Veritas Storage Foundation™ and High Availability 6.0.4 Release Notes - Linux

October 2013

Symantec™

Veritas Storage Foundation and High Availability Release Notes

Product version: 6.0.4

Document version: 6.0.4 Rev 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

Veritas Storage Foundation and High Availability Release Notes

This document includes the following topics:

About this document

This document provides important information about Veritas Storage Foundation and High Availability (SFHA) version 6.0.4 for Linux. Review this entire document before you install or upgrade SFHA.

The information in the Release Notes supersedes the information provided in the product documents for SFHA.

This is "Document version: 6.0.4 Rev 0" of the *Veritas Storage Foundation and High Availability Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

https://sort.symantec.com/documents

Component product release notes

In addition to reading this Release Notes document, review the component product release notes before installing the product.

Product guides are available at the following location on the software media in PDF formats:

`/docs/product_name`

Symantec recommends copying the files to the `/opt/VRTS/docs` directory on your system.

This release includes the following component product release notes:

- *Veritas Storage Foundation Release Notes* (6.0.4)
- *Veritas Cluster Server Release Notes* (6.0.4)

About Veritas Storage Foundation High Availability

Storage Foundation High Availability includes Veritas Storage Foundation and Veritas Cluster Server. Veritas Cluster Server adds high availability functionality to Storage Foundation products.

Before you install the product, read the *Veritas Storage Foundation and High Availability Release Notes*.

To install the product, follow the instructions in the *Veritas Storage Foundation and High Availability Installation Guide*.

For HA installations, also read the *Veritas Cluster Server Release Notes*.

About Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Symantec products.

SORT can help you do the following:

Prepare for your next installation or upgrade	■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture. ■ Analyze systems to determine if they are ready to install or upgrade Symantec products. ■ Download the latest patches, documentation, and high availability agents from a central repository. ■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems.
Manage risks	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository. ■ Identify and mitigate system and environmental risks. ■ Display descriptions and solutions for hundreds of Symantec error codes.
Improve efficiency	■ Find and download patches based on product version and platform. ■ List installed Symantec products and license keys. ■ Tune and optimize your environment.

Note: Certain features of SORT are not available for all products. Access to SORT is available at no extra cost.

To access SORT, go to:

https://sort.symantec.com

Important release information

- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website: http://www.symantec.com/docs/TECH164885

- For the latest patches available for this release, go to:
  https://sort.symantec.com/
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:
  http://www.symantec.com/docs/TECH170013
  Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

Changes in this release

This section describes the changes introduced in this release.

Support for SLES11 SP3

SFHA now supports SUSE Linux Enterprise Server 11 Service Pack 3.

See "Supported Linux operating systems " on page 11.

Supported Oracle configurations

In 6.0.4 release, SFDB tools support Oracle 12c release for Oracle databases on Linux platform.

Note: SFDB does not support the Multitenant database feature for Oracle 12c.

Support for Oracle 12c installation using the Symantec script-based installer

You can now use the Symantec script-based installer to install or upgrade to Oracle 12c.

Note: SFHA Oracle supports basic installation of Oracle 12c. Oracle 12c features are not yet supported.

No longer supported

The following features are not supported in this release of SFHA products:

- The `fsppmk` command is deprecated and can no longer be used to create SmartTier placement policies.

Veritas Storage Foundation for Databases (SFDB) tools features which are no longer supported

The following Storage Foundation for Databases (SFDB) tools features are not supported in this release:

- FlashSnap reverse resync
- Checkpoint policy and Checkpoint quotas
- Interactive modes in clone and rollback

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems

This section lists the supported operating systems for this release of Veritas products. For current updates, visit the Symantec Operation Readiness Tools Installation and Upgrade page: https://sort.symantec.com/land/install_and_upgrade.

Table 1-1 shows the supported operating systems for this release.

Table 1-1 Supported operating systems

Operating systems	Levels	Kernel version	Chipsets
SUSE Linux Enterprise 11	SP2, SP3	3.0.13-0.27 3.0.76-0.11	64-bit x86, EMT*/Opteron 4.1 64-bit only

* Extended Memory Technology

Note: Only 64-bit operating systems are supported.

If your system is running an older version of SUSE Linux Enterprise Server, upgrade it before attempting to install the Veritas software. Consult the SUSE documentation for more information on upgrading or reinstalling your operating system.

Symantec supports only SUSE distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel Application Binary Interface (ABI) compatibility.

Required Linux RPMs for SFHA

Make sure you install the following operating system-specific RPMs on the systems where you want to install or upgrade SFHA. SFHA will support any updates made to the following RPMs, provided the RPMs maintain the ABI compatibility.

Table 1-2 lists the RPMs that SFHA requires for a given Linux operating system.

Table 1-2 Required RPMs

Operating system	Required RPMs
SLES 11 SP2	coreutils-8.12-6.19.1.x86_64.rpm ed-0.2-1001.30.1.x86_64.rpm findutils-4.4.0-38.26.1.x86_64.rpm glibc-2.11.3-17.31.1.x86_64.rpm glibc-32bit-2.11.3-17.31.1.x86_64.rpm ksh-93u-0.6.1.x86_64.rpm libacl-2.2.47-30.34.29.x86_64.rpm libacl-32bit-2.2.47-30.34.29.x86_64.rpm libgcc46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm libgcc46-4.6.1_20110701-0.13.9.x86_64.rpm libncurses5-5.6-90.55.x86_64.rpm libstdc++46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm libstdc++46-4.6.1_20110701-0.13.9.x86_64.rpm module-init-tools-3.11.1-1.21.1.x86_64.rpm pam-32bit-1.1.5-0.10.17.x86_64.rpm parted-2.3-10.21.18.x86_64.rpm

Table 1-2 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 11 SP3	coreutils-8.12-6.25.27.1.x86_64.rpm
	ed-0.2-1001.30.1.x86_64.rpm
	findutils-4.4.0-38.26.1.x86_64.rpm
	glibc-2.11.3-17.54.1.x86_64.rpm
	glibc-32bit-2.11.3-17.54.1.x86_64.rpm
	ksh-93u-0.18.1.x86_64.rpm
	libacl-2.2.47-30.34.29.x86_64.rpm
	libacl-32bit-2.2.47-30.34.29.x86_64.rpm
	libgcc_s1-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libgcc_s1-4.7.2_20130108-0.15.45.x86_64.rpm
	libncurses5-5.6-90.55.x86_64.rpm
	libstdc++6-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libstdc++6-4.7.2_20130108-0.15.45.x86_64.rpm
	module-init-tools-3.11.1-1.28.5.x86_64.rpm
	pam-32bit-1.1.5-0.10.17.x86_64.rpm
	parted-2.3-10.38.16.x86_64.rpm

Mandatory patch required for Oracle Bug 4130116

If you are running Oracle versions 9.2.0.6 or 9.2.0.7, you must apply the Oracle patch for Oracle Bug 4130116. Contact Oracle to obtain this patch, and for details on how to apply it.

For more information, refer to the following TechNote:

http://www.symantec.com/docs/HOWTO19718

Storage Foundation for Databases features supported in database environments

Storage Foundation for Databases (SFDB) product features are supported for the following database environments:

Table 1-3 SFDB features supported in database environments

Veritas Storage Foundations feature	DB2	Oracle	Oracle RAC	Sybase	Sybase ASE CE
Oracle Disk Manager	No	Yes	Yes	No	No
Cached Oracle Disk Manager	No	Yes	No	No	No
Concurrent I/O	Yes	Yes	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes	Yes	Yes
Database Storage Checkpoints **Note:** Requires Enterprise license	Yes	Yes	Yes	No	No
Database Flashsnap **Note:** Requires Enterprise license	Yes	Yes	Yes	No	No
SmartTier for Oracle **Note:** Requires Enterprise license	No	Yes	Yes	No	No

Notes:

- SmartTier is an expanded and renamed version of Dynamic Storage Tiering (DST).
- Storage Foundation for Databases (SFDB) tools Database Checkpoints, Database Flashsnap, and SmartTier for Oracle are supported with an Enterprise product license.

For the most current information on Storage Foundation products and single instance Oracle versions supported, see:

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

Review the current Oracle documentation to confirm the compatibility of your hardware and software.

Veritas Storage Foundation memory requirements

Symantec recommends 2 GB of memory over the minimum requirement for the operating system.

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

Fixed issues

This section includes the issues fixed since the previous major release. The fixed issues are presented in separate tables for each applicable minor release.

Installation and upgrades fixed issues

This section describes the installation and upgrade issues fixed since the previous major release.

Installation and upgrades: issues fixed in 6.0.4

In this release, there were no fixed issues related to installation and upgrades.

Installation and upgrades: issues fixed in 6.0.3

This section describes the installation and upgrade issues fixed in 6.0.3.

Table 1-4 Installation and upgrades 6.0.3 fixed issues

Incident	Description
2967125	Eval injection vulnerability in the Digest module before 1.17 for Perl allows context-dependent attackers to execute arbitrary commands via the new constructor.
2851403	Veritas File System modules may fail to unload if SmartMove is enabled and a break-off snapshot volume has been reattached. **Note:** This issue is not reproducible on 6.0.3

Veritas Storage Foundation and High Availability fixed issues

Issues fixed for Veritas Storage Foundation and High Availability (SFHA) include issues fixed for Veritas File System and Veritas Volume Manager.

See "Veritas File System fixed issues" on page 16.

See "Veritas Volume Manager fixed issues" on page 18.

Veritas File System fixed issues

This section describes Veritas File System issues fixed since the previous major release.

Veritas File System: issues fixed in 6.0.4

Table 1-5 describes the incidents that are fixed in Veritas File System (VxFS) in 6.0.4.

Table 1-5 Veritas File System 6.0.4 fixed issues

Incident	Description
3312030	The default quota support on Veritas File System version 6.0.4 is changed to 32 bit.
3270210	On SUSE Linux Enterprise Server 11 (SLES 11) SP3, no support is provided for `VRTSfsadv`, `VRTSfssdk`, `VRTSvxfs` and `VRTSvxvs`.
3270200	Support for SLES 11 SP3 is added.
3268931	Support for SLES 11 SP3 is added.
3268916	Support for SLES 11 SP3 is added.
3268908	Support for SLES 11 SP3 is added.
3257467	Support for SLES 11 SP3 is added.
3248955	The system fails to build modules on SLES 11 SP3 for Veritas Oracle Disk Manager (ODM), because it cannot find the `utsrelease.h` file needed to build the ODM module for SP3 on SLES 11.
3247752	The system panics with the following message during the internal stress tests: `Unable to handle kernel paging request at <addr>`
3214816	With the DELICACHE feature enabled, frequent creation and deletion of the inodes of a user may result in corruption of the user quota file.
3140990	Requirement for the ability to turn off VxFS's invalidation of pages for some Network File System (NFS) workloads.
3121933	The `pwrite(2)` function fails with the EOPNOTSUPP error.
3101418	The current time returned by the operating system (Oracle error code `ORA-01513`) during Oracle startup is invalid.

Table 1-5 Veritas File System 6.0.4 fixed issues *(continued)*

Incident	Description
3089210	The message `V-2-17: vx_iread_1 filesystem file system inode inode number marked bad incore` is displayed in the system log.
3068902	In case of stale NFS mounts, the `statfs()` function calls on non-VxFS file systems may cause df commands to hang.
3042485	During internal stress testing, the `f:vx_purge_nattr:1` assert fails.
2999493	The file system check validation fails after a successful full fsck during the internal testing with the message: `run_fsck : First full fsck pass failed, exiting`
2991880	In low memory conditions on a VxFS, certain file system activities may seem to be non-responsive.
2983248	The `vxrepquota`(1M) command dumps core.
2977828	The file system is marked bad after an inode table overflow error occurs.
2966277	Systems with high file system activity like read/write/open/lookup may panic the system.
2963763	When the `thin_friendly_alloc()` and `deliache_enable()` functionality is enabled, VxFS may enter a deadlock.
2926684	In rare cases, the system may panic while performing a logged write.
2908391	It takes a long time to remove checkpoints from the VxFS file system, when there are a large number of files present.
2839871	On a system with DELICACHE enabled, several file system operations may hang.
2834192	You are unable to mount the file system after the full `fsck`(1M) utility is run.
2756779	The code is modified to improve the fix for the read and write performance concerns on Cluster File System (CFS) when it runs applications that rely on the POSIX file-record using the fcntl lock.

