over a global service group.
2731133	When the NFSRestart resource is taken offline, it forcefully stops the automountd process.
2732228	VCS does not shut down if you use the `init` script.
2735410	HAD floods the engine logs and after sometime dumps core and restarts.
2741299	The CmdSlave process dumps core with a SIGSEGV signal.
2746802	A failover service group is not brought online after VCS starts.
2746816	GAB terminates the HAD process.
2788059	The system does not panic on loss of storage connectivity even after the `PanicSystemOnDGLoss` attribute of the DiskGroup resource is set to 1.
2804891	When you configure the Low Latency Transport (LLT) module in the verbose mode, an incorrect error message may appear on the console with the following message ID: `V-14-2-15300` LLT then encounters a segmentation fault and the I18N module dumps a core file.
2812400	The VXFEN module fails to start if agile disk naming scheme is used.
2818567	During LLT configuration, an internal check for duplicate cluster IDs may lead to either a private network slowdown or panic in the VCS cluster.

Table 1-19 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Incident	Description
2831283	Unconfiguring GAB immediately after a network failure may cause the system to panic. On a system with Solaris operating system you may see the following error message: `BAD TRAP: type=31 rp=2a10d4cf530 addr=28mmu_fsr=0 occurred in module "gab" due to a NULL pointer dereference` On other operating systems, you can see `NULL pointer dereference` as part of the error message.
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, the hagrp command line utility gives incorrect output when you use the following options: ■ `-clear` ■ `-flush` ■ `-state`
2855755	The VXFEN module may fail to start or the online coordination point migration may fail if the CP server is used as a coordination point for the first time for a node.
2896402	After you perform an online or offline operation on the resource, the agent logs an error message in the agent log or engine log.
2411651	In the communication between GAB and GAB clients, GAB drops certain packets due to size restrictions and may cause the cluster to hang.

Veritas Cluster Server 5.1 SP1 RP1 fixed issues

Table 1-20 lists the issues fixed in 5.1 SP1 Rolling Patch 1.

Table 1-20 Veritas Cluster Server 5.1 SP1 RP1 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.

Table 1-20 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Incident	Description
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
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.

Table 1-20 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Incident	Description
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).
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.

Table 1-20 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Incident	Description
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).
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).

Table 1-20 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Incident	Description
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.
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".

Table 1-20 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Incident	Description

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.
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.

Table 1-20 Veritas Cluster Server 5.1 SP1 RP1 fixed issues *(continued)*

Incident	Description
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).
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 Cluster Server 5.1 SP1 fixed issues

The issues fixed in VCS 5.1 SP1 are as follows:

Table 1-21 VCS 5.1 SP1 fixed issues

Incident	Description
1142970	VCS logs an error "SSL Handshake failed" if the client creates a channel and then disconnects.
1179782	The haclus -state command displays inconsistent output in four-node cluster configurations
1214140	HAD reports incorrect CPU utilization
1404384	HAD crashes while switching over Global group and PreSwitch is set to TRUE.
1456724	Group switch/failover logic does not complete if the parent group gets autodisabled in between.
1458338	Even though /var/VRTSvcs/diag/had directory was empty, VCS was moving it to /var/VRTSvcs/diag/had.<*timestamp*>, and printed the following kind of a message on the console. VCS NOTICE V-16-1-53021 Diagnostics directory moved to /var/VRTSvcs/diag/had.1227075736, please check its contents and contact VERITAS Support. This was due to incorrect handling of directories.
1479349	On disabling the array side switch port, the system panics intermittently.
1531512	The Oracle agent picks up only the last corresponding action from oraerror.dat ignoring the previous actions. Even though the log shows the errors, the resource does not move to FAULT state.
1531720	Seeding of a port a does not seed other ports.

Table 1-21 VCS 5.1 SP1 fixed issues *(continued)*

Incident	Description
1537433 1403471	In secure global cluster environment, VCS took a long time to detect a cluster fault of a remote cluster. This was because in secure connection, a socket was opened as a blocking socket.
1539089	The agent framework leaked memory if there is continuous logging into agent's log file.
1587173	The LC_ALL value was set to empty string by the hastart script even though it was not present in the environment.
1588784	VCS engine does not support system names starting with numbers.
1634031	If a resource faults in a planned offline of a global group on node1 in a primary cluster, followed by planned online of the group in a remote cluster, the group can go online in the primary cluster if node1 gets rebooted, resulting in global concurrency violation. This may lead to data corruption.
1710470	Had may crash in a global cluster environment. In a global cluster environment, if the SystemList of a global group is modified to add new system in C1 then ResourceInfo attribute for all remote resources of this group, should get set to default value in C2 . This was not happening and hence hares -display for remote resources in C2 was causing _had to get SEGV.
1780722	For a SAMBA GROUP, "netbios" resource, with CIDR address for interface, fails to come ONLINE.
1789808	Cluster does not accept HA commands without reboot of whole cluster.
1795151	Global group fails to come online on the DR site with a message that it is in the middle of a group operation
1829180	The VCSAG_SU() function from vcsag_i18n_inc.sh file, has incorrect options to execute the `su` command.
1830978	HAD may crash while sending notifications to notifier if the send fails. This happens due to incorrect data-access.
1851078	RemoteGroup agent faults when set up for monitor only the local service group is taken offline
1852521	DiskGroupSnap agent assumes all nodes are part of a campus cluster configuration

Table 1-21 VCS 5.1 SP1 fixed issues *(continued)*

Incident	Description
1862229	Symptom: Resource coordpoint keeps fluctuating between online/faulted status without actual failure. Description: The coordpoint resource reports a faulted status only when the number of coordination points with missing keys (of the local node) exceeds a predefined fault tolerance value. Though no coordpoint with missing keys was found, the agent reports as one or more of the coordination points are not accessible/missing out registrations. The fault is not with the cpagent functionality, but is observed that the `cpsat listpd`command was failing sometimes when more than the predefined processes execute the command simultaneously (using threads). Resolution: Moved out the code to execute the command only once much before forking the threads and store the value to be used by all them later. Thus, preventing each thread to simultaneously execute it which is leading to a failure sometimes.
2097935	Need strict host name matching in coordination point installer.
2130967	Health check monitoring is performed in ASMInst agent, if newly added MonitorOption attribute is set to 0.

Table 1-22 Bundled agents 5.1 SP1 fixed issues

Incident	Description
251704	The index number of the virtual interfaces changes after a failover. As a result applications using the interfaces may face communication failure.
252354	Application agent does not pass the value of CleanReason to the script that the Agent executes as per the value of the CleanProgram attribute.
2001039	In a frozen local service group, if a resource goes offline outside VCS control, a RemoteGroup resource also goes offline.
2019904	The OS determines the interface that the notifier uses to send packets. The user must be optionally able to specify the interface. This is an enhancement.
2045972	If you upgrade to VCS 5.1, Application resources may go to the FAULTED state.

Table 1-23 VCS engine 5.1 SP1 fixed issues

Incident	Description
254693	If the OnlineRetryLimit and PreOnline attributes are set for a group, then the group does not fail over in case of a fault.
970971	At unfreeze of a failover group, VCS does not evaluate the group for Concurrency violation.
1074707	A global group goes online on a remote site despite a concurrency violation. This behavior is observed if a related group resource goes intentionally online on the remote site.
1472734	VCS must log an alert message to increase the value of the ShutdownTimeout attribute on a multi-CPU computer.
1634494	If the GlobalCounter attribute does not increase at the configured interval, VCS must report the failure.
1859387	When a parent group completely faults in a system zone, an online-local-hard (OLH) child group fails over to another system. This behavior is observed only if the child group is marked for manual failover in campus cluster. That is, the value of the AutoFailover attribute for the group is equal to 2.
1861740	The Subject line of the SMTP notification email must contain the Entity name. This is an enhancement.
1866434	When hashadow attempts to restart the High Availability Daemon (HAD) module, if the `/var/VRTSvcs/lock/.hadargs` file does not exist, hashadow gets the SIGSEGV signal.
1874261	VCS sets the MonitorOnly attribute of the resources in a frozen service group to 0. This behavior occurs even if the ExternalStateChange attribute is not set to OfflineGroup for a resource that goes intentionally offline.
1874533	If a resource, with the ExternalStateChange attribute set to OfflineGroup, goes intentionally offline, VCS sets the MonitorOnly attribute of all resources in a frozen service group to 0. VCS must set the MonitorOnly attribute only when required.
1915908	The `hares` command lets you have the "." special character in a resource name.
1958245	Notifier Agent is unable to get the local IP address in the linked-based IPMP.
1971256	If a service group faults during failover, then VCS may not honor the Prerequisites/Limits attributes on the target system.
2022123	The `notifier` command must let you specify the originating source IP address. This is an enhancement.

Table 1-23 VCS engine 5.1 SP1 fixed issues *(continued)*

Incident	Description
2061932	If you override a static attribute for a resource with an empty type definition, then VCS dumps duplicate entries for the attribute in the configuration file.
2081725	After a child group autostarts, VCS does not bring a partially-online parent to ONLINE state.
2083232	If you configure an online-local dependency between two parallel groups, then the parent service group does not AutoStart on all nodes.
2084961	If the VCS configuration includes a file with an absolute pathname, then you cannot load templates from the cluster manager GUI.
2111296	In rare cases, VCS tries to bring online a failover service group that is already online on another node.
2128658	If you add a script-based WAN heartbeat, the heartbeat agent generates a core file and fails to report status.

Table 1-24 Enterprise agents 5.1 SP1 fixed issues

Incident	Description
776230	The Sybase agent offline script must support the databases that require a longer time to shut down. This is an enhancement.
2117378	When the monitor entry function encounters an invalid process ID (PID) in a PID file, the function fails to detect the exit status of the process. This behavior occurs if the PID has an empty line below it. The function also accordingly fails to detect the correct state of a failed DB2 Connect instance.

Table 1-25 LLT, GAB, and I/O fencing 5.1 SP1 fixed issues

Incident	Description
1908938	[GAB] In a large cluster, cascaded lowest node failures result in GAB panic during sequence space recovery.
1840826	[GAB] Prevent 'gabconfig -c' while port 'a' is in the middle of iofence processing.
1989645	[GAB] Whenever there is a disparity in connected memberships, GAB master must wait for maximum `gab_conn_wait` time. However, when the GAB master is a new joinee it does not wait which caused wrong iofence message being sent to healthy nodes in the cluster.

Table 1-25 LLT, GAB, and I/O fencing 5.1 SP1 fixed issues *(continued)*

Incident	Description
2066020	[LLT] The `dlpiping` utility exits with an error similar to "dlpiping: send ECHO_REQ failed."
2005045	[LLT] The `hastart` command fails to start HAD on one of the nodes with message “GabHandle::open failed errno = 16” in syslog after HAD is stopped on all the nodes in the cluster simultaneously.
1859023	[LLT] The `lltconfig -T query` command displays a partially incorrect output
1846387 2084121	[Fencing] The `vxfenswap` and the `vxfentsthdw` utilities fail when rsh or ssh communication is not set to the same node.
1922413	[Fencing] The `vxfentsthdw` utility should detect storage arrays which interpret NULL keys as valid for registrations/reservations.
1847517	[Fencing] The `vxfenswap` utility has an incorrect usage message for -n option
1992560	[Fencing] The `vxfentsthdw` utility uses scp to communicate with the local host.
1512956	[Fencing] The `vxfenclearpre` utility displays error messages

Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 fixed issues

Table 1-26 lists the Veritas Storage Foundation Cluster File System (SFCFS) issues fixed in this release.