Veritas File System: issues fixed in 6.0.3

Table 1-6 describes the incidents that are fixed in Veritas File System in 6.0.3.

Table 1-6 Veritas File System 6.0.3 fixed issues

Incident	Description
3004466	Installation of 5.1SP1RP3 fails on RHEL 6.3.
2895743	Accessing named attributes for some files seems to be slow.
2893551	When nfs connections are under load , file attribute information is replaced by question marks.
2881211	File ACLs not preserved in checkpoints properly if file has hardlink.
2858683	Reserve extent attributes changed after vxrestore, only for files greater than 8192bytes.
2857751	The internal testing hits the assert "f:vx_cbdnlc_enter:1a".
2857629	File system corruption can occur requiring a full fsck of the system.
2821152	Internal Stress test hit an assert “f:vx_dio_physio:4,1” on locally mounter file system.
2806466	fsadm -R resulting in panic at LVM layer due to vx_ts.ts_length set to 2GB.
2773383	Read/Write operation on a memory mapped files seems to be hung.
2756779	Write and read performance concerns on CFS when running applications that rely on posix file-record locking (fcntl).
2641438	Modifications to user name space extended attributes are lost after a system reboot.
2624262	fsdedup.bin hit oops at vx_bc_do_brelse.
2611279	Filesystem with shared extents may panic.
2417858	VxFS quotas do not support 64 bit limits.

Veritas Volume Manager fixed issues

This section describes Veritas Volume Manager issues fixed since the previous major release.

Veritas Volume Manager: issues fixed in 6.0.4

Table 1-7 describes the incidents that are fixed in Veritas Volume Manager in 6.0.4.

Table 1-7 Veritas Volume Manager 6.0.4 fixed issues

Incident	Description
3321269	The command `vxunroot` may hang during un-encapsulation of root disk.
3240858	The `/etc/vx/vxesd/.udev_lock` file may have different permissions at different instances.
3199056	Veritas Volume Replicator (VVR) primary system panics in the `vol_cmn_err` function due to the VVR corrupted queue.
3186971	The logical volume manager (LVM) configuration file is not correctly set after turning on DMP native support. As a result, the system is unbootable.
3186149	On Linux System with LVM version 2.02.85, on enabling `dmp_native_support` LVM volume Groups will disappear
3182350	If there are more than 8192 paths in the system, the `vxassist(1M)` command hangs when you create a new VxVM volume or increase the existing volume's size.
3182175	The `vxdisk -o thin,fssize list` command can report incorrect File System usage data.
3177758	Performance degradation is seen after upgrade from SF 5.1SP1RP3 to SF 6.0.1 on Linux.
3162418	The `vxconfigd(1M)` command dumps core due to wrong check in `ddl_find_cdevno()` function.
3146715	Rlinks do not connect with Network Address Translation (NAT) configurations on Little Endian Architecture.
3130353	Continuous disable or enable path messages are seen on the console for EMC Not Ready (NR) devices.
3121380	I/O of Replicated Volume Group (RVG) hangs after one data volume is disabled.
3091916	The Small Computer System Interface (SCSI) I/O errors overflow the syslog.
3063378	Some VxVM commands run slowly when EMC PowerPath presents and manages "read only" devices such as EMC SRDF-WD or BCV-NR.
3015181	I/O can hang on all the nodes of a cluster when the complete non-A or A class of storage is disconnected.
3012929	The `vxconfigbackup(1M)` command gives errors when disk names are changed.

Table 1-7 Veritas Volume Manager 6.0.4 fixed issues *(continued)*

Incident	Description
3010191	Previously excluded paths are not excluded after upgrade to VxVM 5.1SP1RP3.
2986596	The disk groups imported with mix of standard and clone Logical Unit Numbers (LUNs) may lead to data corruption.
2972513	Cluster Volume Manager (CVM) and PGR keys from shared data disks are not removed after stopping VCS.
2959733	Handling the device path reconfiguration in case the device paths are moved across LUNs or enclosures to prevent the `vxconfigd(1M)` daemon coredump.
2882312	If an SRL fault occurs in the middle of an I/O load, and you immediately issue a read operation on data written during the SRL fault, the system returns old data.
2857360	The `vxconfigd(1M)` command hangs when the `vol_use_rq` tunable of VxVM is changed from 1 to 0.
2857044	System crashes while resizing a volume with data change object (DCO) version 30.
3052770	The `vradmin syncrvg` operation with a volume set fails to synchronize the secondary RVG with the primary RVG.

Veritas Volume Manager: issues fixed in 6.0.3

Table 1-8 describes the incidents that are fixed in Veritas Volume Manager in 6.0.3.

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues

Incident	Description
3002770	Accessing NULL pointer in dmp_aa_recv_inquiry() caused system panic.
2971746	For single-path device, bdget() function is being called for each I/O, which cause high cpu usage and leads to I/O performance degradation.
2970368	Enhancing handling of SRDF-R2 WD devices in DMP.
2965910	vxassist dump core with the `-o ordered` option.
2964169	In multiple CPUs environment, I/O performance degradation is seen when I/O is done through VxFS and VxVM specific private interface.

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2962262	Uninstallation of DMP fails in presence of other multi-pathing solutions.
2948172	Executing the `vxdisk -o thin,fssize list` command can result in panic.
2943637	DMP IO statistic thread may cause out of memory issue so that OOM(Out Of Memory) killer is invoked and causes system panic.
2942609	Message displayed when user quits from Dynamic Reconfiguration Operations is shown as error message.
2940446	Full fsck hangs on I/O in VxVM when cache object size is very large
2935771	In the VVR environment, RLINK disconnects after the master is switched.
2933138	panic in voldco_update_itemq_chunk() due to accessing invalid buffer
2930569	The LUNs in 'error' state in output of 'vxdisk list' cannot be removed through DR(Dynamic Reconfiguration) Tool.
2928764	SCSI3 PGR registrations fail when dmp_fast_recovery is disabled.
2919720	vxconfigd core in rec_lock1_5()
2919714	exit code from vxevac is zero when migrating on thin luns but FS is not mounted
2919627	Dynamic Reconfiguration tool should be enhanced to remove LUNs feasibly in bulk.
2919318	The I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing.
2916094	Enhancements have been made to the Dynamic Reconfiguration Tool(DR Tool) to create a separate log file every time DR Tool is started, display a message if a command takes longer time, and not to list the devices controlled by TPD (Third Party Driver) in 'Remove Luns' option of DR Tool.
2915063	Rebooting VIS array having mirror volumes, master node panicked and other nodes CVM FAULTED
2911040	Restore from a cascaded snapshot when its source is DETACHED leaves the volume in unusable state
2910043	Avoid order 8 allocation by vxconfigd in node reconfig.
2899173	vxconfigd hang after executing the `vradmin stoprep` comand.

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2898547	vradmind on VVR Secondary Site dumps core, when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2892983	vxvol dumps core if new links are added while the operation is in progress.
2886402	vxconfigd hang while executing tc ./scripts/ddl/dmpapm.tc#11.
2886333	The `vxdg(1M) join` command should not allow mixing clone and non-clone disks in a DiskGroup.
2878876	vxconfigd dumps core in vol_cbr_dolog() due to race between two threads processing requests from the same client.
2869594	Master node panics due to corruption if space optimized snapshots are refreshed and "vxclustadm setmaster" is used to select master.
2866059	Improving error messages hit during the `vxdisk resize` operation.
2859470	SRDF R2 with EFI label is not recognized by VxVM and showing in error state
2858853	vxconfigd coredumps in dbf_fmt_tbl on the slave node after a Master Switch if you try to remove a disk from the DG.
2851403	The `vxportal` and `vxfs` processes are failed to stop during first phase of rolling upgrade.
2851085	DMP doesn't detect implicit LUN ownership changes for some of the dmpnodes.
2839059	vxconfigd logged warning `cannot open /dev/vx/rdmp/cciss/c0d device to check for ASM disk format.`
2837717	The `vxdisk(1M) resize` command fails if `da name` is specified.
2836798	Prevent DLE on simple/sliced disk with EFI label
2834046	NFS migration failed due to device reminoring.
2833498	vxconfigd hangs while reclaim operation is in progress on volumes having instant snapshots
2826125	VxVM script daemon is terminated abnormally when it is invoking with exact the same process id of the last invocation.
2815517	vxdg adddisk should not allow mixing clone & non-clone disks in a DiskGroup

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
2801962	Grow of a volume takes significantly large time when the volume has version 20 DCO (Data Change Object) attached to it
2798673	System panics in voldco_alloc_layout() while creating volume with instant DCO.
2779580	Secondary node gives configuration error (no Primary RVG) after reboot of master node on Primary site.
2753954	At cable disconnect on port1 of dual-port FC HBA, paths via port2 are also marked SUSPECT.
2744004	vxconfigd is hung on the VVR secondary node during VVR configuration.
2715129	vxconfigd hangs during Master takeover in a CVM (Clustered Volume Manager) environment.
2692012	The vxevac move error message needs to be enhanced to be less generic and give clear message for failure..
2619600	Live migration of virtual machine having SFHA/SFCFSHA stack with data disks fencing enabled, causes service groups configured on virtual machine to fault.
2567618	VRTSexplorer coredumps in vxcheckhbaapi/print_target_map_entry
2510928	Extended attributes for SRDF luns reported as Mirror with EMC (VMAX array)
2398416	vxassist dumps core while creating volume after adding attribute "wantmirror=ctlr" in default vxassist rulefile
2273190	Incorrect setting of the UNDISCOVERED flag can lead to database inconsistency.
2149922	Record the diskgroup import and deport events in syslog.
2000585	The `vxrecover -s` command does not start any volumes if a volume is removed whilst it is running.
1982965	vxdg import DG fails if da-name is based on naming scheme which is different from the prevailing naming scheme on the host.
1973983	`vxunreloc` fails when dco plex is in DISABLED state.
1903700	vxassist remove mirror does not work if nmirror and alloc is specified on VxVM 3.5

Table 1-8 Veritas Volume Manager 6.0.3 fixed issues *(continued)*

Incident	Description
1901838	Incorrect setting of the Nolicense flag can lead to dmp database inconsistency.
1859018	The `link detached from volume` warnings are displayed when a linked-breakoff snapshot is created.
1765916	VxVM socket files don't have proper write protection
1725593	The `vxdmpadm listctlr` command has to be enhanced to print the count of device paths seen through the controller.

LLT, GAB, and I/O fencing fixed issues

This section describes LLT, GAB, and I/O fencing issues fixed since the previous major release.

LLT, GAB, and I/O fencing fixed issues in 6.0.4

Table 1-10 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-9 LLT, GAB, and I/O fencing fixed issues

Incident	Description
3106493	If for some reason, kernel components of the Veritas Cluster Server (VCS) software stack are stopped and restarted in quick succession, then during a restart, the cluster communication may fail.
3137520	Low Latency Transport (LLT) detects a duplicate node ID incorrectly, even if the nodes are using different ethernet SAP values.

LLT, GAB, and I/O fencing fixed issues in 6.0.1

Table 1-10 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-10 LLT, GAB, and I/O fencing fixed issues

Incident	Description
2708619	If you set the scsi3_disk_policy attribute to dmp, you cannot enable the Veritas fencing module (VxFEN). The VxFEN source code is updated to pick up the dmp device path that contains the full disk name instead of a partition or slice.

Table 1-10 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
2845244	`vxfen` startup script gives error `grep: can't open /etc/vxfen.d/data/cp_uid_db`. The error comes because vxfen startup script tries to read a file that might not be present. This error is typically seen when starting vxfen for the very first time after installation.
2554167	Setting peerinact value to 0 in the `/etc/llttab` file floods the system log file with large number of log messages.

Storage Foundation for Databases (SFDB) tools fixed issues

This section describes Storage Foundation for Databases (SFDB) tools issues fixed since the previous major release.

Storage Foundation for Databases (SFDB) tools: issues fixed in 6.0.4

Table 1-11 describes the incidents that are fixed in SFDB tools in 6.0.4.

Table 1-11 Storage Foundation for Databases (SFDB) tools 6.0.4 fixed issues

Incident	Description
3211391 (3211388)	During a Veritas Database Edition (DBED) instant checkpoint cloning, if you enable the Block Change Tracking feature, it results in an `ORA-00600` error.
3277940 (3290416)	Some DBED operations may fail with the following error message: `ORA-01406: fetched column value was truncated`

Storage Foundation for Databases (SFDB) tools: issues fixed in 6.0.3

Table 1-12 describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in 6.0.3.

Table 1-12 Storage Foundation for Databases (SFDB) tools 6.0.3 fixed issues

Incident	Description
3030663	`dbed_vmclonedb` does not read pfile supplied by `-p 'pfile_modification_file'` option.

Known issues

This section covers the known issues in this release.

VMware virtual environment specific issues

VMwareDisks agent may fail to start or storage discovery may fail if SELinux is running in enforcing mode [2854053]

The VMwareDisks agent and discFinder binaries refer to the libvmwarevcs.so shared library. SELinux security checks prevent discFinder from loading the libvmwarevcs.so library, which requires text relocation. If SELinux is running in enforcing mode, these two executables may report the "`Permission denied`" error and may fail to execute.

Workaround: Enter the following command and relax the security check enforcement on the Symantec libvmwarevcs.so library:

```
# chcon -t textrel_shlib_t '/opt/VRTSvcs/lib/libvmwarevcs.so'
```

vCenter Integrated Installer may fail to install VCS on OEL 6.x guests [2905806]

In a VMware environment, if you create a virtual machine by setting **Guest Operating System Version** to **Oracle Linux 4/5/6 (64-bit)**, you cannot use the vSphere Client to install Veritas Cluster Server (VCS) on the virtual machine. The installation wizard is unable to distinguish between OEL 5.x and OEL 6.x guests and therefore does not support installation on OEL 6.x guests.

Workaround: When you create the virtual machine, set the **Guest Operating System Version** to **RedHat Enterprise Linux 6 (64-bit)** and then use vCenter Integrated Installer. Alternatively, install VCS from the command line, using one of the installer-based options. For more information, see the *Veritas Cluster Server Installation Guide*.

Oracle configuration wizard may not allow to add or remove listeners [2868770]

While configuring an Oracle application, the configuration wizard fails to accept new listeners or to remove the existing listeners if you navigate back from the virtual IPs page to the listener page.

Workaround: No workaround.

Symantec High Availability tab does not report LVMVolumeGroup resources as online [2909417]

The Symantec High Availability tab does not automatically report the online status of activated LVMVolumeGroup resources in the following case:

- If you created the VCS cluster as part of the Symantec High Availability Configuration Wizard workflow.

Workaround: Start the LVMVolumeGroup resources from the Symantec High Availability tab. For more information, see the Symantec High Availability Solutions Guide for VMware.

Unable to configure application when a new LSI Logic SCSI controller is added dynamically [2937623]

In case of SLES operating system, it is possible that hot plug drivers do not get loaded and hence the operating system is unable to recognize the newly added SCSI controllers.

Workaround: Load the drivers into the kernel before attempting to configure application or add a new SCSI controller. You can add the drivers by using the command:

```
# modprobe acpiphp
# modprobe pci_hotplug
```

During snapshot reversal, SFHA commands fail, and an unrelated error message may appear [2939177]

If a virtual machine is under SFHA control, and you revert to the virtual machine configuration snapshot, the SFHA configuration file moves to the read- only mode. During this time, if you run a SFHA command, the command fails. Also, the following unrelated error message appears in the log entries:

```
Failed to delete the attribute key <value of the key>
 of attribute DiskPaths in the VCS configuration with error 11335
(Configuration must be ReadWrite : Use haconf -makerw)
```

Workaround: No workaround. This message is safe to ignore.

vSphere UI may not show applications and their status [2938892]

The vSphere UI may show that it cannot connect to the virtual machine and may prompt for username and password of the virtual machine. Even after entering correct username and password, vSphere UI may not show the list of applications

monitored by VCS. This can happen if the VOM MH configuration has authorization problems.

Workaround: Restart the xprtld daemon on the virtual machine with the following commands:

```
/etc/init.d/xprtld stop
/etc/init.d/xprtld start
```

SSO configuration fails if the system name contains non-English locale characters

If you install the Symantec High Availability guest components on a system that has non-English locale characters in its name, then the SSO configuration between such a system and the Symantec High Availability Console host fails. (2910613)

Workaround: Ensure that the system name does not contain any non-English locale characters.

Error while unconfiguring the VCS cluster from the Symantec High Availability tab

This issue may occur if you try to unconfigure the VCS cluster from the Symantec High Availability tab in VMware vSphere Client. (3011461)

The VCS cluster unconfiguration task fails with the following error:

```
Failed to stop the HAD service. Node='systemname', Error=00000425."
```

Workaround: Perform the following steps to unconfigure the cluster:

1. Forcefully stop the VCS High Availability Daemon (HAD) process on the system using the following command:

   ```
   taskkill /f had.exe
   ```

2. If HAD starts again, stop HAD using the following command:

   ```
   hastop -local
   ```

3. Ensure that the HAD process is in the stopped state.
4. Launch the VCS Cluster Configuration Wizard (VCW) and then delete the cluster using the VCW wizard flow.

   Refer to the *VCS Administrator's Guide* for more information about VCW.

Application instances are ignored if its dependent storage is configured for other instance [2966123]

Application instances are ignored if its dependent storage is already configure under VCS for other instances. The following message is logged in healthview_A.log when you try to configure a new application instance.

```
Skipping the <Application> instance: <instance_name>,
because the storage is already configured or an application
running under VCS control.
```

Workaround: Reconfigure all the instances that use common storage in a single run of the wizard.

Installation and upgrade known issues

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

Error messages in syslog (1630188)

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

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

Web installer does not ask for authentication after the first session if the browser is still open (2509330)

If you install or configure SFHA and then close the Web installer, if you have other browser windows open, the Web installer does not ask for authentication in the subsequent sessions. Since there is no option to log out of the Web installer, the session remains open as long as the browser is open on the system.

Workaround: Make sure that all browser windows are closed to end the browser session and subsequently log in again.

Stopping the Web installer causes Device Busy error messages (2633924)

If you start the Web installer, and then perform an operation (such as prechecking, configuring, or uninstalling), you may get an error message saying the device is busy.

Workaround: Do one of the following:

- Kill the start.pl process.
- Start the webinstaller again. On the first Web page you see that the session is still active. Either take over this session and finish it or terminate it directly.

The uninstaller does not remove all scripts (2696033)

After removing SFHA, some of the RC scripts remain in the `/etc/rc*.d/` folder. This is due to an issue with the chkconfig rpm in RHEL6 and updates. You can manually remove the scripts from the `/etc/rc*.d/` folder after removing the VxVM packages.

Workaround:

Install the chkconfig-1.3.49.3-1 chkconfig rpm from the RedHat portal. Refer to the following links:

http://grokbase.com/t/centos/centos/117pfhe4zz/centos-6-0-chkconfig-strange-behavior

http://rhn.redhat.com/errata/RHBA-2012-0415.html

Installing DMP with a keyless license or DMP-only license does not enable DMP native support for LVM root volumes (2874810)

When you install DMP with a keyless license or DMP-only license, the tunable parameter dmp_native_support is set to on. However, the DMP native support is not enabled for LVM root volumes. The DMP native support is enabled for non-root LVM volumes.

Workaround:

After package installation, use the following command to enable the DMP support for root LVM volumes.

```
# vxdmpadm settune dmp_native_support=on
```

Then reboot the system.

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

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to Veritas Storage Foundation (SF) 6.0.1, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/openv`), then while upgrading to SF 6.0.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure RPMs `VRTSpbx`, `VRTSat`, and `VRTSicsco`. This causes NetBackup to stop working.

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the `/usr/openv/netbackup/bin/version` file and `/usr/openv/netbackup/version` file to the `/tmp` directory. If you have clustered NetBackup installed, you must also copy the `/usr/openv/netbackup/bin/cluster/NBU_RSP` file to the `/tmp` directory. After you unmount the NetBackup file system, manually copy these two version files from `/tmp` to their original directories. If you have clustered NetBackup installed, you must also copy the `/usr/openv/netbackup/bin/cluster/NBU_RSP` file from `/tmp` to its original directory.

If the `version` files' directories do not exist, create the directories:

```
# mkdir -p /usr/openv/netbackup/bin
# mkdir -p  /usr/openv/netbackup/bin
```

Run the installer to finish the upgrade process. After upgrade process completes, remove the two version files and their directories.

If your system is already affected by this issue, then you must manually install the `VRTSpbx`, `VRTSat`, and `VRTSicsco` RPMs after the upgrade process completes.

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

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

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

```
The partioning on disk /dev/sda is not readable by
The partioning tool parted, which is used to change the
partition table.
You can use the partitions on disk /dev/sda as they are.
You can format them and assign mount points to them, but you
```

```
cannot add, edit, resize, or remove partitions from that
disk with this tool.
```

Or

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

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

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

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

Upgrading from Veritas Storage Foundation 5.1 Service Pack 1 Rolling Patch 2 to 6.0.1 with rootability enabled fails (2581313)

Upgrading from Veritas Storage Foundation (SF) 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 to 6.0.1 while using an encapsulated root disk fails because the post installation scripts of Veritas Volume Manager (VxVM) are unable to start the `initrd` daemon.

Workaround: To upgrade from 5.1 SP1 RP2 to 6.0.1 while using an encapsulated root disk, you must reinstall the nash utility on the system prior to the upgrade.

To upgrade from 5.1 SP1 RP2 to 6.0.1 while using an encapsulated root disk

1. Encapsulate the root disk.
2. Reinstall the nash utility.
3. Upgrade to the SF 6.0.1 release.

During upgrade from 5.1SP1 to 6.0.1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name is used by a deported disk group (2280560)

During an upgrade from SFHA 5.1 SP1 to SFHA 6.0.1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name for the split operation is used by an existing deported disk group.

Workaround:

Specify a different disk group name as a target for the split operation.

After finishing a kernel upgrade on a master node the cvm group on a slave node does not come online (2439439)

After successfully finishing a kernel upgrade on one node, the cvm group does not come online on the second node.

Workaround: Check that your cluster is not in a jeopardy state before you perform a rolling upgrade.

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups [2574731]

The service groups freeze due to upgrading using the product installer if you stopped the installer after the installer already stopped some of the processes and then resumed the upgrade.

Workaround:

You must unfreeze the service groups manually after the upgrade completes.

To unfreeze the service groups manually

1 List all the frozen service groups

```
# hagrp -list Frozen=1
```

2 Unfreeze all the frozen service groups:

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

After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon you change the locale of nodes that use it. The vxconfig daemon starts at boot. If you have changed locale, you need to restart the daemon.

Workaround: See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more information on vxconfigd daemon recovery.

Issue with soft links getting deleted in a manual upgrade

While performing a manual upgrade (from 5.1 to 6.0) of the VRTSvlic RPM, some of the soft links created during your previous installation are deleted. As a result, `vxkeyless` binary is not found in its specified path.

To prevent this, use the `--nopreun` option.

For example: `rpm -Uvh --nopreun VRTSvlic-3.02.61.003-0.x86_64.rpm`

While upgrading the VCS stack from a version prior to VCS 5.1, reconfiguration of MultiNICA IPv4RouteOptions attribute is required (2003864)

The 5.1SP1 MultiNICA agent now uses `ip` command by default. Due to behavioral differences in `ip` and `ifconfig` commands in regards to route configuration, MultiNICA flushes routes and sets them back for the new active device. If the MultiNICA resource configuration is not intended to make use of `ifconfig` command (see table below), you must configure IPv4RouteOptions attribute in MultiNICA resource definition.

Note: RouteOptions values are used by the `route` command where as the IPv4RouteOptions value is used by the `ip route` command. The values to be configured for these two attribute are very specific to their respective commands.