Table 1-26 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
2910827	Installation of `VRTScavf` patch failed on PA machine.
2684573	Enhancement request for force option of `cfsumount` command.
2161660	CFS stress S1 on SPARC hit `assert vx_hlock_getpbdata_return:1 .`
2433934	VirtualStore performance discrepency between CFS and standalone VxFS using NFS.

Table 1-26 Veritas Storage Foundation Cluster File System fixed issues *(continued)*

Incident	Description
2565400	Poor read performance with DSMC (TSM) backup on CFS filesystems.
2630954	The `fsck(1M)` command exits during an internal CFS stress reconfiguration testing.
2669724	CFSMountAgent core dump due to assertion failure in `VCSAgThreadTbl::add()`
2824895	VCSCVMqa "cfsumount" test getting failed.

Veritas Storage Foundation Cluster File System 5.1 SP1 RP1 fixed issues

This section describes fixed issues in 5.1 SP1 RP1 release.

Veritas Storage Foundation Cluster File System fixed issues

Table 1-27 lists the Veritas Storage Foundation Cluster File System (SFCFS) issues fixed in this release.

Table 1-27 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.
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.

Table 1-27 Veritas Storage Foundation Cluster File System fixed issues *(continued)*

Incident	Description
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 1-28 lists the Veritas Group Lock Manager issues fixed in this release.

Table 1-28 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 1-29 lists the Veritas Cluster Server Agents for Cluster File System issues fixed in this release.

Table 1-29 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 Cluster File System 5.1 SP1 fixed issues

This section lists the Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 release.

Veritas Storage Foundation Cluster File System fixed issues

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

Table 1-30 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
1856719	Fixed an smap update issue due to which 'df' and 'rm' commands can hang.
1544221	Optimize getattr call to speedup the case when binaries are mmaped from multiple nodes.

Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 fixed issues

Table 1-31 lists the issues fixed in SF Oracle RAC 5.1 SP1 RP2.

Table 1-31 SF Oracle RAC 5.1 SP1 RP2 fixed issues

Incident number	Description
2853857	Setting the `usevirtualIP` attribute to 1 overwrites the IP address on the virtual interface of some of the nodes of in the cluster.
2853860	The PrivNIC resource faults when the `useVirtualIP` attribute is set to 1.
2332314	Nodes running `LM noise.fullfsck` with ODM report stack trace errors.
2603511	Database operations can fail on nodes running Oracle RAC 11.2.0.3 and later. The following message is reported in the system logs: `ODM ERROR V-41-4-1-105-22 Invalid argument`

The issues fixed in previous releases are as follows:

5.1 SP1 RP1 There are no fixes specific to this release.

5.1 SP1 See "Veritas Storage Foundation for Oracle RAC 5.1 SP1 fixed issues" on page 70.

Veritas Storage Foundation for Oracle RAC 5.1 SP1 fixed issues

The issues fixed in this release are as follows:

SF Oracle RAC issues See "Issues fixed in SF Oracle RAC 5.1 SP1" on page 70.

SFDB tools issues See "Storage Foundation for Databases (SFDB) tools fixed issues" on page 71.

Issues fixed in SF Oracle RAC 5.1 SP1

Table 1-32 lists the issues fixed in SF Oracle RAC 5.1 SP1.

Table 1-32 Fixed issues in SF Oracle RAC 5.1 SP1

Incident number	Description
1439223	The `lmxpollport` function times out with an incorrect timeout value on systems that have been running for more than 410 days. The issue was caused by the lbolt variable, which resets after 410 days.
1844422	The CSSD agent configuration fails if the OCR files are placed in a directory on CFS.
1855800	The SF Oracle RAC installer may fail to configure the CSSD resource if the `/etc/hosts` file contains commented IP address and host name entries.
1879412	The Low Latency Transport Multiplexer (LMX) module may cause the system to panic with the following message: `kernel heap corruption detected` Incorrect manipulation of the request queue corrupts the memory and causes the system to panic. When the last request is removed from the queue, the queue pointers are not updated correctly.
1927920	The PrivNIC and MultiPrivNIC agents do not support MTU size settings in the `main.cf` configuration file.
1938799	When LMX registers with LLT, LLT calls the `lmxlltxcanput` function before delivering any packet to LMX, causing performance overheads.

Table 1-32 Fixed issues in SF Oracle RAC 5.1 SP1 *(continued)*

Incident number	Description
2038617	The installation and configuration check "LLT links' speed and auto negotiation settings" fails when LLT is configured over UDP/TCP.
2042817	Oracle Clusterware fails to restart due to incorrect regisration with VCSMM.
2045700	The SF Oracle RAC installer fails to validate the length of the disk group and volumes names used for OCR and voting disk.
2089351	The CSSD agent incorrectly reports OFFLINE even when one of the cssd, crsd, or evmd daemons is still running, causing the nodes to panic with the following message: `Oracle CRS failure.  Rebooting for cluster integrity.`
2138574	In a node with 16 clusters, the PrivNIC agent fails to fail over the IP address for nodes with the NodeID value greater than 10.
2237829	When a node panics on a heavily loaded cluster, the Oracle instance on another node crashes.

Storage Foundation for Databases (SFDB) tools fixed issues

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

Table 1-33 Veritas Storage Foundation for Databases tools fixed issues

Incident	Description
1873738	The dbed_vmchecksnap command may fail
1399393	Clone command fails on an Oracle RAC database
1736516	Clone command fails for instant checkpoint on Logical Standby database
1789290	dbed_vmclonedb -o recoverdb for offhost fails for Oracle 10gr2 and prior versions
1810711	Flashsnap reverse resync command fails on offhost flashsnap cloning

Software limitations in this release

This section describes the software limitations in this release.

- Veritas Storage Foundation:
  See "Veritas Storage Foundation 5.1 SP1 RP2 software limitations" on page 72.
- Veritas Storage Foundation for Databases (SFDB) Tools:
  See "Veritas Storage Foundation for Databases tools 5.1 SP1 RP2 software limitations" on page 76.
- Veritas Cluster Server:
  See "Veritas Cluster Server 5.1 SP1 RP2 software limitations" on page 77.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 software limitations" on page 84.
- Veritas Storage Foundation for Oracle RAC:
  See "Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 software limitations" on page 85.

Veritas Storage Foundation 5.1 SP1 RP2 software limitations

There are no software limitations in this release.

The software limitations for previous releases are as follows:

5.1 SP1 RP1	See "Veritas Storage Foundation 5.1 SP1 RP1 software limitations" on page 72.
5.1 SP1	See "Veritas Storage Foundation 5.1 SP1 software limitations" on page 74.

Veritas Storage Foundation 5.1 SP1 RP1 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 1-34

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 5.1 SP1 software limitations

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

Veritas Volume Manager 5.1 SP1 software limitations

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

DMP settings for NetApp storage attached environment

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

Table 1-35

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
```

Veritas Volume Replicator software limitations

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

Replication in a shared environment

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

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

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

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 5.1SP1 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 File System software limitations

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

Recommended limit of number of files in a directory

To maximize VxFS performance, do not exceed 100,000 files in the same directory. Use multiple directories instead.

Veritas Storage Foundation for Databases tools 5.1 SP1 RP2 software limitations

There are no software limitations in this release.

The software limitations for previous releases are as follows:

5.1 SP1 RP1	See "Veritas Storage Foundation for Databases tools 5.1 SP1 RP1 software limitations" on page 76.
5.1 SP1	See "Veritas Storage Foundation for Databases tools 5.1 SP1 software limitations" on page 77.

Veritas Storage Foundation for Databases tools 5.1 SP1 RP1 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 RP2, upgrade the Oracle binaries and database to version 11.1.0.7, and move to SP1 and then upgrade to SP1 RP2.

Veritas Storage Foundation for Databases tools 5.1 SP1 software limitations

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

Veritas Storage Foundation for Databases tools software limitations

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

Oracle Data Guard in an Oracle RAC environment

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

Upgrading if using Oracle 11.1.0.6

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

Veritas Cluster Server 5.1 SP1 RP2 software limitations

The are no software limitations in this release.

The software limitations from previous releases are as follows:

5.1 SP1 RP1	See "Veritas Cluster Server software limitations" on page 77.
5.1 SP1	See "Veritas Cluster Server 5.1 SP1 software limitations" on page 78.

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 Cluster Server 5.1 SP1 software limitations

This section describes the software limitations in Veritas Cluster Server 5.1 SP1.

Limitations related to the bundled agents

Limitations related to global clusters

Systems in a cluster must have same system locale setting

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

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

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

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

GAB panics the systems while VCS gets diagnostic data

On receiving a SIGABRT signal from GAB, VCS engine forks off vcs_diag script. When VCS engine fails to heartbeat with GAB, often due to heavy load on the system, the vcs_diag script does a sys req to dump the stack trace of all processes in the system to collect diagnostic information. The dump of stack trace is intended to give useful information for finding out which processes puts heavy load. However, the dumping puts extra load on the system that causes GAB to panic

the system in such heavy loads. See *Veritas Cluster Server User's Guide* for more information. [383970]

Workaround: Disable the vcs_diag script. To disable, rename the file /opt/VRTSvcs/bin/vcs_diag to /opt/VRTSvcs/bin/vcs_diag.backup.

Using agents in NIS

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can hang if the host is disconnected from the network. If such a program is used as an agent entry point, a network disconnect can cause the entry point to hang and possibly time out. For example, if the host is configured to use NIS maps as a client, basic commands such as ps -ef can hang if there is network disconnect. Symantec recommends creating users locally. To reflect local users, configure:

/etc/nsswitch.conf

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

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

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

Workaround:

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

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

Limitations with DiskGroupSnap agent

The DiskGroupSnap agent has the following limitations:

- The DiskGroupSnap agent does not support layered volumes. [1368385]

- If you use the Bronze configuration for the DiskGroupSnap resource, you could end up with inconsistent data at the secondary site in the following cases [1391445]:
  - After the fire drill service group is brought online, a disaster occurs at the primary site during the fire drill.
  - After the fire drill service group is taken offline, a disaster occurs at the primary while the disks at the secondary are resynchronizing.

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

Fire drill does not support volume sets

The fire drill feature for testing fault readiness of a VCS configuration supports only regular Volume Manager volumes. Volume sets are not supported in this release.

Manually removing VRTSat package erases user credentials

Symantec recommends saving user credentials before manually removing the VRTSat package. If you need the credentials again, you can restore them to their original locations.

To save user credentials

1 Run the vssat showbackuplist command. The command displays the data files and backs them up into the SnapShot directory /var/VRTSatSnapShot. Output resembles the following:

```
vssat showbackuplist
B| /var/VRTSat/.VRTSat/profile/VRTSatlocal.conf
B| /var/VRTSat/.VRTSat/profile/certstore
B| /var/VRTSat/RBAuthSource
B| /var/VRTSat/ABAuthSource
B| /etc/vx/vss/VRTSat.conf
Quiescing ...
Snapshot Directory :/var/VRTSatSnapShot
```

2 Move the credentials to a safe location. Preserving the directory structure makes restoring the files easier.

To restore user credentials

1 Navigate to the SnapShot directory or the safe location where you previously saved credentials:

```
cd /var/VRTSatSnapShot/
```

2 Restore the files:

```
cp ABAuthSource /var/VRTSat/
cp RBAuthSource /var/VRTSat/
cp VRTSat.conf /etc/vx/vss/
cd /var/VRTSatSnapShot/
cp -rp profile /var/VRTSat/.VRTSat/
```

Bundled agent limitations

This section covers the software limitations for VCS 5.0 bundled agents.

NFS wizard limitation

The NFS wizard allows only one NFS service group to be created. You need to create additional groups manually.

Volume agent clean may forcibly stop volume resources

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

NFS failover

If the NFS share is exported to the world (*) and the NFS server fails over, NFS client displays the following error, “Permission denied”.

To avoid this error, export NFS shares explicitly using FQDN hostnames.

False concurrency violation when using PidFiles to monitor application resources

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

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

Networking agents do not support IPv6 protocol

The bundled IP, NIC, IPMultiNIC, MultiNICA, IPMultiNICB, and MultiNICB agents for do not support the IPv6 enhanced IP protocol.

VCS does not provide a bundled agent for volume sets

does not provide a bundled agent to detect Volume Manager volume sets, Problems with volumes and volume sets can only be detected at the DiskGroup and Mount resource levels.

Workaround: Set StartVolumes and StopVolumes attributes of the DiskGroup resource that contains volume set to 1. If a file system is created on the volume set, use a Mount resource to mount the volume set.

Cluster Management Console limitations

This section covers the software limitations for Cluster Management Console.

Cluster connector not supported on some OS versions

Cluster Management Console does not support cluster connector on AIX 5.1, Solaris 7, and RHEL 3.0. If your cluster runs on any of these platforms, you must use direct connection to manage the cluster from a management server.

Limited peer management server support

Peer management server support is limited to a configuration of two management servers in an enterprise. An enterprise of three or more management servers is not supported in this release.

Management server cannot coexist with GCM 3.5 Master

The Cluster Management Console management server should not be installed on the same system with a GCM 3.5 Master. These two products will conflict with each other and are not supported running on the same system.

Agent info files needed for Agent Inventory report

By design, the Agent Inventory report requires agent info files that supply the information reported on individual agents. These files are shipped with agents in VCS.

Global clusters must be CMC-managed clusters

All clusters forming a global cluster (using the VCS 4.0 Global Cluster Option) must be managed clusters in order for Veritas Cluster Management Console views to display correct and consistent information. Managed clusters are running the cluster connector or have a direct connection with the management server.

HP-UX cluster connector install fails if filesystem mount fails

If an HP-UX system has a mount that is not in /etc/fstab or /etc/checklist, the HP-UX installer swinstall will not work. Be sure the mount has an entry in these files.

Be sure the HP-UX system is configured to not attempt mounting of all the filesystems when performing an install or uninstall. This can be accomplished by adding the following lines to the /var/adm/sw/defaults file:

```
swinstall.mount_all_filesystems=false
swremove.mount_all_filesystems=false
```

Windows Active Directory installation requires NetBIOS

If you install Cluster Management Console management server in a Windows Active Directory domain, NetBIOS must be turned on. A native (non-NetBIOS) Active Directory environment is not supported in this release.

Remote root broker not supported on Windows

If you set up a management server on a Windows system, you must configure a root broker on the management server system. This release does not support specifying a remote root broker during management server install [841739].

The root broker can be changed after install using the configureRemoteRoot.exe installed in C:\Program Files\VERITAS\Cluster Management Console\bin (default install directory).

Cluster Manager (Java console) limitations

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

Use the VCS 5.0 Java Console to manage clusters

Cluster Manager (Java Console) from previous VCS versions cannot be used to manage VCS 5.0 clusters. Symantec recommends using the latest version of Cluster Manager.

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

Run Java Console on a non-cluster system

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

Cluster Manager and wizards do not work if the hosts file contains IPv6 entries

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

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

VCS Simulator does not support I/O fencing

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

Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 software limitations

There are no software limitations for this release.

The software limitations for the previous releases are as follows:

5.1 SP1 RP1	See "Veritas Storage Foundation Cluster File System 5.1 SP1 RP1 software limitations" on page 84.
5.1 SP1	See "Veritas Storage Foundation Cluster File System 5.1 SP1 software limitations" on page 84.

Veritas Storage Foundation Cluster File System 5.1 SP1 RP1 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 Cluster File System 5.1 SP1 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.

Limitations related to I/O fencing

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

Stopping systems in clusters with I/O fencing configured

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

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

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

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

Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 software limitations

There are no software limitations in this release.

The software limitations from previous releases are as follows:

5.1 SP1 RP1	There are no software limitations in this release.
5.1 SP1	See "Veritas Storage Foundation for Oracle RAC 5.1 SP1 software limitations" on page 86.

Veritas Storage Foundation for Oracle RAC 5.1 SP1 software limitations

This section describes the software limitations in SF Oracle RAC 5.1 SP1.

Stopping systems in clusters with I/O fencing configured

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

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

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

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

Limited global clustering capabilities in Veritas Cluster Server Management Console

The Veritas Cluster Server Management Console (VCS MC) management server currently includes limited capabilities for global clustering. Use the VCS graphical user interface (hagui) to take advantage of global clustering capabilities such as Fire Drill - Readiness and Remote Group Agent for remote groups.

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.

Health checks may fail on clusters that have more than 10 nodes

If there are more than 10 nodes in a cluster, the health check may fail with the following error:

```
vxgettext ERROR V-33-1000-10038
Arguments exceed the maximum limit of 10
```

The health check script uses the `vxgettext` command, which does not support more than 10 arguments.[2142234]

Cached ODM not supported in SF Oracle RAC environments

Cached ODM is not supported for files on Veritas local file systems and on Cluster File System.

Performance recommendation for space-optimized volume snapshots

For minimal performance impact, Symantec recommends that the Space Optimized Snapshots (SOS) be created only of data volumes. A mirror breakoff snapshot should be created of the Oracle log volume. The log volumes are typically small in size and do not have significant space overhead.

Unsupported volume location scenarios

The ocrvol volume and votevol volume cannot exist in the same shared disk group as that of the Oracle datafiles. However, you can allow for this scenario when you manually configure Oracle service groups.

Oracle Disk Manager (ODM) limitation

Oracle Disk Manager (ODM) uses the Quick I/O driver for asynchronous I/O. Do not turn off the Quick I/O mount option, which is the default.

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.

Replication in a shared environment

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

Storage Checkpoint and Database FlashSnap limitation

The following are the limitations of Storage Checkpoint and Database FlashSnap:

- You cannot create a clone database using a mounted Storage Checkpoint.
- If you create an Oracle instance using the spfile option, you must run the `dbed_update` command before you can successfully perform any Storage Checkpoint or Database FlashSnap functions.
- Storage Checkpoints require file system layout version 6 or version 7. Use the vxupgrade(1M) command to check the current layout version and to change the layout version, if necessary. When upgrading a CFS file system, issue the command from the primary node. Note that after you upgrade a system to layout version 6 or version 7, the file system is no longer compatible with the older VxFS file systems.
- When cloning a database using Database FlashSnap, the Oracle database must have at least one mandatory archive destination. For more information about Oracle parameters for archiving redo logs, see your Oracle documentation.
- Only online snapshots are supported for an Oracle RAC database, when using the `dbed_vmsnap`, `dbed_vmclonedb`, and `dbed_vmchecksnap` commands.
- After running `dbed_vmsnap -o reverse_resync_commit`, your primary database is started using a pfile. If your original primary database used an spfile, you need to shut down the database and restart it using spfile. Then, run `dbed_update` to update the repository.

- The Storage Checkpoint and Database FlashSnap features of SF Oracle RAC do not support the graphical user interface of the Veritas Storage Foundation for Oracle product.
- The Database FlashSnap feature does not support RAID-5 volumes.
- SF Oracle RAC does not support the Veritas FlashSnap agent for Symmetrix (EMC TimeFinder) mapping functionality (package:`VRTSfas`).

Known issues in this release

This section describes the known issues in this release.

- Issues related to installation:
  See "Issues related to installation 5.1 SP1 RP2" on page 89.
- Veritas Storage Foundation:
  See "Veritas Storage Foundation 5.1 SP1 RP2 known issues" on page 96.
- Veritas Storage Foundation for Databases (SFDB) Tools:
  See "Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP2 known issues" on page 121.
- Veritas Cluster Server:
  See "Veritas Cluster Server 5.1 SP1 RP2 known issues" on page 129.
- Veritas Storage Foundation Cluster File System:
  See "Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 known issues" on page 157.
- Veritas Storage Foundation for Oracle RAC:
  See "Veritas Storage Foundation for Oracle RAC known issues" on page 163.

Issues related to installation 5.1 SP1 RP2

The issues related to installation for previous releases are as follows:

5.1 SP1 RP1	See "Issues related to installation 5.1 SP1 RP1" on page 90.
5.1 SP1	See "Issues related to installation 5.1 SP1" on page 93.

This section lists issues related to installation in this release.

Warnings are getting generated in the logs when uninstalling patch PHKL_43127 and patch PHCO_43065 (2965619)

When using installer to uninstall the patches PHKL_43127 and PHCO_43065, you may see the following warnings in the output of swjob:

```
WARNING: The dependencies for fileset "PHKL_43127.VXFS-KRN,l=/,r=1.0"
         cannot be resolved (see previous lines).
         The operation on this fileset will still be attempted even
         though there are unresolved dependencies because the
         "enforce_dependencies" option is set to "false".

       * Summary of Analysis Phase:
WARNING:     Remove        PHKL_43127.VXFS-KRN,l=/,r=1.0
WARNING: 1 of 1 filesets had Warnings.
WARNING: The Analysis Phase had warnings.  See the above output for
         details.


WARNING: The dependencies for fileset "PHCO_43065.VXVM-RUN,l=/,r=1.0"
         cannot be resolved (see previous lines).
         The operation on this fileset will still be attempted even
         though there are unresolved dependencies because the
         "enforce_dependencies" option is set to "false".

       * Summary of Analysis Phase:
WARNING:     Remove        PHCO_43065.VXVM-RUN,l=/,r=1.0
WARNING: 1 of 3 filesets had Warnings.
       * 2 of 3 filesets had no Errors or Warnings.
WARNING: The Analysis Phase had warnings.  See the above output for
         details.
```

Workaround: The warnings can be safely ignored.

Issues related to installation 5.1 SP1 RP1

This section describes the known issues during installation and upgrade.

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

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

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

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.

Issues related to installation 5.1 SP1

This section lists known issues for 5.1 SP1 release.

Issues related to installation

This section describes the known issues during installation and upgrade.

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

The Web-based installer fails to run from DVD to install 5.1SP1.

Workarounds:

For this first 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.

If you have limited local disk space, use the second workaround.

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.

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

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

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

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

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

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

```
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dg1.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dg1 failed.
```

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

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

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

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

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

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

To resolve this issue

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

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

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

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

Installer assigns duplicate node ID during `-addnode` procedure

While performing an `-addnode` 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.

Veritas Storage Foundation 5.1 SP1 RP2 known issues

The known issues for Veritas Storage Foundation for previous releases are as follows:

5.1 SP1RP1	See "Veritas Storage Foundation 5.1 SP1 RP1 known issues" on page 98.
5.1 SP1	See "Veritas Storage Foundation 5.1 SP1 known issues" on page 110.

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.

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

Node join fails if the recovery for the leaving node is not completed.

Workaround: Node join should be retried after recovery is done. The reason being that when a node leaves a cluster, there is associated volume recovery for the leaving node.During this recovery phase, if any node tries to join the cluster, it is not allowed because we cannot assign the node id of the leaving node to the joining node. Hence the node join fails. The operation will be successful after the recovery is completed.

When making SANBOOT disks on the system by "vxcp_lvmroot" command , it failed with the error:VxVM vxcp_lvmroot (2920800)

If Operating System (OS) legacy device special files are not available, the VxVM operation of creating VXVM BOOT disk by running `vxdisksetup -B <device name>` may fail.

Workaround: There is no known workaround to fix the problem. You will have to upgrade to VxVM versions which contain fixes for this problem.

vxautoconvert(1M) fails to covert an LVM VG to a VxVM DG with new namingscheme. Convert fails and not able to roll back LVM VG's original state (2826905)

`vxautoconvert(1M)` fails to convert an LVM VG to a VxVM DG. The convert process aborts at a time when the recovery of LVM VG is not possible.

Workaround: Change to different naming scheme before convert procedure using vxautoconvert(1M) and once the conversion is finished, return back to the original(new namingscheme).

If vxconfigd is under heavy load, “vxassist settag” can make volume tagging information inconsistent (2484764)

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

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

Workaround: There is no workaround for this issue.

On HPIVM 6.1, VxVM can not identify "thin disks" (2942692)

On HPIVM 6.1, VxVM can not identify "thin disks". When "vxdisk -o thin list" command is run, the follwing error message is displayed:

```
"VxVM vxdisk INFO V-5-1-14413  No Thin Provisioned disk are attached to
the system."
```

Workaround: There is no workaround for this issue.

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

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

- Disabled site A storage from all the four nodes and also shutdown site A nodes.
- Enabled Site A storage and activated Site A nodes.
- Site B nodes panic.

After reboot(Site A nodes), when nodes try to join the cluster, vxconfigd dumps core.

Workaround: There is no workaround for this issue.

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 tunable`maxdsiz`. 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 unmanage`command. 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` or`vxautoconvert` 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 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.

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 1-36shows the Hitachi arrays that have new array names.

Table 1-36 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.

Veritas Storage Foundation 5.1 SP1 known issues

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

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.

Workarounds:

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

Veritas Volume Manager known issues

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

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.

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 tunable `maxdsiz`. 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.

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

Pending decision about whether to include.

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

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

Work-around:

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

Thin reclamation on disks with the hpdisk format is not supported (2136238)

Thin reclamation on disks with the hpdisk format is not supported. An attempt to perform reclamation on such disks automatically aborts.

Work-around:

There is no workaround for this issue.

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
```

Work-around:

There is no workaround for this issue.

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.

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.

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

The Veritas Volume Manager (VxVM) 5.1 SP1 RP2 includes several array names that differ from the array names in previous releases. Therefore, if you upgrade from a previous release to VxVM 5.1 SP1 RP2, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 5.1 SP1 RP2. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-37 shows the Hitachi arrays that have new array names.

Table 1-37 Hitachi arrays with new array names

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

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. The persistence of the enclosure name is achieved with the `/etc/vx/array.info` file, which stores the mapping between cabinet serial number and array name. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 5.1 SP1 RP2. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

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.

Work around:

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
```

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.

Issues with ALUA arrays that support standby Asymmetric Access State (AAS) when using EFI disks on HP 11.31 (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.

Veritas File System known issues

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

Storage Checkpoints can exceed the quota limit (2102201)

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

- In certain cases, if a FileSnap is truncated in the presence of a Storage Checkpoint, the `i_nblocks` field of the inode, which tracks the total number of blocks used by the file, can be miscalculated, resulting in inode being marked bad on the disk.
- In certain cases, when more than one FileSnap is truncated simultaneously in the presence of a Storage Checkpoint, the file system can end up in a deadlock state.

This issue causes the following error to display:

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

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

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

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

Workaround: There is no workaround for this issue.

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

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

```
# vxfilesnap source_file
     target_dir
```

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

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

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

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

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

Workaround: There is no workaround for this issue.

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.

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
```

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

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

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

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

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.

Workaround: There is no workaround for this issue.

vxassist relayout removes the DCM (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

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

3. Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7. Start the RVG. Enter the following:

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

8. Resume or start the applications.

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

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

Workaround:

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

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

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP2 known issues

The known issues for Veritas Storage Foundation for Databases (SFDB) tools for previous releases are as follows:

5.1 SP1 RP1	See "Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP1 known issues" on page 122.
5.1 SP1	See "Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 known issues." on page 126.

Veritas Storage Foundation for Databases (SFDB) tools 5.1 SP1 RP2 known issues

The following are the Veritas Foundation for Databases (SFDB) tools known issues for this release

Executing the /opt/VRTSdbed/bin/dbdst_obj_move command failed on 10gRAC env (2927308)

The dbdst_obj_move command fails with FSPPADM error:

```
/opt/VRTS/bin/dbdst_obj_move -S $ORACLE_SID -H $ORACLE_HOME \
-v -t tab_part4 -s 0 -e 10 -c SLOW
FSPPADM err : Not enough space
SFORA dst_obj_adm ERROR V-81-6414 Internal Error at fsppadm_err
```

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

Workaround: Upgrade the filesystem layout to version 8.

To upgrade the filesystem layout to version 8:

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

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

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

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

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 when`sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

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

Workaround:

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

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

This section lists the Veritas Storage Foundation for Databases (SFDB) tools known issues for 5.1 SP1 release.

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

While upgrading from 50mp2 to 51SP1 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:

   ```
   # sfua_rept_migrate
   ```

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

When upgrading from SF Oracle RAC version 5.0 or 5.0.1 to SF Oracle RAC 5.1SP1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the 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 an SF Oracle RAC environment, if the database fails over during Flashsnap operations such as the `dbed_vmsnap -o resync` command and various error messages appear. This issue occurs because Flashsnap commands do not create a VCS resource for the SNAP disk group. As such, when the database fails over, only the primary disk group is moved to another node.

Workaround

There is no workaround for this issue.

The error messages depend on the timing of the database failover. To fix the problem, you need to bring the FlashSnap state to SNAP_READY. Depending on

the failure, you may have to use base VxVM commands to reattach mirrors. After mirrors are attached, you need to wait until the mirrors are in SNAPDONE state. Re-validate the snapplan again.

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

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

Workaround

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

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

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

   ```
   # vxdg join snapshop_disk_group_name
         primary_disk_group_name
   ```

2. Start all the volumes in primary disk group.

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

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

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

   Repeat this step for all the volumes.

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

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

Workaround

There is no workaround for this issue.

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

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

Workaround: Set MonitorOption attribute for Oracle resource to 0.

Veritas Cluster Server 5.1 SP1 RP2 known issues

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

Failure messages of resource un-registration with IMF appear in agent or engine logs after performing online or offline operations on the resource (2909184)

When a resource is registered with IMF for monitoring, any online or offline operation triggers un-registration of the resource from IMF. During such operations, the agent may log an error message in the agent or engine logs stating that the un-registration failed. This issue is also observed for multiple resources.

Workaround:

The failure messages can be safely ignored. The agent re-registers the resources with IMF after sometime.

Issues with the amfstat output (2926158)

The `amfstat` output displays an extra column in the Registered Reapers list and the `amfstat -n` output displays the header twice.

Workaround:

This issue does not impact the functionality of Asynchronous Monitoring Framework (AMF). This issue has been fixed in VCS 6.0 and onwards.

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 (2556350)

This issue occurs if you have a resource dependency, such as r1 -> r2 -> r3. While resources r2 and r3 are online and you initiate bringing resource r1 online, before the OnlineTimeout occurs, resources r2 and r3 suffer a fault. Resource r2 faults first, and then r3 faults. After the fault of both resources is detected, the group

is in an OFFLINE|FAULTED state and resource r1 is stuck waiting to become online. If you execute the `hagrp -clear` command to clear the fault, then HAD dumps core on all nodes due to assertion failure.

Workaround:

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

Sometimes parent group will not restart with OnlineRetryLimit set (2279845)

With OnlineRetryLimit set, 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.

This scenario may happen typically when:

- You have single system in the SystemList of child and parent groups.
- The group has a mix of faulted persistent and non-persistent resources.

Workaround:

If the group with OnlineRetryLimit does not restart or failover, manually clear the fault and run the `online` command.

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 (2586230).

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.

Installer is unable to split a cluster that is registered with one or more CP servers (2110148)

Splitting a cluster that uses server-based fencing is currently not supported. You can split a cluster into two and reconfigure Storage Foundation Cluster File System HA on the two clusters using the installer.

For example, you can split a cluster clus1 into clus1A and clus1B. However, if you use the installer to reconfigure the Storage Foundation Cluster File System HA, the installer retains the same cluster UUID of clus1 in both clus1A and clus1B.

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

Workaround: None.

Veritas Cluster Server 5.1 SP1 known issues

This section covers the known issues in the VCS 5.1 SP1 release.

Issues related to installation

This section describes the known issues during installation and upgrade.

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

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

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

To manually upgrade the VRTSvlic package

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

   ```
   # vxkeyless display
   ```

2. Set the product level to NONE.

   ```
   # vxkeyless set NONE
   ```

3. Upgrade the VRTSvlic package.

   ```
   # swremove VRTSvlic
   ```

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

   ```
   swinstall -s 'pwd'
   ```

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

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

Installer assigns duplicate node ID during `-addnode` procedure

While performing an `-addnode` 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.

Issues with keyless licensing reminders after upgrading VRTSvlic (2141446)

After upgrading from 5.0.1 to 5.1SP1, some keyless licenses may be left in the system. As a result, you may see periodic reminders being logged if the VOM server is not configured.

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

To resolve this issue, perform the following steps:

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

   ```
   # vxkeyless set NONE
   ```

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

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

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

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

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

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

The installer crashes when you add a node using the –addnode option

If you manually remove a node from a VCS cluster or SFRAC cluster and then add the node back to the cluster by using the installer, a duplicate node ID is created. The installer terminates abruptly.

Resolution: Add the node manually.

Errors recorded in the swremove logs of VRTSgab during VCS upgrade from 4.1 to 5.0.1

When VCS is upgraded from 4.1 to 5.0.1 on HP-UX 11i v3 using the Veritas product installer, the installer reports errors for GAB and errors are recorded in the swremove logs related to VRTSgab. [1719136]

You can safely ignore these error messages.

VCS agents dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the update-ux command

On PA-RISC architecture, the VCS agents (Oracle, Netlsnr, Sybase, SybaseBk, MultiNICB, and so on) may dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the `update-ux` command.[1630968]

This is because on HP-UX PA-RISC systems, the default thread stack size is limited to 64k. When the agent requires more than 64k stack memory, it may dump core due to SIGBUS error.

Workaround: Before running the `update-ux` command, edit the /opt/VRTSvcs/bin/vcsenv file to append following lines to it:

```
PLATFORM=`uname -s`
ARCHITECTURE=`uname -m`
if [ "${PLATFORM}" = "HP-UX" ] && [ "${ARCHITECTURE}" = "9000/800" ]; then
        PTHREAD_DEFAULT_STACK_SIZE=524288
        export PTHREAD_DEFAULT_STACK_SIZE
fi
```

Installer cannot split a cluster registered with one or more CP servers [2110148]

Splitting a cluster that uses server-based fencing is currently not supported. You can split a cluster into two and configure VCS on the two clusters using the installer.