Table 1-13 Whether attributes are configured and required actions that you need to perform during upgrade

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
Configured	May or may not be configured	May or may not be configured	In this case the `ifconfig` command is used. If RouteOptions is set, attribute value is used to add/delete routes using command `route`. As the Options attribute is configured, IPv4RouteOptions values are ignored.	No need to configure IPv4RouteOptions .

Table 1-13 Whether attributes are configured and required actions that you need to perform during upgrade *(continued)*

Options	RouteOptions and/or IPv4AddrOptions	IPv4RouteOptions	Comment	Actions that you need to perform during upgrade
Not configured	May or may not be configured	Must be configured	In this case the `ip` command is used. IPv4RouteOptions must be configured and are used to add/delete routes using the `ip route` command. As Options attribute is not configured, RouteOptions value is ignored.	Configure IPv4RouteOptions and set the IP of default gateway. The value of this attribute typically resembles: IPv4RouteOptions = "default via *gateway_ip*" For example: IPv4RouteOptions = "default via 192.168.1.1"

Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

After upgrading from 5.1 to higher versions of VCS, 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 are using keyless licenses before upgrading to 5.1SP1 or higher versions of VCS. 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]
   ```

SELinux error during installation of VRTSvcsag RPM

During the installation of VRTSvcsag RPM on RHEL 5 SELinux enabled machine, you may observe following SELinux error:

```
/usr/sbin/semodule: Failed on /opt/VRTSvcs/bin/selinux/vcsag.pp!
```

This error occurs due to improper installation of the SELinux package. As a result, SELinux commands may not function properly.

Workaround: Reinstall the SELinux package and relabel filesystem by either init or fixfiles method.

Web installer has no option to remove node from a cluster

Web Installer does not provide the option to remove node from a cluster.

Workaround: Manually remove nodes from a cluster. There is no option to remove nodes available from Web Installer or CPI.

CPI installer throws CPI WARNING V-9-40-4493 if enabling DMP root support when upgrading SF to 6.0.4 (3064919)

If you upgrade SF to 6.0.4 with the dmp_native_support option enabled, CPI installer may report the following message:

```
CPI WARNING V-9-40-4493 Failed to turn on
dmp_native_support tunable on <SERVER_NAME>. Refer to
Dynamic Multi-Pathing Administrator's guide to determine
the reason for the failure and take corrective action.
```

This message is inappropriate and should not be displayed. It does not affect the upgrading functions.

Workaround:

There is no workaround for this issue. You can safely ignore the message.

SF failed to upgrade when Application HA is installed (3088810)

If you have installed both ApplicationHA 6.0 and SF 6.0.1, the installer can't upgrade SF 6.0.1 to 6.0.4. The following error message is displayed:

```
CPI ERROR V-9-30-1303 SFHA 6.0.100 does not appear to be installed
on <your system>
```

Workaround:

Use the following command to specify the exact product for the upgrade:

```
# ./installmr -prod SF60
```

Ignore the warning message that SFHA is already installed after the pre-check and continue the upgrade.

Enabling or installing DMP for native support might not migrate LVM volumes to DMP (2737452)

The `lvm.conf` file has the `obtain_device_list_from_udev` variable. If `obtain_device_list_from_udev` is set to 1, then LVM uses devices and symlinks specified by udev database, only.

Workaround:

In the `lvm.conf` file, set the `obtain_device_list_from_udev` variable to `0`. This enables LVM to recognize `/dev/vx/dmp/` devices when performing a `vgscan`, which enables volume group migration to succeed.

LLT known issues

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

LLT connections are not formed when a vlan is configured on a NIC (2484856)

LLT connections are not formed when a vlan is configured on a NIC that is already used to configure an LLT link.

Workaround: Do not specify the MAC address of a NIC in the `lltab` file while configuring LLT if you want to configure a vlan later. If you have already specified

the MAC address of a NIC, then delete the MAC address from the `lltab` file, and update the file before you restart LLT.

LLT port stats sometimes shows recvcnt larger than recvbytes (1907228)

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.

This does not impact the LLT functionality.

LLT may incorrectly declare port-level connection for nodes in large cluster configurations [1810217]

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.

LLT may fail to detect when bonded NICs come up (2604437)

When LLT is configured over a bonded NIC and that bonded NIC is DOWN with the `ifconfig` command, LLT marks the corresponding link down. When the bonded NIC is UP again using the `ifconfig` command, LLT fails to detect this change and marks the link up.

Workaround: Close all the ports and restart LLT, then open the ports again.

Cannot use CPI response files to add nodes to a cluster that is using LLT over UDP (2869763)

When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/lltab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files.

Workaround: None

GAB known issues

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

While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373)

After you unregister the gtx port with `-nodeinit` option, the `gabconfig -C` command shows refcount as 1. But when forceful `deinit` option (`gabdebug -R GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinited on user request
```

The `refcount` value is incremented by 1 internally. However, the refcount value is shown as 2 which conflicts with the `gabconfig -C` command output.

Workaround: There is no workaround for this issue.

Cluster panics during reconfiguration (2590413)

While a cluster is reconfiguring, GAB broadcast protocol encounters a race condition in the sequence request path. This condition occurs in an extremely narrow window which eventually causes the GAB master to panic.

Workaround: There is no workaround for this issue.

I/O fencing known issues

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

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 Veritas Storage Foundation 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 Veritas Storage Foundation 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: There is no workaround for this issue.

After you run the vxfenswap utility the CoordPoint agent may fault (2846389)

After you run the `vxfenswap` utility, if the value of the `FaultTolerance` attribute of the CoordPoint agent is more than the majority (more than 50%) of the coordination points then the Coordination Point agent faults.

Workaround: Manually set the value of the `FaultTolerance` attribute of CoordPoint agent to be less than the majority (more than 50%) of the coordination points.

CP server repetitively logs unavailable IP addresses (2530864)

If coordination point server (CP server) fails to listen on any of the IP addresses that are mentioned in the `vxcps.conf` file or that are dynamically added using the command line, then CP server logs an error at regular intervals to indicate the failure. The logging continues until the IP address is bound to successfully.

```
CPS ERROR V-97-51-103 Could not create socket for host
10.209.79.60 on port 14250
CPS ERROR V-97-1400-791 Coordination point server could not
open listening port = [10.209.79.60]:14250
Check if port is already in use.
```

Workaround: Remove the offending IP address from the listening IP addresses list using the `rm_port` action of the `cpsadm` command.

See the *Veritas Storage Foundation and High Availability Administrator's Guide* for more details.

Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)

Even if the cluster nodes have no registration on the CP server and if you provide coordination point server (CP server) information in the `vxfenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

Workaround: Manually add the cluster information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster information to the CP server during configuration.

The cpsadm command fails if LLT is not configured on the application cluster (2583685)

The `cpsadm` command fails to communicate with the coordination point server (CP server) if LLT is not configured on the application cluster node where you run the `cpsadm` command. You may see errors similar to the following:

```
# cpsadm -s 10.209.125.200 -a ping_cps
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using
environment variable
CPS_NODEID
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

However, if you run the `cpsadm` command on the CP server, this issue does not arise even if LLT is not configured on the node that hosts CP server. The `cpsadm` command on the CP server node always assumes the LLT node ID as 0 if LLT is not configured.

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily, or if the node is a single-node VCS configuration where LLT is not configured, then the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

Workaround: Set the value of the `CPS_NODEID` environment variable to 255. The `cpsadm` command reads the `CPS_NODEID` variable and proceeds if the command is unable to get LLT node ID from LLT.

In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)

When you start server-based I/O fencing, the node may not join the cluster and prints error messages in logs similar to the following:

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1043
Detected a preexisting split brain. Unable to join cluster.
```

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
operation failed.
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@sys1,
domaintype vx; not allowing action
```

The vxfend daemon on the application cluster queries the coordination point server (CP server) to check if the cluster members as seen in the GAB membership are registered with the CP server. If the application cluster fails to contact the CP server due to some reason, then fencing cannot determine the registrations on the CP server and conservatively assumes a pre-existing split-brain.

Workaround: Before you attempt to start VxFEN on the application cluster, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

The vxfenswap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)

The vxfenswap utility runs the vxfenconfig -o modify command over RSH or SSH on each cluster node for validation of coordination points. If you run the vxfenswap command using RSH (with the -n option), then RSH does not detect the failure of validation of coordination points on a node. From this point, vxfenswap proceeds as if the validation was successful on all the nodes. But, it fails at a later stage when it tries to commit the new coordination points to the VxFEN driver. After the failure, it rolls back the entire operation, and exits cleanly with a non-zero error code. If you run vxfenswap using SSH (without the -n option), then SSH detects the failure of validation of coordination of points correctly and rolls back the entire operation immediately.

Workaround: Use the vxfenswap utility with SSH (without the -n option).

Fencing does not come up on one of the nodes after a reboot (2573599)

If VxFEN unconfiguration has not finished its processing in the kernel and in the meantime if you attempt to start VxFEN, you may see the following error in the /var/VRTSvcs/log/vxfen/vxfen.log file:

```
VXFEN vxfenconfig ERROR V-11-2-1007 Vxfen already configured
```

However, the output of the gabconfig -a command does not list port b. The vxfenadm -d command displays the following error:

```
VXFEN vxfenadm ERROR V-11-2-1115 Local node is not a member of cluster!
```

Workaround: Start VxFEN again after some time.

Server-based fencing comes up incorrectly if default port is not mentioned (2403453)

When you configure fencing in customized mode and do no provide default port, fencing comes up. However, the `vxfenconfig -l` command output does not list the port numbers.

Workaround: Retain the "port=<port_value>" setting in the `/etc/vxfenmode` file, when using customized fencing with at least one CP server. The default port value is 14250.

Secure CP server does not connect from localhost using 127.0.0.1 as the IP address (2554981)

The `cpsadm` command does not connect to the secure CP server on the localhost using 127.0.0.1 as the IP address

Workaround: Connect the secure CP server using any of the virtual IPs that is configured with the CP server and is plumbed on the local node.

Unable to customize the 30-second duration (2551621)

When the vxcpserv process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: There is no workaround for this issue.

NIC resource gets created with incorrect name while configuring CPSSG with the configure_cps.pl script (2585229)

The name of the NIC resource created by the `configure_cps.pl` script does not come out correct when, for example, m^{th} VIP is mapped to n^{th} NIC and every m is not equal to n. In this case, although CPSSG continues to function without any problem, when you unconfigure CPSSG using `configure_cps.pl`, it fails.

Workaround: To unconfigure CPSSG, you must remove the CPSSG configuration from the VCS configuration.

The cpsadm command fails after upgrading CP server to 6.0 or above in secure mode (2846727)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSat RPM is not removed from the system, the `cpsadm` command loads the old security libraries present on the system.

As the installer runs the `cpsadm` command on the CP server to add or upgrade the SFHA cluster (application cluster), the installer also fails.

Workaround: Perform the following procedure on all of the nodes of the CP server.

To resolve this issue

1 Rename `cpsadm` to `cpsadmbin`:

```
# mv /opt/VRTScps/bin/cpsadm /opt/VRTScps/bin/cpsadmbin
```

2 Create a file `/opt/VRTScps/bin/cpsadm` with the following content:

```
#!/bin/sh
EAT_USE_LIBPATH="/opt/VRTScps/lib"
export EAT_USE_LIBPATH
/opt/VRTScps/bin/cpsadmbin "$@"
```

3 Change the permissions of the new file to 775:

```
# chmod 755 /opt/VRTScps/bin/cpsadm
```

Server-based fencing may fail to start after reinstalling the stack (2802682)

Server-based fencing may fail to start if you use the existing configuration files after reinstalling the stack.

Workaround:

After reinstalling the stack, add the client cluster information on the coordination point server because the client cluster information is removed when the stack is uninstalled. For more details, see the Setting up server-based I/O Fencing manually section in the Veritas Storage Foundation and High Availability Installation Guide. Alternatively, you can manually modify the `/etc/vxfenmode` file and the `main.cf` file to start fencing in disable mode and then configure fencing.

Hostname and username are case sensitive in CP server (2846392)

The hostname and username on the CP server are case sensitive. The hostname and username used by fencing to communicate with CP server must be in same case as present in CP server database, else fencing fails to start.

Workaround: Make sure that the same case is used in the hostname and username on the CP server.

Virtual machine may return the not-responding state when the storage domain is inactive and the data center is down (2848003)

In a Red Hat Enterprise Virtualization Environment, if the storage domain is in an inactive state and the data center is in down state, the virtual machine may return a not-responding state and the KVMGuest resource in OFFLINE state.

Workaround: To resolve this issue:

1 Activate the storage domain in RHEV-M.

2 Check that the data center is in the up state.

Fencing may show the RFSM state as replaying for some nodes in the cluster (2555191)

Fencing based on coordination point clients in Campus cluster environment may show the RFSM state as replaying for some nodes in the cluster.

Workaround:

Restart fencing on the node that shows RFSM state as replaying.

Cannot run the vxfentsthdw utility directly from the install media if VRTSvxfen package is not installed on the system (2858190)

If VRTSvxfen package is not installed on the system, then certain script files that are needed for the vxfentsthdw utility to function are not available. So, without the VRTSvxfen package installed on the system you cannot run the utility from the install media.

Workaround: Install VRTSvxfen package, then run the utility from either the install media or from the `/opt/VRTSvcs/vxfen/bin/` location.

CoordPoint agent does not report the addition of new disks to a Coordinator disk group [2727672]

The LevelTwo monitoring of the CoordPoint agent does not report a fault even if the constituent of a coordinator disk group changes due to addition of new disks in the cooridnator disk group

Workaround: There is no workaround for this issue.

Coordination point server-based fencing may fail if it is configured on 5.1SP1RP1 using 6.0.1 coordination point servers (2824472)

The 5.1SP1 installer (CPI) cannot set up trust between a 5.1SP1 client and a 6.0 or later server, because there are no separate directories for truststores in the 5.1SP1. When trust cannot be setup, the 5.1SP1 installer cannot configure 5.1SP1 clients to work with 6.0 or later CPS in secure mode.

Workaround:

Set up trust manually between the CPS and clients using the cpsat or the vcsat command. After that, CPS and client will be able to communicate properly in the secure mode.

Veritas Storage Foundation and High Availability known issues

Known issues for Veritas Storage Foundation and High Availability (SFHA) include known issues of Veritas File System and Veritas Volume Manager.

See "Veritas File System known issues" on page 46.

See "Veritas Volume Manager known issues" on page 51.

Veritas File System known issues

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

Taking a FileSnap over NFS multiple times with the same target name can result in the 'File exists' error (2353352)

The "File exists" error occurs as a result of the caching behavior of the NFS client. Because the link operation is successful, the NFS client assumes that a file with the specified target name, such as `file2::snap:vxfs:`, was created.. As a result, the NFS client caches a file with this name.

Workaround: Remove the target file after a snapshot is created. This forces the NFS client to remove the name from the cache. For example:

```
# ln file1 file2::snap:vxfs:
# rm file2::snap:vxfs:
```

Rolling upgrade from version 6.0.3 may cause corruption in the Veritas File System external quota file

The 6.0.3 release supported 64-bit quota, which allowed quota limits higher than 1 terabyte. However, there is a possibility of corruption during rolling upgrade when two nodes may be on different patch levels—one node using 64-bit quota limit and the other node using 32-bit quota limit.

Workaround: Disable quotas before you start rolling upgrade. Back up the external quota file. This will help to restore the original file in the event that the quota file becomes corrupted during the upgrade.

For more information, see the technote:
http://www.symantec.com/docs/TECH211178

Enabling delayed allocation on a small file system sometimes disables the file system (2389318)

When you enable delayed allocation on a small file system, such as around 100 MB, the file system can get disabled. In this case, the following error message ,displays in the system console log:

```
mesg 001: V-2-1: vx_nospace - file_system file system full
(size block extent)
```

Workaround:

Use the `vxtunefs` command to turn off delayed allocation for the file system.

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system nears 100% usage even if other volumes have free space (2438368)

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system is nearing 100% usage even if other volumes in the file system have free space.

Workaround:

After sufficient space is freed from the volume, delayed allocation automatically resumes.

Task blocked messages display in the console for RHEL6 (2560357)

On RHEL6, the kernel occasionally displays messages in the console similar to the following example:

```
INFO: task seq:16957 blocked for more than 120 seconds.
```

These messages display because the task is blocked for a long time on sleep locks. However, the task is not hung and the messages can be safely ignored.

Workaround: You can disable these messages by using the following command:

```
# echo 0 > /proc/sys/kernel/hung_task_timeout_secs
```

Deduplication can fail with error 110 (2591473)

In some cases, data deduplication fails with a message similar to the following example:

```
Saving    Status    Node            Type        Filesystem
-------------------------------------------------------------------
00%       FAILED    node01          MANUAL      /data/fs1
        2011/10/26 01:38:58 End full scan with error
```

In addition, the deduplication log contains an error similar to the following example:

```
2011/10/26 01:35:09 DEDUP_ERROR AddBlock failed. Error = 110
```

These errors indicate that the deduplication process is running low on space and needs more free space to complete.

Workaround:

Make more space available on the file system.

vxresize fails while shrinking a file system with the "blocks are currently in use" error (2437138)

The `vxresize` shrink operation may fail when active I/Os are in progress on the file system and the file system is being shrunk to a size closer to its current usage. You see a message similar to the following example:

```
UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink /dev/vx/rdsk/dg1/vol1 -
blocks are currently in use.
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume
vol1, in diskgroup dg1
```

Workaround:

Rerun the shrink operation after stopping the I/Os.

Possible assertion failure in vx_freeze_block_threads_all() (2244932)

There is a possible assertion failure in the `vx_freeze_block_threads_all`() call when the `pdir_threshold` tunable is set to 1.

Workaround:

There is no workaround for this issue.

Severe impact in read performance (sequential and random) on compressed files compared to uncompressed files (2609152)

The read throughput is highly degraded for compressed files. The difference is seen for sequential I/O and random I/O. For sequential reads, the degrdataion is visbile even when the amount of data read compressed files is one-third of the uncompressed files (compression ratio).

Workaround:

There is no workaround for this issue.

fsppadm operations issued on multi-volume file system fail if there are other mounted file systems with a disk layout Version less than 6 (2909206)

The `fsppadm` command checks all mounted file systems, and if it finds any file systems with a disk layout Version that is less than 6, then it exits with the following error message:

```
# fsppadm assign /dst_vset /tmp/pol_test.xml

UX:vxfs fsppadm: ERROR: V-3-26510: Low level Volume enumeration failure
on / with message  Function not implemented
```

This error occurs because the `fsppadm` command functionality is not supported on a disk layout Version that is less than 6.

Workaround:

There is no workaround for this issue.

System unable to select ext4 from the file system (2691654)

The system is unable to select ext4 from the file system.

Workaround: There is no workaround.

A low memory machine with RHEL 6.2 or higher may panic with General Protection Fault (3046544)

If a machine with low memory (less than 3GB) is installed with RHEL 6.2 or higher system, it may hit the General Protection Fault assert with the following stack trace:

```
machine_kexec at ffffffff8103284b
crash_kexec at ffffffff810ba972
oops_end at ffffffff81501860
die at ffffffff8100f26b
do_general_protection at ffffffff815013f2
general_protection at ffffffff81500bc5
shrink_page_list.clone.0 at ffffffff8112db97
shrink_inactive_list at ffffffff8112e10c
shrink_zone at ffffffff8112ee8f
balance_pgdat at ffffffff811303d9
kswapd at ffffffff81130606
kthread at ffffffff81091df6
kernel_thread at ffffffff8100c14a
```

Workaround:

There is no workaround for this issue.

The replication operation fails (3010202)

On SLES 10 SP3 or SLES 11 SP2, the replication operation fails with the following message:

```
btree.c     1827:       ASSERT(0) failed
```

This issue occurs if you use the `bcopy` function because it is deprecated in these releases.

Workaround:

You can turn off the shared extent optimization on these machines. Contact Symantec support to get the exact commands.

The de-duplication operation may seem to be hung (2753687)

After the de-duplication operation completes, some of the temporary files are not cleaned. Hence the shutdown script fails to kill all the processes and the de-duplication operation may hang.

Workaround:

Use the `kill` command to manually kill the de-duplication operation.

The `fsdedupadm` command does not show the full name of the machine (3015478)

The length of the name for a node is limited to 15 characters. So whenever the name of a machine exceeds this length, it shows only the first 15 characters of name.

Workaround:

There is no workaround for this issue.

Incorrect option for VxFS file system in /etc/fstab (3059189)

When you create a VxFS file system, VOM sets incorrect option 'suid' for it in the `/etc/fstab` file. This would make the operating system unable to startup normally after reboot:

```
# NOTE: When adding or modifying VxFS or VxVM entries, add '_netdev'
# to the mount options to ensure the filesystems are mounted after
# VxVM and VxFS have started.
/dev/vx/dsk/dg3/dg3vol2 /dg3vol2 vxfs suid 0 2
```

Workaround:

Before reboot, manually edit the `/etc/fstab` file and change 'suid' to '_netdev'.

Veritas Volume Manager known issues

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

The vxrecover command does not handle RAID5 volumes correctly (2715124)

The `vxrecover` command calls the recovery process for the top-level volume, which internally takes care of recovering its subvolumes. The `vxrecover` command does not handle RAID5 volumes correctly. The recovery process fails to recover the subvolumes, which remain in the NEEDSYNC state.

Workaround:

Manually recover the RAID5 volumes using the `vxvol` utility, as follows:

```
# vxvol -g diskgroup resync volume
```

The vxcdsconvert utility is supported only on the master node (2616422)

The `vxcdsconvert` utility should be run only from the master node, not from the slave nodes of the cluster.

device.map must be up to date before doing root disk encapsulation (2202047)

If you perform root disk encapsulation while the `device.map` file is not up to date, the `vxdiskadm` command displays the following error:

```
VxVM vxencap INFO V-5-2-6098 Missing entry for root disk <rootdisk name>
in /boot/grub/device.map
```

Workaround:

Before you perform root disk encapsulation, run the the following command to regenerate the device.map file:

```
grub-install --recheck /dev/sdb
```

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

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

Workaround

Reboot the system or run the following command.

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

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

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

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

vxrestored daemon fails to restore disabled paths (1663167)

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

Workaround:

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

To enable the mpt_disable_hotplug_remove tunable

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

   ```
   options mptsas mpt_disable_hotplug_remove=0
   ```

2. Rebuild the initrd image:

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

3. Reboot the system.

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

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

Workaround:

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

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

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

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

Workaround:

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

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

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

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

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

- One or more arrays that provide the shared storage for the cluster are being powered off
- At the same time when the arrays are being powered off, an operation that requires an internal transaction is initiated (such as VxVM configuration commands)

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

Workaround:

To recover from this situation

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

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

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

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

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

Root disk encapsulation is not supported if the root disk (DMP node) has customized name or user-defined name (1603309)

Encapsulation of the root disk fails if it has been assigned a customized name with vxdmpadm(1M) command. If you want to encapsulate the root disk, make sure that you have not assigned a customized name to its corresponding DMP node.

See the vxdmpadm(1M) manual page.

See the "Setting customized names for DMP nodes" section of the *Veritas Storage Foundation Administrator's Guide*.

VxVM starts before OS device scan is done (1635274)

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

Workaround:

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

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

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