For example, you can split a cluster Clus1 into clus1A and clus1B. However, if you use the installer to reconfigure the VCS, the installer retains the same cluster UUID of Clus1 in clus1A and clus1B. If both clust1A and clus1B use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that

attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Installer enters a loop when cluster is running in secure mode while configuring server-based fencing [2166599]

During server-based fencing configuration with a secure cluster, if `vxfen` fails to start and you retry server-based fencing configuration, the installer keeps asking to enter another system to enable security after you manually start VCS.

Workaround: When `vxfen` fails to start in customized mode for server-based fencing with a secure cluster, do not choose to retry configuring fencing. Select the default option and vxfen starts in disabled mode. You can also retry fencing configuration using `-fencing` option.

Issues related to any OS or supported technology

NFS cluster I/O fails when storage is disabled

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

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

Operational issues for VCS

Volumes outside of VCS control that are mount locked cannot be unmounted without specifying the key

If a VxFS file system has "mntlock=key" in its mount options, then you cannot unmount the file system without specifying the key. Groups having DiskGroup resources configured with UmountVolumes set, may fail to switch or failover if the volumes are mount locked. [1276594]

Saving large configuration results in very large file size for main.cf

If your service groups have a large number resources or resource dependencies, and if the PrintTree attribute is set to 1, saving the configuration may cause the configuration file to become excessively large in size and may impact performance. [616818]

Workaround: Disable printing of resource trees in regenerated configuration files by setting the PrintTree attribute to 0.

AutoStart may violate limits and prerequisites load policy

The load failover policy of Service Group Workload Management may be violated during AutoStart when all of the following conditions are met: [251660]

- More than one autostart group uses the same Prerequisites.
- One group, G2, is already online on a node outside of VCS control, and the other group, G1, is offline when VCS is started on the node.
- The offline group is probed before the online group is probed.

In this scenario, VCS may choose the node where group G2 is online as the AutoStart node for group G1 even though the Prerequisites load policy for group G1 is not satisfied on that node.

Workaround: Persistently freeze all groups that share the same Prerequisites before using the `hastop -force` command to stop the cluster or node where any such group is online. This workaround is not required if the cluster or node is stopped without the force option.

Trigger not invoked in REMOTE_BUILD state

In some situations, VCS does not invoke the injeopardy trigger if the system is a REMOTE_BUILD state. VCS fires the trigger when the system goes to the RUNNING state.

The hagetcf script reports an error

Running the hagetcf script to gather information about the VCS cluster generates the following error:

```
tar: cannot stat ./var/VRTSvcs/log/.*A.log.  Not dumped.
```

Workaround: This message may be safely ignored.

Node cannot join cluster because port v is not ready for configuration

This behavior is observed when a node leaves a cluster and another node tries to join the cluster at the same time. If the GAB thread is stuck in another process, the new node cannot join the cluster and GAB logs the following warning:

```
GAB WARNING V-15-1-20126 Port v not ready
for reconfiguration, will retry.
```

Using the coordinator attribute

This release contains an attribute for disk groups called coordinator, which configures disks as coordinator disks by the I/O fencing driver. Setting the attribute prevents the coordinator disks from being reassigned to other disk groups. See

the Veritas Volume Manager documentation for additional information about the coordinator attribute.

The attribute requires that the disk group contain an odd number of disks. Symantec recommends that you use only three coordinator disks. Using more (five or seven) disks may result in different subclusters.

Some alert messages do not display correctly

The following alert messages do not display correctly [612268]:

51033	Global group %s is unable to failover within cluster %s and AutoFailOver is %s. Administrative action is required.
51032	Parallel global group %s faulted on system %s and is unable to failover within cluster %s. However, group is still online/partial on one or more systems in the cluster
51031	Unable to automatically fail over global group %s remotely because local cluster does not have Authority for the group.
51030	Unable to find a suitable remote failover target for global group %s. Administrative action is required
50916	Unable to automatically failover global group %s remotely due to inability to communicate with remote clusters. Please check WAN connection and state of wide area connector.
50914	Global group %s is unable to failover within cluster %s and ClusterFailOverPolicy is set to %s. Administrative action is required.
50913	Unable to automatically fail over global group %s remotely because clusters are disconnected and ClusterFailOverPolicy is set to %s. Administrative action is required.
50836	Remote cluster %s has faulted. Administrative action is required.
50761	Unable to automatically fail over global group %s remotely because ClusterList values for the group differ between the clusters. Administrative action is required.

Issues with configuration of resource values

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

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

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

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

Issues with bunker replay

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

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

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

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

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

Volume group can cause concurrency violation under VCS control

When you restart the system, it causes automatic activation of the LVM volume groups. The LVM volume can cause concurrency violation issue for VCS if they are under VCS control.

To avoid this issue, you must disable auto-activation of the volume groups. Set the `AUTO_VG_ACTIVATE` variable to 0 in `/etc/lvmrc` file, using the following command:

```
# cat /etc/lvmrc |grep AUTO_VG_ACTIVATE

# AUTO_VG_ACTIVATE and RESYNC which are required by the script in /sbin/lvm
# AUTO_VG_ACTIVATE flag to 0 and customizing the function
#    set AUTO_VG_ACTIVATE to 0.
AUTO_VG_ACTIVATE=0
```

Note: This routine is executed only if `AUTO_VG_ACTIVATE` is set to 1.

The `swverify` command displays a note

When you run `swverify` on HP-UX systems with VCS 5.1SP1 installation, the system displays the following note:

```
Note: Volatile file "/var/VRTSat/.VRTSat/Profile/vxatdlog.conf"
```

You can ignore this note as it does not affect the working of VCS. Do not regard it as an error or a warning.

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

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

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

```
::1         localhost
```

Saving large configuration results in very large file size for main.cf [616818]

If your service groups have a large number resources or resource dependencies, and if the PrintTree attribute is set to 1, saving the configuration may cause the configuration file to become excessively large in size and may impact performance.

Workaround: Disable printing of resource trees in regenerated configuration files by setting the PrintTree attribute to 0.

Issues related to the VCS engine

LinkHbStatus does not reflect the link status correctly

After disabling the LLT links of a node, the LinkHbStatus does not reflect the 'DOWN' flag for that node in 'hasys -disp'.[1831129]

Engine may hang in LEAVING state

When the `hares -online` command is issued for a parent resource when a child resource faults, and the `hares -online` command is followed by the `hastop -local` command on the same node, then the engine transitions to the leaving state and hangs.

Workaround: Issue the `hastop -local -force` command.

Timing issues with AutoStart policy

Consider a case where the service group is offline and engine is not running on node 1. If you restart the engine on node 1 after HAD is killed on node 2 and before the engine is restarted on node 2, then VCS does not initiate the autostart policy of the group.

On a default OEL4U4 install, VCS kernel components cannot start up

By default, OEL4U4 systems boot up in Xen-enabled kernels.

```
# uname -a

Linux host1 2.6.18-164.el5xen #1 SMP Thu March 4 04:41:04 EDT 2010
x86_64 x86_64 x86_64 GNU/Linux
```

However, VCS kernel modules are built only for the non-Xen kernels:

```
# cat kvers.lst

2.6.18-8.el5v

2.6.18-8.el5
```

Workaround: Set up your system for booting into the non-Xen kernels. For instructions, refer to the OS vendor's documentation.

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

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

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

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

Workaround

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

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

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

VCS fails to go to the running state on HP-UX 11.31 with March 2011 release

Due to a regression caused by the patch PHKL_41700 (QXCR1001078659) that went into HP-UX 11.31 March 2011 release, the select() call takes long time to return from 'timeout sleep'. Due to this, _had misses the heartbeat with GAB resulting in SIGABRT by GAB. [2287383]

Workaround: You must tune 'hires_timeout_enable' kernel parameter to 1 before starting the cluster. Run the following command to set this variable to 1:

```
# kctune hires_timeout_enable=1
```

Note: HP has delivered the resolution for this issue via PHKL_41967 patch post the March 2011 release.

Issues related to the bundled agents

Application agent cannot monitor kernel processes

Application agent cannot monitor processes which have wildcard characters that give a special meaning to `grep` command. [1232043]

RemoteGroup agent's monitor function may time out when remote system is down

If a RemoteGroup agent tries to connect to a system (specified as IpAddress) that is down, the monitor function of the RemoteGroup agent times out for the resource. [1397692]

LVMVolumeGroup resources do not depend on DiskReservation resources

An LVMVolumeGroup resource does not depend on a DiskReservation resource. [1179518]

Problem in failing over the IP resource

When a system panics, the IP address remains plumbed to the system for a while. In such a case, VCS may not succeed in failing over the IP resource to another system. This can be observed when a system panics during I/O Fencing.

Workaround: Increase the value of the OnlineRetryLimit attribute for the IP resource type.

LVMLogicalVolume agent may hang

The LVMLogicalVolume agent may hang in some situations, depending on the value of the IOTimeout attribute. Symantec recommends using the LVMCombo agent instead of the LVMLogicalVolume and LVMVolumeGroup agents.

LVM agents do not detect disconnected cable

LVM commands continue to function correctly when the cable to disks is pulled. The LVM agent does not detect a fault in this situation.

MultiNICB agent on fails with IPv6 protocol if no network is specified

Description: If you configure MultiNICB agent with IPv6 protocol without specifying a host or by only specifying non-reachable hosts in NetworkHosts, the agent keeps switching the active interface.

Workaround: You must specify at least one reachable host in the NetworkHosts attribute.

Could not write IPMultiNICB Options to file

If you have not specified the Options attribute in IPMultiNICB resource, the following message is logged:

```
Could not write IPMultiNICB Options to file.
```

However, there is no functionality loss. [2234686]

Workaround: Either specify the Options attribute or ignore the log massage.

MultiNICB resource goes to faulted state if you do not set the NetworkHosts attribute for IPv6 protocol

While using MultiNICB resource with interfaces configured with IPv6 protocol, the resource goes into faulted state if NetworkHosts attribute is not configured. [2132685]

Workaround: Set the NetworkHosts attribute for IPv6.

Issues related to global service groups

This section covers the issues related to global service groups.

Fault detection takes time in a global cluster running in secure mode

For global clusters running in secure mode, VCS may take a long time to detect a cluster fault on a remote cluster. [1403471]

Switch across clusters may cause concurrency violation

If you try to switch a global group across clusters while the group is in the process of switching across systems within the local cluster, then the group may go online on both the local and remote clusters. This issue affects only global groups. Local groups do not experience this behavior.

Workaround: Ensure that the group is not switching locally before attempting to switch the group remotely.

Global service group does not go online on AutoStart node

At cluster startup, if the last system where the global group is probed is not part of the group's AutoStartList, then the group does not AutoStart in the cluster. This issue affects only global groups. Local groups do not display this behavior.

Workaround: Ensure that the last system to join the cluster is a system in the group's AutoStartList.

Declare cluster dialog may not display highest priority cluster as failover target

When a global cluster fault occurs, the Declare Cluster dialog enables you to fail groups over to the local cluster. However, the local cluster may not be the cluster assigned highest priority in the cluster list.

Workaround: To bring a global group online on a remote cluster, do one of the following:

- From the Java Console, right-click the global group in the Cluster Explorer tree or Service Group View, and use the Remote Online operation to bring the group online on a remote cluster.
- From the Web Console, use the Operations links available on the Service Groups page to bring the global group online on a remote cluster.

Issues related to the VCS database agents

Issues related to the VCS Agent for DB2

This section covers issues related to the VCS agent for DB2.

awk error message

On IA-64, the default awk command may produce this error: Input line /usr/bin:/bin:/usr/s cannot be longer than 3,000 bytes. The source line number is 1.

Workaround: Install GNU awk.

Issues related to the VCS Agent for Oracle

This section covers the issues related to the VCS agent for Oracle.

NOFAILOVER action specified for certain Oracle errors

The Veritas High Availability agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file oraerror.dat, which consists of a list of Oracle errors and the actions to be taken.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for a description of the actions.

Currently, the reference file specifies the NOFAILOVER action when the following Oracle errors are encountered:

```
ORA-00061, ORA-02726, ORA-6108, ORA-06114
```

The NOFAILOVER action means that the agent sets the resource's state to OFFLINE and freezes the service group. You may stop the agent, edit the oraerror.dat file, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

Health check may not work for Oracle 10g R1 and 10g R2

For Oracle 10g R1 and 10g R2, if you set MonitorOption to 1, health check monitoring may not function when the following message is displayed [589934]:

```
Warning message - Output after executing Oracle Health
Check is: GIM-00105: Shared memory region is corrupted.
```

Workaround: Set MonitorOption to 0 to continue monitoring the resource.

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

The Oracle agent with 11g R1 and 11g R2 does not support Health check monitoring using the MonitorOption attribute. If the database is 11g R1 or 11g R2, the MonitorOption attribute for Oracle resource should be set to 0.

Intentional Offline feature is not supported for Oracle 11g R1 and 11g R2

The Oracle agent with 11g R1 and 11g R2 database does not support the Intentional Offline feature.

Pfile or SPfile is not supported on ASM diskgroups

The ASMInst agent does not support pfile or spfile for ASM Instance on ASM diskgroups in 11g R2. Symantec recommends you to store the file on the local file system.

ASM instance does not unmount VxVM volumes after ASMDG resource is offline

In configurations where ASMInstance resource is part of a separate parallel service group, the ASM instance does not unmount the volumes even after the ASMDG resource is taken offline. Therefore, the Volume resource cannot be taken offline. This issue occurs when you use VxVM volumes as ASM disk groups. [918022]

Workaround: Configure the ASMInstance resource as part of the failover service group where ASMDG resource is configured.

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.

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

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

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

VCS agent for Oracle: Intentional Offline does not work

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

Issues related to the Cluster Manager (Java Console)

This section covers the issues related to the Cluster Manager (Java Console).

32-bit JRE requirement

This release requires the installation of the 32-bit JRE ibm-java-ppc-jre-6.0-6.0.ppc. (1870929)

Cluster Manager (Java Console) may display an error while loading templates

You can access the Template View in the Cluster Manager from the Tools > Templates menu. If you have Storage Foundation configured in a VCS cluster setup, the following error may occur while the Cluster Manager loads the templates. (1433844)

```
VCS ERROR V-16-10-65 Could not load :-
/etc/VRTSvcs/Templates/DB2udbGroup.tf
```

Workaround: Ignore the error.

Exception when selecting preferences

On Windows systems, selecting the Java (Metal) look and feel of the Java Console may cause a Java exception. [585532]

Workaround: After customizing the look and feel, close restart the Java Console.

Java Console errors in a localized environment

When connected to cluster systems using locales other than English, the Java Console does not allow importing resource types or loading templates from localized directories. [585532]

Workaround: The workaround is to copy the types files or templates to directories with english names and then perform the operation.

Printing to file from the VCS Java Console throws exception

VCS Java Console and Help throw an exception while printing to a file from a system that does not have a printer configured. Also, the content is not written to the file.

Workaround: Before printing, make sure at least one printer is configured on the system where the VCS Java Console is launched.

Common system names in a global cluster setup

If both local and remote systems have a common system name in a global cluster setup, group operations cannot be performed on those systems using the Java console.

Workaround: Use command-line interface to perform group operations.

Some Cluster Manager features fail to work in a firewall setup

In certain environments with firewall configurations between the Cluster Manager and the VCS cluster, the Cluster Manager fails with the following error message [1392406]:

```
V-16-10-13 Could not create CmdClient. Command Server
may not be running on this system.
```

Workaround: You must open port 14150 on all the cluster nodes.

Issues related to VCS Simulator

This section covers the issues related to VCS Simulator.

Simulator clusters with Windows configurations fail to start on UNIX host platforms

The following clusters are affected: Win_Exch_2K3_primary, Win_Exch_2K3_secondary, Win_Exch_2K7_primary, Win_Exch_2K7_secondary, WIN_NTAP_EXCH_CL1, WIN_NTAP_EXCH_CL2, Win_SQL_Exch_SiteA, Win_SQL_Exch_SiteB, WIN_SQL_VVR_C1, WIN_SQL_VVR_C2. [1363167]

Workaround: For each of these clusters, there is a separate directory named after the cluster under the VCS Simulator installation directory

C:\Program Files\VERITAS\VCS Simulator on Windows

/opt/VRTScssim on Unix

Perform the following steps:

- Navigate to the conf/config directory under this cluster specific directory.
- Open the types.cf file in an editor and change all instances of the string "i18nstr" to "str".
- Open the SFWTypes.cf file in an editor if it exists in this directory and change all instances of the string "i18nstr" to "str".
- Repeat these steps for the following files if they exist: MSSearchTypes.cf, SQLServer2000Types.cf, ExchTypes.cf, SRDFTypes.cf.

VCS Simulator does not start on Windows systems

On Windows systems, starting VCS Simulator displays an error that the required MSVCR70.DLL is not found on the system. [859388]

Workaround: Run the following command:

```
set PATH=%PATH%;%VCS_SIMULATOR_HOME%\bin;
```

Or append %VCS_SIMULATOR_HOME%\bin; to PATH environment variable.

Error in LVMVolumeNFSGroup template for AIX

In the VCS Simulator, the AIX_NFS cluster gives error while loading the LVMVolumeGroupNFS template. [1363967]

This problem can also affect real AIX clusters if they try to load this template.

Workaround: For the Simulator, navigate to the Templates/aix directory under the VCS Simulator installation directory (C:\Program Files\VERITAS\VCS Simulator on Windows, /opt/VRTScssim on Unix). Open the

LVMVolumeNFSGroup.tf file and look for all instances of the MajorNumber = "". Remove the empty double-quotes and set the correct integer value for MajorNumber.

For real clusters, make identical changes to /etc/VRTSvcs/Templates/LVMVolumeNFSGroup.tf.

VCS 5.0.1 Rolling Patch 1 known issues

The VCS issues in this release are as follows:

- The ASMInst agent does not support pfile or spfile for the ASM Instance on the ASM diskgroups in 11g Release 2. Symantec recommends that you store the file on the local file system. [1975010]
- The VRTSperl patch takes more than 10 minutes to install on an HP Integrity system node:
  On an HP Integrity system node, installing the VRTSperl patch takes more than 10 minutes and requires that VCS is offline during this period. The installation time may vary based on the configuration of the machine on which the VRTSperl patch is being installed.

Issues related to AMF driver

AMF driver fails to unload with the Mount Agent running

If Mount Agent uses IMF to monitor mounts of type VxFS, then you cannot unload AMF driver as long as Mount Agent is running. [2262747]

Workaround: Stop the mount agent before you unload the AMF driver.

Startup or shutdown failure messages reported for LLT, GAB, VXFEN, and VCSMM

If you need to reboot the system when you install SF Oracle RAC, the init scripts for LLT, GAB, VXFEN, and VCSMM report start or stop failure messages. This is because SF Oracle RAC is not yet configured and the required configuration files are not yet generated for these components. These messages may be ignored. [1666327]

Issues related to I/O fencing

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

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

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

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

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

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

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

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

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

See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for more information on preferred fencing.

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

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

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

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

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

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

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

CP server cannot bind to multiple IPs (2085941)

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

Resolution: No known resolution for this issue.

Issues related to Symantec Product Authetication Service with

This section covers the issues related to Symantec Product Authentication Service with VCS.

The atldapconf command fails if the user in the Active Directory does not belong to any group

While using the atldapconf command, the user group must be specified. [1596332]

Output of addldapdomain returns error

The output of addldapdomain returns an error and the help contains incorrect information [1589886]

The vcsat and cpsat commands may appear to be hung

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

- /opt/VRTScps/bin/cpsat

- /opt/VRTSvcs/bin/vcsat

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

Workaround:

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

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

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

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

Verification for VRTSat package or patch returns errors

If you run swverify command on VRTSat package or patch, the command returns errors for missing files on VRTSat.CLIENT-PA32. [1244204]

Workaround: This message may be safely ignored.

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

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

Workaround:

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

Errors observed during partial upgrade of SFHA

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

Workaround

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

Veritas Storage Foundation Cluster File System 5.1 SP1 RP2 known issues

There are no known issues for Veritas Storage Foundation Cluster File System in this release.

The known issues for Veritas Storage Foundation Cluster File System for previous releases are as follows:

5.1 SP1 RP1	See "Veritas Storage Foundation Cluster File System 5.1 SP1 RP1 known issues" on page 157.
5.1 SP1	See "Veritas Storage Foundation Cluster File System 5.1 SP1 known issues" on page 159.

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 Cluster File System 5.1 SP1 known issues

This section lists the Veritas Storage Foundation Cluster File System known issues for 5.1 SP1 release.

Veritas Storage Foundation Cluster File System known issues

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

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.

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

If you run swverify command on VRTSat package or patch, the command returns errors for missing files on VRTSat.CLIENT-PA32. [1244204]

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

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

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

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

Issues related to I/O fencing

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

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

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

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

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

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

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

- The fencing setup uses customized mode with one or more CP servers.
- The application cluster has more than eight nodes.

- The node weight for a single node (say galaxy with node id 0) is more than the sum total of node weights for the rest of the nodes.
- A network fault occurs and the cluster partitions into two with the single node (galaxy) on one part and the rest of the nodes on the other part.

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

See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for more information on preferred fencing.

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

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

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

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

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

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

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

CP server cannot bind to multiple IPs (2085941)

Coordination point server (CP server) binds only to a single virtual IP and listens on the same. Application clusters cannot access the CP server if it fails to establish connection to this virtual IP. Therefore, if the connection fails because of the subnet in which the virtual IP of the CP server exists, you cannot access the CP

server even if there is another subnet through which the client can connect to the CP server over a different IP.

Resolution: No known resolution for this issue.

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

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

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

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

Workaround: None.

Veritas Storage Foundation for Oracle RAC known issues

The known issues in SF Oracle RAC are as follows:

5.1 SP1 RP2	See "Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 known issues" on page 163.
5.1 SP1 RP1	There are no known issues in this release.
5.1 SP1	See "Veritas Storage Foundation for Oracle RAC 5.1 SP1 known issues" on page 164.

Veritas Storage Foundation for Oracle RAC 5.1 SP1 RP2 known issues

This section describes the known issues in SF Oracle RAC 5.1 SP1 RP2.