```
Cannot setup space
```

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

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

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

Workaround: There is no workaround for this issue.

Co-existence check might fail for CDS disks (2214952)

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

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

In layouts where GPT is used for initialization (typically, when the disk size is currently greater than or had earlier exceeded 1 TB), the AIX co-existence label is

placed at sector 55 and VxVM ID block (also known as HP co-existence label) is placed at sector 64. Consequently, AIX utilities would not be able to recognize a cdsdisk initialized using GPT to be a valid VxVM disk. Symantec is working with IBM and third party OEMs to enhance the co-existence check in these utilities.

Workaround:

There is no workaround for this issue.

The "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set (2354560)

In Cluster Volume Manager environment, "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set. This can happen on master/slave nodes.

Workaround:

Administrator has to run "vxdisk scandisks" or "vxdisk -o alldgs list" followed by "vxdg listclone" to get all the disks containing "clone_disk" or "udid_mismatch" flag on respective host.

The vxsnap print command shows incorrect value for percentage dirty (2360780)

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In SFHA 6.0, if this command is run while the volumes are online and being actively used, the shown %dirty may lag from actual percentage dirty for instant snap data cache object (DCO) volumes. That is, the command output may show less %dirty than actual.

Issues with the disk state on the CVM slave node when vxconfigd is restarted on all nodes (2615680)

When a CVM master node and a slave node have lost storage access, and `vxconfigd` is restarted on all nodes, the disk state on the CVM slave node shows as invalid.

Workaround:

To work around this issue

1. Restore storage connectivity.
2. Deport the disk group.
3. Import the disk group.

During system boot, some VxVM volumes fail to mount (2622979)

During system boot, some VxVM volumes that exist in the `/etc/fstab` file fail to mount with the following error messages:

```
# fsck
Checking all file systems.
  error on stat() /dev/vx/dsk//volume: No such
file or directory
```

The load order of kernel modules in Linux results in the VxFS file system driver loading late in the boot process. Since the driver is not loaded when the `/etc/fstab` file is read by the operating system, file systems of the type vxfs will not mount.

Workaround:

To resolve the failure to mount VxFS file systems at boot, specify additional options in the `/etc/fstab` file. These options allow the filesystems to mount later in the boot process. An example of an entry for a VxFS file system:

```
/dev/vx/dsk/testdg/testvolume /mountpoint  vxfs  _netdev,hotplug  1 1
```

To resolve the issue, the fstab entry for VxVM data volumes should be as per following template:

```
/dev/vx/dsk/testdg/testvol    /testmnt    vxfs   _netdev     0 0
```

Removing an array node from an IBM Storwize V7000 storage system also removes the controller (2816589)

When using an IBM Storwize V7000 storage system, after removing one array node, the corresponding controller is also removed.

Workaround: The following procedure resolves this issue.

To resolve this issue

1 Set the `iotimeout` tunable to 600:

```
# vxdmpadm setattr enclosure encl1 recoveryoption=throttle \
 iotimeout=600
```

2 After you re-add the SAN VC node, run the `vxdctl enable` command for Dynamic Multi-Pathing (DMP) to detect the added paths:

```
# vxdctl enable
```

Diskgroup import of BCV luns using -o updateid and -o useclonedev options is not supported if the diskgroup has mirrored volumes with DCO or has snapshots. (2831658)

VxVM uses guid stored in configuration to uniquely identify all objects. The DCO volume stores the guid of mirrors and snapshots. If the diskgroup is imported with -o updateid and -o useclonedev, it changes the guid of objects in VxVM configuration database and the guids stored in DCO volume are not updated. So the operations involving DCO will not be able to find objects with the stored guid and this could lead to failure of certain operations involving DCO or could lead to unexpected behaviour.

Workaround:

No workaround available.

Hardware paths for operating system paths have changed in DMP 6.0 (2410716)

In DMP 6.0, the hardware paths for operating system paths have changed. After upgrading to DMP 6.0, path attributes are reset to the default values. You must reconfigure any path-level attributes that were defined in the /etc/vx/dmppolicy.info file.

Workaround:

To configure path-level attributes

1. Remove the path entries from the `/etc/vx/dmppolicy.info` file.
2. Reset the path attributes.

Upgrading from Veritas Storage Foundation and High Availability 5.x to 6.0.1 may fail for IBM XIV Series arrays (2715119)

Starting in the Veritas Storage Foundation and High Availability 5.1 SP1 release, the Array Support Library (ASL) for the IBM XIV enclosures converts the LUN Serial Number from Hexadecimal to Decimal. Because of this change, the enclosure names differ from releases prior to the 5.1 SP1 releases. When you upgrade Veritas Storage Foundation and High Availability from a release prior to that release to the current 6.0.1 release, XIV LUNs may go into an error state. Note that the latest RPs on 5.1/5.1SP1 are already modified to use the same logic for enclosure naming.

Workaround:

After the upgrade, run `vxddladm assign names`.

Cannot grow Veritas Volume Manager (VxVM) disk using the vxdisk resize command during Dynamic LUN Expansion operation (2064510)

The following error message is displayed during the Dynamic LUN Expansion operation of a LUN with the SIMPLE format:

```
VxVM vxdisk ERROR V-5-1-8643 Device <device name>: resize failed:
Invalid data in request
```

The `vxdisk resize` command keeps the cylinder size (number of the heads * total number of the sectors per track) constant before and after the resize operation, unless the number of cylinders go beyond 2^16-1 (65535) . Because of the VTOC limitation of storing geometry values only till 2^16 -1, if the number of cylinders increases beyond the limit, `vxdisk resize` increases the cylinder size. If this happens, the private region will overlap with the public region data and corrupt the user data.

As a result of this LUN geometry change, VxVM is unable to complete `vxdisk resize` on simple format disks. VxVM was not designed to handle such geometry changes during Dynamic LUN Expansion operations on simple disks.

Workaround:

The VxVM `vxdisk resize` command behaves differently depending on whether the disk is simple, sliced, or CDS format.

The problem shown above only occurs on simple disk configurations. As a result of this difference in behavior, if the geometry changes during a Dynamic LUN Expansion operation at the LUN level, you can convert the disk to a CDS format disk. Use the `vxcdsconvert` command on the disk. Then you can issue the `vxdisk resize` command.

See http://www.symantec.com/docs/TECH136240 for more information.

After devices that are managed by EMC PowerPath lose access to storage, Veritas Volume Manager commands are delayed (2757198)

In an enviroment which includes devices that are managed by EMC PowerPath, a storage loss causes Veritas Volume Manager commands to be delayed. In the event of storage loss, VxVM sends SCSI inquiry from each LUN path to check the health of path, which are delayed by the presence of EMC PowerPath.

Complete site is detached, if plex detach operation is performed even after site consistency off (2845383)

By design, you cannot detach the last plex of a site on a site consistent volume without detaching the complete site. By default, attempting to detach the last plex causes an error. If you use the force detach option, then the complete site is detached to ensure site consistency. This behavior is seen even if you turn off the site consistent flag if the allsites flag is on.

Performance impact when a large number of disks are reconnected (2802698)

If the storage connectivity is lost to part of the storage, the disk group configuration copy is rebalanced to the disks that have connectivity. For example, if the storage for an entire enclosure is removed from a disk group with muliple enclosures. The rebalancing process takes time, during which time the `vxconfigd` daemon is busy and does not respond to commands.

Plex synchronization is not completed after resuming synchronization on a new master when the original master lost connectivity (2788077)

When you run `vxrecover -o force`, it recovers only one subvolume and it cannot detect that the rest of the volume needs recovery.

When you run the `vxassist mirror` command, you run the `vxplex att`command serially on each subvolume. If the failure happens before you start the `attach`operation (need to mark the concerned plex as the attach operation is in progress), `vxrecover` will not redo the attach operation because it cannot find any record of the attach operation in progress.

Workaround:

Run the following command on each subvolume to manually recover the complete volume:

```
# /usr/lib/vxvm/type/fsgen/vxplex -U fsgen -g diskgroup \
 -o force useopt att volume plex
```

cvm_clus resource goes into faulted state after the resource is manually panicked and rebooted in a 32 node cluster (2278894)

The `cvm_clus` resource goes into faulted state after the resource is manually panicked and rebooted in a 32 node cluster.

Workaround: There is no workaround for this issue.

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

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

Workaround:

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

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.

The SCSI registration keys are not removed even if you stop VCS engine for the second time (3037620)

If you stop VCS engine for the first time, the SCSI registration keys can be removed. But if you stop VCS engine for the second time, the keys are not removed.

Workaround:

There is no workaround for this issue.

Using `vxsnap` to addmir for a volume is very slow and the plex is detached after the operation (3037712)

If you use vxsnap to addmir for a volume on RHEL6.3, the process is very slow and the plex becomes detached after the operation.

Workaround:

There is no workaround for this issue.

Upgrading VxVM package (in-place upgrade) on an encapsulated boot disk does not work on SLES11 (3061521)

On SLES11, after the in-place upgrade on root encapsulated system, the following message is displayed on reboot:

```
Waiting for device /dev/vx/dsk/bootdg/rootvol to appear:
..............could not find /dev/vx/dsk/bootdg/rootvol.
Want me to fall back to <resume device>?(Y/n)
```

Workaround:

For an SLES11 machine with encapsulated root disk, perform the following steps to upgrade VxVM.

To upgrade VxVM:

1. Unroot the encapsulated root disk:

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

2. Upgrade the VxVM.

   You can perform this using the CPI installer or using the `rpm` command.

   ```
   # rpm -Uvh VRTSvxvm-version.rpm
   ```

3. Reboot the system.
4. Re-encapsulate the root disk:

   ```
   # /etc/vx/bin/vxencap -c -g root_diskgroup rootdisk=root_disk
   ```

Nodes with root encapsulated disks fail to restart after you perform an operating system upgrade for SLES11SPx on 6.0.1 systems. (2612301)

Upgrading the kernel or operating system on an encapsulated boot disk does not work on SUSE Linux Enterprise Server (SLES) 11.

Note: This issue though fixed in release 6.0.1 on Red Hat Enterprise Linux (RHEL) 5, 6 and SLES 10. It occurs on SLES 11 because of a program issue.

Workaround:

Perform the following procedure on the system with the encapsulated root disk to upgrade the kernel or operating system.

To upgrade the kernel or operating system on a system with an encapsulated root disk

1 Unroot the encapsulated root disk:

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

2 Upgrade the kernel or the operating system.

You may upgrade the kernel using the following rpm command:

```
# rpm -Uvh Kernel-upgrade_version
```

OR

You may upgrade the operating system using tools like YaST or Zypper.

3 Reboot the system.

4 Re-encapsulate the root disk:

```
# /etc/vx/bin/vxencap -c -g root_diskgroup rootdisk=root_disk
```

Nodes with encapsulated root disks fail to reboot after you upgrade VxVM from version 6.0.3 to 6.0.4. (3300580)

On an SUSE Linux Enterprise Server (SLES) 11 system with encapsulated root disks, a reboot after upgrading Veritas Volume Manager (VxVM) from 6.0.3 to 6.0.4 prompts you to fall back to the resume device. The prompt message is similar to following:

```
could not find /dev/vx/dsk/bootdg/rootvol, want me to fall back to
/dev/sda2? (Y/N)
```

Workaround:

To resolve this issue

1 Select **Y** on the prompt to continue booting up.

2 Backup the existing `/boot/VxVM_initrd.img` file.

3 Regenerate the `VxVM_initrd.img` file using the following command:

```
# /etc/vx/bin/vxinitrd  /boot/VxVM_initrd.img `uname -r`
```

In SUSE Linux Enterprise Server 11 (SLES11) SPx KVM guest, Veritas Volume Manager is not able to discover the disks exported from host as virtio-disk to guest. (3325022)

In SLES11 SPx installed KVM guest, during device discovery, Veritas Volume Manager fails to claim the devices exported from host using virtio-disk interface.

Workaround:

Export the device from host using virtio-lun interface.

Replication known issues

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

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

When VVR is used for replicating shared disk groups in a Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) or Veritas Storage Foundation for Oracle RAC (SFRAC) environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg` commands:

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

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

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

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

Workaround:

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

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

Replication hang when VVR logowner is on CVM slave node (2405943)

When VVR is used for asynchronous replication in shared disk group environment, one of the nodes of the cluster at the primary site is chosen as the logowner. When the logowner node is on a node which is a slave node for the underlying CVM cluster, in the presence of heavy I/O from a node that is not the logowner, it is possible to get into a replication hang. This is due to an internal defect which will be fixed in later releases.

Workaround:

Enable the PreOnline trigger of the RVGLogOwner agent so that the VVR logowner will always reside on the CVM master node. For the detailed procedure, refer to the RVGLogowner agent notes section in the *Veritas Cluster Server Bundled Agents Reference Guide*.

vradmin syncvol command compatibility with IPv6 addresses (2075307)

The `vradmin syncvol` command does not work with the compressed form of IPv6 addresses if the target disk group and volume names are not specified.

Workaround:

In IPv6 environments, if you run the `vradmin syncvol` command and identify the target host using the compressed form of the IPv6 address, then you also need to specify the target disk group and volume names.

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.

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

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

Workaround:

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround:

The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

Running SUSE Linux and using Novell's YaST tool to configure an IPv6 address may result in an error (1679261)

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

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

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

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

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

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

Workaround:

The following procedure resolves this issue.

To resolve this issue

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

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

In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon (1672410, 1672417, 1825031)

Issue: After upgrading VVR to an IPv6-only environment in 6.0 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 commands are running, vradmind may temporarily lose heart beats (2071568, 2275444)

This issue may occasionally occur when you use `vradmin` commands to administer VVR. While the `vradmin` commands run, `vradmind` may temporarily lose heartbeats, and the commands terminate with the following error message:

```
VxVM VVR vradmin ERROR V-5-52-803 Lost connection to host host;
terminating command execution.
```

Workaround:

To resolve this issue

1. Depending on the application I/O workload and network environment, 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` on all the hosts of the RDS to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following on all the hosts of the RDS:

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

vxassist relayout removes the DCM (145413)

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.

vradmin verifydata operation fails when replicating between versions 5.1 and 6.0 (2360713)

When replicating in a cross-version VVR environment consisting of hosts running Storage Foundation 5.1 and hosts running Storage Foundation 6.0, the `vradmin verifydata` command fails with the following error:

`VxVM VVR vxrsync ERROR V-5-52-2222 [from host]: VxVM in.vxrsyncd ERROR V-5-36-2125 Server volume access error during [assign volids] volume path: [/dev/vx/dsk/dg/snapshot_volume] reason: [this could be because a target volume is disabled or an rlink associated with a target volume is not detached during sync operation].`

Workaround:

There are two workarounds for this issue.

- Upgrade the hosts running Storage Foundation 5.1 to Storage Foundation 5.1SP1 or later and re-run the `vradmin verifydata` command.
- Follow the offline verification procedure in the "Verifying the data on the Secondary" section of the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*. This process requires ensuring that the secondary is up-to-date, pausing replication, and running the `vradmin syncrvg` command with the `-verify` option.

Cannot relayout data volumes in an RVG from concat to striped-mirror (2129601)

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 verifydata may report differences in a cross-endian environment (2834424)

When replicating between two nodes in a cross-platform environment, and performing an autosync or replication, the vradmin verifydata command may report differences. This is due to different endianness between the platforms. However, the file system on the secondary node will be consistent and up to date.

The vxrecover command does not automatically recover layered volumes in an RVG (2866299)

The `vxrecover` command calls the recovery process for the top-level volume, which internally takes care of recovering its subvolumes. The `vxrecover` command does not handle layered volumes correctly. The recovery process fails to recover the subvolumes, which remain in the NEEDSYNC state.

Workaround:

Manually recover the layered volumes using the `vxvol` utility, as follows:

```
# vxvol -g diskgroup resync volume
```

RVG monitor script may display command not found messages (1709034)

On VCS hosts with VVR resources configured, the following error message displayed in engine_A.log indicates a script error:

```
/opt/VRTSvcs/bin/RVG/monitor: line 124: {print $6}: command not found
/opt/VRTSvcs/bin/RVG/monitor: line 124: {print $6}: command not found
/opt/VRTSvcs/bin/RVG/monitor: line 124: {print $6}: command not found
```

This may fail online/monitor the bunker RVG resources, when they are configured.

Workaround:

Manually edit the following files to update the script:

```
/opt/VRTSvcs/bin/RVG/monitor
/opt/VRTSvcs/bin/RVG/online
/opt/VRTSvcs/bin/RVG/offline
```

In each file, modify the following line:

```
sys=`LC_ALL=C; export LC_ALL; $hasys -nodeid | $awk '{print $6}'`
```

to

```
sys=`LC_ALL=C; export LC_ALL; $hasys -nodeid | awk '{print $6}'`
```

Veritas Storage Foundation for Databases (SFDB) tools known issues

The following are known issues in this release of Veritas Storage Foundation for Databases (SFDB) tools.

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SFHA. There is no workaround at this point of time.

Database Storage Checkpoint unmount may fail with device busy (2591463)

In some cases, when a database that is cloned using a Database Storage Checkpoint is shut down, an error similar to the following may occur:

```
SFAE Error:0457: Failed to unmount device
/dev/vx/dsk/datadg/datavol:Ckpt_1317707593_rw_1317708154.
Reason: VxFS returned error : umount: /tmp/clonedb/data: device is
busy
```

Workaround:

As an Oracle user, force shut down the clone database if it is up and then retry the unmount operation.

Attempt to use SmartTier commands fails (2332973)

The attempts to run SmartTier commands such as `dbdst_preset_policy` or`dbdst_file_move` fail with the following error:

```
fsppadm: ERROR: V-3-26551: VxFS failure on low level mechanism
with message - Device or resource busy
```

This error occurs if a sub-file SmartTier command such as `dbdst_obj_move` has been previously run on the file system.

There is no workaround for this issue. You cannot use file-based SmartTier and sub-file SmartTier simultaneously.

Attempt to use certain names for tiers results in error (2581390)

If you attempt to use certain names for tiers, the following error message is displayed:

```
SFORA dbdst_classify ERROR V-81-6107 Invalid Classname BALANCE
```

This error occurs because the following names are reserved and are not permitted as tier names for SmartTier:

- BALANCE
- CHECKPOINT
- METADATA

Workaround:

Use a name for SmartTier classes that is not a reserved name.

Clone operation failure might leave clone database in unexpected state (2512664)

If the clone operation fails, it may leave the clone database in an unexpected state. Retrying the clone operation might not work.

Workaround:

If retrying does not work, perform one the following actions depending on the point-in-time copy method you are using:

- For FlashSnap, resync the snapshot and try the clone operation again.
- For FileSnap and Database Storage Checkpoints, destroy the clone and create the clone again.
- For space-optimized snapshots, destroy the snapshot and create a new snapshot.

Contact Symantec support if retrying using the workaround does not succeed.

FlashSnap resync fails if there is an existing space-optimized snapshot (2479901)

If you try a FlashSnap resync operation when there is an existing space-optimized snapshot, the resync operation fails with the following error:

```
Error: VxVM vxdg ERROR V-5-1-4597 vxdg join FS_oradg oradg failed
datavol_snp : Record already exists in disk group
archvol_snp : Record already exists in disk group
```

Workaround:

Destroy the space-optimized snapshot first and then perform the FlashSnap resync operation.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0.4 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.0.4.

When upgrading from SFHA version 5.0 to SFHA 6.0.4 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.

Clone command fails if PFILE entries have their values spread across multiple lines (2844247)

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.

SFDB commands do not work with the ZHS16GBK character set (2715323)

SFDB commands do not work if the character set of the Oracle database is set to ZHS16GBK. This occurs because SFDB commands are not supported with multi-byte character sets except AL32UTF8 and ZHS16GBK is a multi-byte character set.

There is no workaround for this issue.

Clone fails with error "ORA-01513: invalid current time returned by operating system" with Oracle 11.2.0.3 (2804452)

While creating a clone database using any of the point-in-time copy services such as Flashsnap, SOS, Storage Checkpoint, or Filesnap, the clone fails. This problem appears to affect Oracle versions 11.2.0.2 as well as 11.2.0.3.

You might encounter an Oracle error such as the following:

```
/opt/VRTSdbed/bin/vxsfadm -s flashsnap -o clone
-a oracle -r dblxx64-16-v1 --flashsnap_name TEST11 --clone_path
/tmp/testRecoverdb --clone_name clone1
USERNAME:  oragrid
STDOUT:
Retrieving snapshot information ...                              Done
Importing snapshot diskgroups ...                                Done
Mounting snapshot volumes ...                                    Done

ORA-01513: invalid current time returned by operating system
```

This is a known Oracle bug documented in the following Oracle bug IDs:

- Bug 14102418: DATABASE DOESNT START DUE TO ORA-1513
- Bug 14036835: SEEING ORA-01513 INTERMITTENTLY

Workaround:

Retry the cloning operation until it succeeds.

Data population fails after datafile corruption, rollback, and restore of offline checkpoint (2869259)

Sometimes when a datafile gets corrupted below its reservation size, the rollback may not pass and the file may not be rolled back correctly.

There is no workround at this point of time.

Checkpoint clone fails if the `archive log` destination is same as the datafiles destination (2869266)

Checkpoint cloning fails if the `archive log` destination is the same as the datafiles destination. The error is similar to:

```
Use of uninitialized value $path in hash element
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 121.
Use of uninitialized value $path in concatenation (.) or string
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 124.
Use of uninitialized value $path in pattern match (m//)
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 126.

SFDB vxsfadm ERROR V-81-0564 Oracle returned error.

Reason: ORA-02236: invalid file name (DBD ERROR: error possibly near
<*> indicator at char 172 in 'CREATE CONTROLFILE REUSE SET DATABASE
'TClone03' RESETLOGS NOARCHIVELOG
```

Workaround:

For the 6.0.4 release, create distinct archive and datafile mounts for the checkpoint service.

FileSnap detail listing does not display the details of a particular snap (2846382)

FileSnap does not support displaying a detailed listing of a snapshot or clone. FileSnap only supports displaying a summary of all the snapshots or clones. For example, for the CLI `vxsfadm -s filesnap -a oracle --name=snap1 -o list`, a summary listing all the snapshots is displayed, instead of a detailed listing of a particular snapshot.

Workaround:

There is no workaround for this issue.

Checkpoint clone fails in CFS environment if cloned using same checkpoint and same clone name on both nodes (2869268)

The Checkpoint clone of an oracle database fails in a CFS environment, if you create a clone with a clone name and checkpoint name same as another clone up on a different CFS node.

Workaround:

There is no workaround. Create a clone with a different clone name.

Very long off-host cloning times for large number of datafiles (2849540)

When cloning off-host in certain Oracle database configurations, particularly with several hundred datafiles, the cloning can take a very long time, upto an hour or more. This problem does not cause the cloning to fail. The problem applies to all services such as FlashSnap, Space-optimized snapshots, FileSnap, and Checkpoint.

Workaround:

There is no workaround at this point of time.

`sfua_rept_migrate` fails after phased SFRAC upgrade from 5.0MP3RP5 to 6.0.1 (2874322)

Command `sfua_rept_migrate` sometimes gives an error when upgrading to 6.0.1, and fails to unmount the repository volume. The error message is similar to:

```
# ./sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
Unmounting SFUA Sybase ASA repository.
UX:vxfs umount: ERROR: V-3-26388: file system /rep has been mount
locked
SFORA sfua_rept_migrate ERROR V-81-5550 umount /dev/vx/dsk/repdg/repvol
failed.
SFORA sfua_rept_migrate ERROR V-81-9162 Failed to umount repository.
```

Workaround:

The error does not hamper the upgrade. The repository migration works fine, but the old repository volume does not get unmounted. Unmount the mount using the manual option.

For example, use `/opt/VRTS/bin/umount -o mntunlock=VCS /rep`.

For more information, see TECH64812.

vxdbd fails to start after upgrade from 6.0.1 to 6.0.4 MR on linux (2969173)

The `vxdbd` daemon fails to start after manual upgrade from 6.0.1 to 6.0.4 Maintenance Release (MR) on Linux platform

Workaround:

Use the `--nopreun` option for the `rpm` upgrade command to start up `vxdbd` properly after manual upgrade.

For example:

```
$rpm -Uvh --nopreun VRTSdbed
```

The dbdst_show_fs(1M) command may fail with a Perl warning message (3289243)

The `dbdst_show_fs`(1M) command may fail with the following Perl warning message:

```
oracle@testbox:~> dbdst_show_fs -S $ORACLE_SID -m /snap_data11r2 -o volume
perl: warning: Setting locale failed.
perl: warning: Please check that your locale settings:
        LANGUAGE = (unset),
        LC_ALL = "",
        LANG = "en_US.UTF-8"
    are supported and installed on your system.
perl: warning: Falling back to the standard locale ("C").
SFORA dbdst_show_fs ERROR V-81-6209 Repository Empty.
```

Note: This issue is observed only in Linux SLES 10.

Workaround: Use the following command for the default locale settings and retry:

```
oracle@testbox:~> export LC_ALL=C
```

Software limitations

This section covers the software limitations of this release.

See the corresponding Release Notes for a complete list of software limitations related to that component or product.

See “Documentation” on page 88.

Limitations related to VMware virtual environment

Symantec High Availability wizard does not support configuration of parallel service groups

Symantec High Availability wizard does not support configuration of parallel service groups. You can, however, configure parallel groups from the command line. Such parallel service groups are visible in the Symantec High Availability tab. You cannot, however, modify the configuration from the GUI.

Symantec High Availability Configuration Wizard does not support nested mount points [2874775]

Symantec High Availability Configuration Wizard does not discover nested mount points. As a result, you cannot use the wizard to monitor applications that access a database via a nested mount point.

Workaround: Use VCS commands to configure the mount resources for monitoring nested mount points. Ensure that you set appropriate resource dependencies to configure the nested mounts. For more information, see the *Veritas Cluster Server Administrator's Guide*.

Symantec High Availability tab does not provide an option to remove a node from a VCS cluster [2944997]

The **Symantec High Availability** tab does not support removal of a system from a VCS cluster.

Workaround: You can use VCS commands to remove the system from the VCS cluster. You can also unconfigure the entire VCS cluster from the **Symantec High Availability** tab.

VMware fault tolerance does not support adding or removing non-shared disks between virtual machines

VMware fault tolerance does not support adding or removing of non-shared disks between virtual machines. During a failover, disks that contain application data cannot be moved to alternate failover systems. Applications that are being monitored, thus cannot be brought online on the failover systems.

Veritas Volume Manager software limitations

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

LVM volume group 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.

The vxlist command cannot correctly display numbers greater than or equal to 1 EB

The `vxlist` command and all of the other commands that use the same library as the `vxlist` command cannot correctly display numbers greater than or equal to 1 EB.

You are unable to specify units in the tunables file

The CPI displays an error when you set the unit type, such as MB or GB, for tunable parameters in the tunables file that you specify with the `-tunablesfile` *`file`* option. To avoid the error, you must set the unit type manually.

SFHA does not support thin reclamation of space on a linked mirror volume (2729563)

The thin reclamation feature does not support thin reclamation for a linked mirror volume.

Thin reclamation requests are not redirected even when the ioship policy is enabled (2755982)

Reclamation requests fail from nodes that do not have local connectivity to the disks, even when the ioship policy is enabled. Reclamation I/Os are not redirected to another node.

Veritas Operations Manager does not support disk, disk group, and volume state information related to CVM I/O shipping feature (2781126)

The Veritas Operations Manager (VOM) does not support disk, disk group, and volume state information related to the I/O shipping feature introduced in this release of Cluster Volume Manager. New states such as lfailed, lmissing or LDISABLED are introduced when I/O shipping is active because of storage disconnectvity.

DMP does not support devices in the same enclosure that are configured in different modes (2643506)

DMP does not support the configuration where two devices in the same enclosure are configured in different modes. For example, if one device is configured as ALUA and another one is configured as Active/Passive (A/P).

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

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

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, change the default values for the DMP tunable parameters.

Table 1-14 describes the DMP tunable parameters and the new values.

Table 1-14 DMP settings for NetApp storage attached environment

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
```

Snapshot configuration with volumes in shared disk groups and private disk groups is not supported

A snapshot configuration with volumes in the shared disk groups and private disk groups is not a recommended configuration. In this release, this configuration is not supported.

Storage reclamation does not happen on volumes with break-off snapshot (2798523)

In this release, storage reclamation on a volume is prevented when it has a break-off type snapshot. If storage reclamation is allowed on such volumes, it can lead to the following undesired situation. Instant snapshot operations, including `vxsnap refresh` and `vxsnap restore` operations, lead to full synchronization of either the snapshot or the primary volume depending on the operation.

In this release, if the volume has a snapshot, the storage reclamation is silently prevented. The physical storage is not reduced. The reclaim command reports that the reclamation is done on the disks but the actual storage is not reclaimed for volumes with snapshots:

```
# vxdisk -o full reclaim dg1
Reclaiming storage on:
Disk xiv0_617 : Done.
Disk xiv0_616 : Done.
Disk xiv0_618 : Done.
Disk xiv0_612 : Done.
Disk xiv0_613 : Done.
Disk xiv0_614 : Done.
Disk xiv0_615 : Done
```

As shown in the following example output, the storage is not actually reclaimed.

```
# vxdisk -o thin list
DEVICE    SIZE(MB) PHYS_ALLOC(MB) GROUP TYPE
xiv0_612  19313    2101           dg1   thinrclm
xiv0_613  19313    2108           dg1   thinrclm
xiv0_614  19313    35             dg1   thinrclm
xiv0_615  19313    32             dg1   thinrclm
xiv0_616  19313    31             dg1   thinrclm
xiv0_617  19313    31             dg1   thinrclm
xiv0_618  19313    31             dg1   thinrclm
```

Veritas File System software limitations

The following are software limitations in the 6.0.4 release of Veritas Storage Foundation.

Linux I/O Scheduler for Database Workloads

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

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

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

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

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

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

Reboot the system once the appropriate file has been modified.

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

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.

Limitations with delayed allocation for extending writes feature

The following limitations apply to the delayed allocation for extending writes feature:

- In the cases where the file data must be written to disk immediately, delayed allocation is disabled on that file. Examples of such cases include Direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.
- Delayed allocation is not supported on memory mapped files.

- Delayed allocation is not supported with BSD quotas. When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.
- Delayed allocation is not supported for shared mounts in a cluster file system.

FlashBackup in NetBackup 7.1 and prior does not support disk layout Version 8 and 9

The FlashBackup feature of NetBackup 7.1 or prior does not support a VxFS file system with disk layout Version 8 or 9.

Compressed files that are backed up using NetBackup 7.1 or prior become uncompressed when you restore the files

The NetBackup 7.1 release and prior does not support the file compression feature. If you back up compressed files using NetBackup 7.1 or a prior release, the files become uncompressed when you restore the files.

Replication software limitations

The following are replication software limitations in this release of Veritas Storage Foundation and High Availability.

VVR Replication in a shared environment

Currently, replication support is limited to 8-node cluster applications.

VVR 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 6.0 and the prior major releases of Storage Foundation (5.1 and 5.1SP1). Replication between versions is supported for disk group versions 150, 160, and 170 only. Both the Primary and Secondary hosts must be using a supported disk group version.

Softlink access and modification times are not replicated on SLES10 for VFR jobs

When running a file replication job on SLES10, softlink access and modification times are not replicated.

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

The preferred fencing feature gives preference to more weighted or larger subclusters by delaying the smaller subcluster. This smaller subcluster delay is effective only if the initial RACER node in the larger subcluster is able to complete the race. If due to some reason the initial RACER node is not able to complete the race and the VxFEN driver activates the racer re-election algorithm, then the smaller subcluster delay is offset by the time taken for the racer re-election and the less weighted or smaller subcluster could win the race. This limitation though not desirable can be tolerated.

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.

In a cluster using SCSI-3 based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on both the data disks and coordinator disks. In a

cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. 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 coordination 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 coordination 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.

Uninstalling VRTSvxvm causes issues when VxFEN is configured in SCSI3 mode with dmp disk policy (2522069)

When VxFEN is configured in SCSI3 mode with dmp disk policy, the DMP nodes for the coordinator disks can be accessed during system shutdown or fencing arbitration. After uninstalling VRTSvxvm RPM, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm RPM is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

Veritas Storage Foundation for Databases (SFDB) tools software limitations

The following are the SFDB tools software limitations in this release.

Oracle Data Guard in an Oracle RAC environment

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

Upgrading to Oracle 10.2.0.5 is required if using SFDB tools

If you are running Oracle version 10.2.0.4 and upgrading a Storage Foundation product with SFDB tools to 6.0.4, you must upgrade the Oracle binaries and database to version 10.2.0.5, before upgrading to 6.0.4.

Parallel execution of `vxsfadm` is not supported (2515442)

Only one instance of the `vxsfadm` command can be run at a time. Running multiple instances of `vxsfadm` at a time is not supported.

Creating point-in-time copies during database structural changes is not supported (2496178)

SFDB tools do not support creating point-in-time copies while structural changes to the database are in progress, such as adding or dropping tablespaces and adding or dropping data files.

However, once a point-in-time copy is taken, you can create a clone at any time, regardless of the status of the database.

Documentation

Product guides are available in the PDF format on the software media in the `/docs/product_name` directory. Additional documentation is available online.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The latest product documentation is available on the Symantec website.

http://sort.symantec.com/documents

Documentation set

Table 1-15 lists the documentation for Veritas Storage Foundation and High Availability.

Table 1-15 Veritas Storage Foundation and High Availability documentation

Document title	File name
Veritas Storage Foundation and High Availability Release Notes	sfha_notes_604_lin.pdf
Veritas Storage Foundation and High Availability Installation and Configuration Guide	sfha_install_604_lin.pdf

Table 1-16 lists the documents for Veritas Cluster Server.

Table 1-16 Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_604_lin.pdf
Veritas Cluster Server Release Notes	vcs_notes_604_lin.pdf

Table 1-16 Veritas Cluster Server documentation *(continued)*

Title	File name
Veritas Cluster Server Administrator's Guide	vcs_admin_604_lin.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_604_lin.pdf
Veritas Cluster Server Agent Developer's Guide (This document is available online, only.)	vcs_agent_dev_604_unix.pdf
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_604_lin.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_604_lin.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_604_lin.pdf

Table 1-17 lists the documentation for Veritas Storage Foundation.

Table 1-17 Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_604_lin.pdf
Veritas Storage Foundation Installation Guide	sf_install_604_lin.pdf
Veritas Storage Foundation Administrator's Guide	sf_admin_604_lin.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sfhas_oracle_admin_604_unix.pdf
Veritas File System Programmer's Reference Guide (This document is available online, only.)	vxfs_ref_604_lin.pdf

Table 1-18 lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-18 Veritas Storage Foundation and High Availability Solutions products documentation

Document title	File name
Veritas Storage Foundation and High Availability Solutions Solutions Guide	sfhas_solutions_604_lin.pdf

Table 1-18 Veritas Storage Foundation and High Availability Solutions products documentation *(continued)*

Document title	File name
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	sfhas_virtualization_604_lin.pdf
Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide	sfhas_replication_admin_604_lin.pdf

If you use Veritas Operations Manager (VOM) to manage Veritas Storage Foundation and High Availability products, refer to the VOM product documentation at:

http://sort.symantec.com/documents

Manual pages

The manual pages for Veritas Storage Foundation and High Availability Solutions products are installed in the `/opt/VRTS/man` directory.

Set the `MANPATH` environment variable so the `man`(1) command can point to the Veritas Storage Foundation manual pages:

- For the Bourne or Korn shell (`sh or ksh`), enter the following commands:

```
MANPATH=$MANPATH:/opt/VRTS/man
  export MANPATH
```

- For C shell (`csh` or `tcsh`), enter the following command:

```
setenv MANPATH ${MANPATH}:/opt/VRTS/man
```

See the `man`(1) manual page.

Manual pages are divided into sections 1, 1M, 3N, 4, and 4M. Edit the `man`(1) configuration file `/etc/man.config` to view these pages.

To edit the man(1) configuration file

1. If you use the man command to access manual pages, set `LC_ALL` to “C” in your shell to ensure that the pages are displayed correctly.

   ```
   export LC_ALL=C
   ```

   See incident 82099 on the Red Hat Linux support website for more information.

2. Add the following line to `/etc/man.config`:

   ```
   MANPATH /opt/VRTS/man
   ```

   where other man paths are specified in the configuration file.

3. Add new section numbers. Change the line:

   ```
   MANSECT         1:8:2:3:4:5:6:7:9:tcl:n:l:p:o
   ```

   to

   ```
   MANSECT         1:8:2:3:4:5:6:7:9:tcl:n:l:p:o:3n:1m
   ```

The latest manual pages are available online in HTML format on the Symantec website at:

https://sort.symantec.com/documents