Oracle Grid Infrastructure processes fail in 11.2.0.1 (2937578)

The Oracle Grid Infrastructure process `crsd.bin` fails with the following symptoms:

```
# $GRID_HOME/log/<node>/alert<node>.log
2010-12-05 09:58:48.176 [/ocw/grid/bin/orarootagent.bin(1372166)]
```

```
CRS-5822:Agent '/ocw/grid/bin/orarootagent_root' disconnected from server.
Details at (:CRSAG F00117:) in /ocw/grid/log/racnode2/agent/crsd/\
orarootagent_root/orarootagent_root.log.
2010-12-05 09:58:48.361 [ohasd(290984)]CRS-2765:
Resource 'ora.crsd' has failed on server 'node2'.
2010-12-05 09:58:51.273 [crsd(1581252)]CRS-1012:
The OCR service started on node node2.
```

Workaround: Apply the Oracle patch 11814167.

For more information, see the Oracle Metalink document: 1326008.1

Veritas Storage Foundation for Oracle RAC 5.1 SP1 known issues

The known issues in 5.1SP1 release are as follows:

Oracle RAC issues	See "Oracle RAC issues" on page 164.
SF Oracle RAC issues	See "SF Oracle RAC issues" on page 166.
SFDB tools issues	See "Veritas Storage Foundation for Databases (SFDB) tools known issues" on page 171.

Oracle RAC issues

This section lists the known issues in Oracle RAC.

Oracle Grid Infrastructure installation may fail with the SF Oracle RAC installer

When you run the `installsfrac -configure` command to install Oracle Grid Infrastructure for Oracle RAC 11g Release 2, the installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

Workaround: Export the `OUI_ARGS` environment variable, before you run the SF Oracle RAC installation program:

```
export OUI_ARGS=-ignoreInternalDriverError
```

For more information, see the Oracle Metalink document: 970166.1

During installation or system startup, Oracle Grid Infrastructure may fail to start

After successful installation of Oracle RAC 11g Release 2 Grid Infrastructure, while executing the `root.sh` script, `ohasd` may fail to start. Similarly, during system startup, Oracle Grid Infrastructure may fail to start though the VCS engine logs may indicate that the cssd resource started Oracle Grid Infrastructure successfully.

The following message may be displayed on running the `strace` command:

```
# /usr/bin/strace -ftt -p pid_of_ohasd.bin
14:05:33.527288 open("/var/tmp/.oracle/npohasd", O_WRONLY <unfinished ...>
```

For possible causes and workarounds, see the Oracle Metalink document: 1069182.1

Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software

The Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software.If the failure indicates that the OCR and vote device locations are not shared, ignore the message.

Oracle VIP Configuration Assistant fails with an error message

During Oracle RAC 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

```
The given interface(s), "" is not public.
Public interfaces should be used to configure virtual IPs.
```

This message appears only when the VIP is not from the regular public IP range (for example, 200.). [1182220]

Workaround: Invoke the `vipca` utility manually as the superuser.

```
# export DISPLAY=nebula:0.0
# $CRS_HOME/bin/vipca
```

Oracle Cluster Verification utility displays a warning message

During the final stage of Oracle RAC 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification utility.

For example:

```
Utility
============================================================
OUI-25031: Some of the configuration assistants failed. It is
strongly recommended that you retry the configuration
```

```
assistants at this time. Not successfully running any "
Recommended" assistants means your system will not be correctly
configured.
1. Check the Details panel on the Configuration Assistant Screen
to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button
to retry them.
============================================================
```

Workaround: You may safely ignore this message if the cluster is operating satisfactorily.

SF Oracle RAC issues

This section lists the known issues in SF Oracle RAC for this release.

Miscalculated file set usage (2123429)

When file set quotas are enabled, it may be possible for VxFS to get into a state where it thinks a very large number of blocks are allocated to checkpoints. This issue can be seen using the `fsckptadm` command:

```
# fsckptadm getquotalimit /mnt1
Filesystem   hardlimit   softlimit   usage   action_flag
/mnt1        10000       10000       18446744073709551614
```

This could cause writes to checkpoints to fail. It could also trigger the removal of removable checkpoints.

Workaround

If this occurs, disabling and re-enabling file set quotas causes VxFS to recalculate the number of blocks used by checkpoints:

```
# fsckptadm quotaoff /mnt1
# fsckptadm quotaon /mnt1
# fsckptadm getquotalimit /mnt1
Filesystem   hardlimit   softlimit   usage   action_flag
/mnt1        10000       10000       99
```

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

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

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

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

Workaround: None.

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.

Deporting issues with shared disk groups

If you manually deport a shared disk group, the CVMVolDg agent does not automatically reimport it as a shared disk group. You must manually reimport it as a shared disk group.

Stopping cluster nodes configured with I/O fencing

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect or “split brain.”

For more information, see *Veritas Cluster Server User’s Guide*.

I/O fencing uses SCSI-3 Persistent Reservation keys to implement data protection. The software places keys on I/O fencing coordinator and data disks. The administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator disks and data disks to prevent possible difficulties with subsequent cluster startup. Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordinator and data disks. Depending on the order of reboot and subsequent startup events, the cluster might warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown command instead of the reboot command to perform a graceful reboot for systems.

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

Stopping VCS does not unregister port f from GAB membership

In an SF Oracle RAC cluster with all the CFS resources under VCS control, when you stop VCS, all the CFS resources must go down cleanly and CFS must unregister port f from GAB membership. Oracle RAC 10g Clusterware does not clean up all its processes when it is stopped. Now, when you stop VCS, all the CFS resources go down. However, due to the left over Oracle processes, CFS does not unregister port f from GAB membership.

Workaround: Perform the following steps to bring down port f.

To bring down port f

1 Kill all the Oracle processes.

```
# kill -9 'ps -u oracle|awk '{print $1}'`
```

2 Verify that all CFS file systems are unmounted.

```
# mount | grep cluster
```

3 Unregister port f from GAB membership.

```
# fsclustadm cfsdeinit
```

DBED features are not integrated with GCO

DBED features are not integrated with Global Cluster Option (GCO). After GCO migration, be aware that DBED features will not be functional. [1241070]

Issue with format of the last 8-bit number in private IP addresses

The PrivNIC/MultiPrivNIC resources fault if the private IP addresses have a leading 0 in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1. or X.0X.X.1 or 0X.X.X.1, where X is an octet of the IP address. [1164506]

When you configure private IP addresses for Oracle Clusterware, ensure that the IP addresses have a format as displayed in the following two-node example:

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

Confirm the correct format by viewing the PrivNIC or MultiPrivNIC resource in the `/etc/VRTSvcs/conf/config/main.cf` file.

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

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

Work-around:

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

When master node loses access to complete storage, detached sites remain in RECOVER state even after reattaching and recovering the sites

In a campus cluster environment, if the master node loses access to complete storage, all but one of the sites is detached and the DCO volumes may get detached if the `dgfailpolicy` is set to `dgdisable`. If the detached sites are reattached and recovered, the site still remains in RECOVER state. [1828142]

Workaround: Change the status of the site as described in the following procedure to resolve the issue.

To change the status of the site

1. Log onto the CVM master node.
2. Reattach the detached sites:

   ```
   # vxdg -g dg_name reattachsite site_name
   ```

   The site remains in RECOVER state.

3 Restore DCO volumes by unpreparing and preparing the volumes.

Unprepare the volumes:

```
# vxsnap -g dg_name -f unprepare vol_name
```

Prepare the volumes:

```
# vxsnap -g dg_name prepare vol_name drl=on
```

4 Reattach the detached sites:

```
# vxdg -g dg_name reattachsite site_name
```

5 Verify that the state of the detached sites is now ACTIVE:

```
# vxprint
```

Issues related to LLT

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

LLT port stats sometimes shows recvcnt larger than recvbytes

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

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

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

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

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

packages are removed. This releases disk space. The installer then claims more space than it actually needs.

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

The VRTSacclib package is deprecated (2032052)

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

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

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

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

While upgrading from 50mp2 to 51SP1 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.
```

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

When upgrading from SF Oracle RAC version 5.0 or 5.0.1 to SF Oracle RAC 5.1SP1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3

startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the 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 an SF Oracle RAC environment, if the database fails over during Flashsnap operations such as the `dbed_vmsnap -o resync` command and various error messages appear. This issue occurs because Flashsnap commands do not create a VCS resource for the SNAP disk group. As such, when the database fails over, only the primary disk group is moved to another node.

Workaround

There is no workaround for this issue.

The error messages depend on the timing of the database failover. To fix the problem, you need to bring the FlashSnap state to SNAP_READY. Depending on

the failure, you may have to use base VxVM commands to reattach mirrors. After mirrors are attached, you need to wait until the mirrors are in SNAPDONE state. Re-validate the snapplan again.

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

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

Workaround

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

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

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

   ```
   # vxdg join snapshop_disk_group_name
         primary_disk_group_name
   ```

2. Start all the volumes in primary disk group.

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

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

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

   Repeat this step for all the volumes.

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

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

Workaround

There is no workaround for this issue.

Downloading the patches

The patches included in Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 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 RP2 download archive and instructions, visit:

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

Chapter 2

Upgrading to version 5.1 SP1 RP2

This chapter includes the following topics:

About the installrp script

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 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
`[ -ignorepatchreqs ]`	The `-ignorepatchreqs` option is used to allow installation or upgrading even if the pre-requisite packages or patches are missed on the system.
`[ -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.

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

Command Line Option	Function
[-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.
[-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.

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

Command Line Option	Function
`[ -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.
`[ -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 RP2 on a cluster

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

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 RP2:

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

Performing a full upgrade to 5.1 SP1 RP2 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 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

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

4 Resolve any issues that the precheck finds.

5 Start the upgrade:

```
# ./installrp 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 RP2 on an SFHA cluster

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

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

1 Log in as superuser.

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

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

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

    The program proceeds in a noninteractive mode to examine the systems for licenses, filesets, disk space, system-to-system communications, etc.
11. Review the output as the program displays the results of the check and saves the results of the check in a log file.

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

13 Start the upgrade.

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

Review the output.

14 Restart the nodes:

```
# shutdown -r now
```

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

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

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

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

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

```
# mount /filesystem
```

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

```
# mount /checkpoint_name
```

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

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

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

1 Log in as superuser.

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

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 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
```

5 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
```

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 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
```

9 If required, apply the OS kernel patches.

10 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
```

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

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

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

12 Restart the nodes:

```
# shutdown -r now
```

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

14 Bring the CVM service group online on each node:

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

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

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

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

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

17 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

18 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

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

Perform the steps in the following procedure to upgrade to version 5.1 SP1 RP2.

Note: If you are upgrading from SF Oracle RAC versions 3.5, 3.5 Update 3, and 3.5 Update 4, you need to first upgrade to version 4.1, then upgrade to 5.1 SP1, and finally upgrade to 5.1 SP1 RP2. For instructions on upgrading to versions 4.1 and 5.1 SP1, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for the corresponding version.

Contact Technical Support to obtain the software media or to download the software for version 4.1.

To perform a full upgrade to version 5.1 SP1 RP2 on an SF Oracle RAC cluster

1 Upgrade the operating system, if required.

See "Upgrading the HP-UX operating system" on page 187.

For instructions, see the operating system documentation.

2 Upgrade to version 5.1 SP1 RP2.

See "Upgrading SF Oracle RAC using the Veritas script-based installation program" on page 187.

Alternatively, use the Web-based installation program to upgrade SF Oracle RAC.

```
# ./webinstaller start
```

Follow the installation prompts to upgrade SF Oracle RAC.

3 Relink the SF Oracle RAC libraries with Oracle.

See "Relinking Oracle RAC libraries with the SF Oracle RAC libraries" on page 189.

4 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

  ```
  # hagrp -online oracle_group -sys node_name
  ```

- If the Oracle database is not managed by VCS:

  ```
  # srvctl start database -d db_name
  ```

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

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

7 - If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

  ```
  # haconf -makerw
  # hagrp -modify oracle_group AutoStart 1
  # haconf -dump -makero
  ```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

  ```
  $ srvctl modify database -d db_name -y AUTOMATIC
  ```

8 Complete other post-upgrade steps.

  For instructions, see the chapter *Performing post-upgrade tasks* in this document.

9 For upgrade scenarios that involve Oracle RAC 9i, start `gsd` as the Oracle user:

  ```
  $ $ORACLE_HOME/bin/gsdctl start
  ```

10 Upgrade Oracle RAC.

11 If you want to upgrade CP server systems that use VCS or SFHA to version 5.1 SP1 RP2, make sure that you upgraded all application clusters to version 5.1 SP1 RP2. Then, upgrade VCS or SFHA on the CP server systems.

  For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Upgrading the HP-UX operating system

If you are on an unsupported version of the operating system, you need to upgrade it to HP-UX B.11.31.1009, HP-UX 11i Version 3 September 2010 Operating Environments Update Release or later.

Note: If you are upgrading from SF Oracle RAC 5.0 or 5.0 RP1 on HP-UX 11i v3, you need to stop VCS before you upgrade the operating system. To stop VCS on all nodes, run the following command as the superuser:

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

If you are upgrading the operating system from HP-UX 11i v2, make sure that you choose the following depots along with the HP-UX 11i v3 September 2010 OEUR release depots:

- `Base-VxFS-50`
- `Base-VxTools-50`
- `Base-VxVM-50`

To upgrade the operating system from HP-UX 11i v2, run the `update-ux` command specifying the Veritas depots along with the HP-UX operating system depots:

```
# swinstall -s  os_path Update-UX
# update-ux -s os_path HPUX11i-DC-OE \
Base-VxFS-50 Base-VxTools-50 Base-VxVM-50
```

where `os_path` is the full path of the directory containing the operating system depots.

To upgrade the operating system from HP-UX 11i v3, run the `update-ux` command as follows:

```
# update-ux -s os_path HPUX11i-DC-OE
```

where `os_path` is the full path of the directory containing the operating system depots.

For detailed instructions on upgrading the operating system, see the operating system documentation.

Upgrading SF Oracle RAC using the Veritas script-based installation program

The product installer performs the following tasks to upgrade SF Oracle RAC:

- Verifies the compatibility of the systems before the upgrade.
- Stops the SF Oracle RAC processes before the upgrade.
- Uninstalls SF Oracle RAC.
- Installs the SF Oracle RAC packages on the nodes.
- Starts SF Oracle RAC on all the nodes.
- Displays the location of the log files, summary file, and response file.

To upgrade to version 5.1 SP1 RP2 using the script-based program

1. Log in as the superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
3. Stop VCS.

   ```
   # hastop -all
   ```

4. Apply any operating system patches, if required.
5. Check whether or not any Storage Checkpoints or VxFS file systems that are not managed by VCS are mounted:

   ```
   # mount | grep vxfs
   ```

6. Unmount all Storage Checkpoints and file systems that are not managed by VCS:

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

7. 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.
8. If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
   - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
   - Stop each RVG individually:

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

   - On the primary node, verify that all RLINKs are up-to-date:

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

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

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

10 Stop all VxVM volumes for each disk group:

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

Verify that no volumes remain open:

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

11 Check if the VEA service is running:

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

If the VEA service is running, stop the service:

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

12 Change to the directory that contains the `installrp`script:

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

where node1 and node2 are nodes to be upgraded.

13 Restart each node in the cluster.

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

14 Complete the remaining tasks to finish the upgrade:

See "Performing a full upgrade to version 5.1 SP1 RP2 on an SF Oracle RAC cluster" on page 185.

Relinking Oracle RAC libraries with the SF Oracle RAC libraries

You must relink the Oracle RAC libraries with the SF Oracle RAC libraries after upgrading SF Oracle RAC.

The steps vary depending on the version of Oracle RAC in use before the upgrade:

- If you upgraded from an SF Oracle RAC version running Oracle RAC 10g Release 2 or Oracle RAC 11g:
  See "To relink Oracle RAC 10g Release 2 or Oracle RAC 11g using the installer" on page 190.
- If you upgraded from an SF Oracle RAC version running Oracle RAC 9i or Oracle RAC 10g Release 1 or if you want to relink the libraries manually for later versions of Oracle RAC:
  See "To relink the Oracle RAC libraries with SF Oracle RAC libraries manually" on page 191.

Note: You must upgrade the database to a supported version after you complete the post-upgrade tasks.

To relink Oracle RAC 10g Release 2 or Oracle RAC 11g using the installer

1 On each node, shut down the Oracle service group:

```
# hagrp -offline oracle_group -sys system_name
```

2 Run the `installsfrac` installer:

```
# cd /opt/VRTS/install
# ./installsfrac -configure sys1 sys2
```

3 Enter **5** to select the option **Post Oracle Installation Tasks**.

```
1)  Configure SF Oracle RAC sub-components
2)  SF Oracle RAC Installation and Configuration Checks
3)  Prepare to Install Oracle
4)  Install Oracle Clusterware/Grid Infrastructure and Database
5)  Post Oracle Installation Tasks
6)  Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 5
```

4 Select the option **Relink Oracle Database Binary**.

```
1)  Configure CSSD agent
2)  Relink Oracle Database Binary
3)  Exit SF Oracle RAC Configuration
b)  Back to previous menu
Choose option: [1-3,b,q] (1) 2
```

5 Provide the following information:

```
Enter Oracle UNIX user name: [b] (oracle)
Enter Oracle UNIX group name: [b] (oinstall)
Enter absolute path of Oracle Clusterware/Grid Infrastructure
Home directory: [b]
Enter absolute path of Oracle Database Home directory: [b]
```

The installer detects the Oracle version.

6 Enter **y** to proceed with relinking.

```
Do you want to continue? [y,n,q] (y)
```

Perform the steps in the following procedure if you upgraded nodes running Oracle RAC 9i or Oracle RAC 10g Release 1 or if you want to relink the libraries manually.

To relink the Oracle RAC libraries with SF Oracle RAC libraries manually

1 On each node, shut down the Oracle service group:

```
# hagrp -offline oracle_grp -sys system_name
```

2 On each node, run the following command as the oracle user to link Oracle with the Veritas IPC, VCSMM, and ODM libraries:

```
$ /opt/VRTSvcs/rac/bin/linkrac oracle_version
```

Replace oracle_version with one of the following values depending on the version you installed:

- 9i (for Oracle RAC 9i)
- 10gR1 (for Oracle RAC 10g Release 1)
- 10gR2 (for Oracle RAC 10g Release 2)
- 11gR1 (for Oracle RAC 11g Release 1)
- 11gR2 (for Oracle RAC 11g Release 2)

Note: If the Oracle binaries are on a cluster file system, perform this step on only one node. If the Oracle binaries are on a local file system of each cluster node, run the command to link Oracle with Vertitas libraries on each cluster node.

If your system uses a bundled C compiler or a compiler other than the ANSI C compiler, you can safely ignore such warnings as:

```
(Bundled) cc: warning 922: "+Oshortdata=8" is unsupported in the
bundled compiler, ignored.
```

Refer to Oracle Metalink Document ID 66442.1 for more information.

Search on: 66442.1 FAQ about C Compiler Issues on HP-UX.

3 On each node, start the Oracle service group:

```
# hagrp -online oracle_grp -sys system_name
```

4 After starting the Oracle instance, confirm Oracle uses the Veritas libraries. Examine the Oracle alert file, `alert_$ORACLE_SID.log`, for the following lines:

```
Oracle instance running with ODM: Veritas 5.1 ODM Library,
Version 2.0
```

Additionally, for Oracle RAC 10g, verify that the cluster interconnect IPC version is `VERITAS IPC '5.0.31.0'`.

Note: If Oracle binaries are on a local file system of each cluster node, examine the Oracle alert file on each of the cluster nodes.

If you see the following message in the `alert_$ORACLE_SID.log` file

```
cluster interconnect IPC version string is not available
Oracle interface version information 2.4
cluster IPC library version information 2.2
```

Perform the following steps:

- Stop the database on all the nodes (for instances under VCS control):

  ```
  # hagrp -offline oracle_grp -sys system_name
  ```

  Stop the database on all the nodes (for instances not under VCS control):

  ```
  # srvctl stop database -d database_name
  ```

- For Oracle RAC 9i/Oracle RAC 10g Release 1: Export the IPC_LIB_PATH variable:
  On HP-UX (IA) for Oracle RAC 9i:

  ```
  $ export IPC_LIB_PATH=/opt/VRTSvcs/rac/lib/hpux64/\
  libskgxp<oracleversion>_ver24_64.sl
  ```

On HP-UX (PA) for Oracle RAC 9i:

```
$ export IPC_LIB_PATH=/opt/VRTSvcs/rac/lib/pa20_64/\
libskgxp<oracleversion>_ver24_64.sl
```

On HP-UX (IA) for Oracle RAC 10g Release 1:

```
$ export IPC_LIB_PATH=/opt/VRTSvcs/rac/lib/hpux64/\
libskgxp<oracleversion>_ver25_64.sl
```

On HP-UX (PA) for Oracle RAC 10g Release 1:

```
$ export IPC_LIB_PATH=/opt/VRTSvcs/rac/lib/pa20_64/\
libskgxp<oracleversion>_ver25_64.sl
```

where *<oracleversion>* is 9 (for Oracle RAC 9i) or 10 (for Oracle RAC 10g)

Performing a full upgrade to 5.1 SP1 RP2 on a standalone system

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

To upgrade to 5.1 SP1 RP2 on a standalone system

1. 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. 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.
5. Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

   ```
   # mount | grep vxfs
   ```

6 Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

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

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

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

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

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

Verify that no volumes remain open:

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

10 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

11 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
```

12 Restart the system.

```
# shutdown -r now
```

Performing a rolling upgrade to 5.1 SP1 RP2 on 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

About rolling upgrades

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

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

You can perform a rolling upgrade for the following products:

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

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

Prerequisites for 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. The following procedure assumes four nodes: node1, node2, node3, node4.

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 Relink the Oracle RAC libraries with SF Oracle RAC libraries.

See "Relinking Oracle RAC libraries with the SF Oracle RAC libraries" on page 189.

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

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

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

16 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
   ```

Chapter 3

Uninstalling version 5.1 SP1 RP2

This chapter includes the following topics:

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

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 RP2 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 RP2 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 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.
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. Stop all VxVM volumes. For each disk group enter:

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

   Verify that no volumes remain open.

   ```
   # vxprint -Aht -e v_open
   ```
5. Run the uninstallrp script to rollback patches, type:

   ```
   # ./uninstallrp
   ```
6. Restart the system:

   ```
   # shutdown -r now
   ```

The uninstallrp script removes 5.1 SP1 RP2 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 SP1RP2 with the Web-based installer

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

To uninstall 5.1 SP1RP2

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.