Veritas Storage Foundation™ Release Notes

HP-UX

6.0.1

Symantec™

Veritas Storage Foundation™ Release Notes

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

Product version: 6.0.1

Document version: 6.0.1 Rev 0

Legal Notice

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

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

Storage Foundation Release Notes

This document includes the following topics:

- About this document
- Component product release notes
- About Veritas Storage Foundation
- About Symantec Operations Readiness Tools
- Important release information
- Changes introduced in 6.0.1
- No longer supported
- System requirements
- Fixed issues
- Known issues
- Software limitations
- Documentation errata
- Documentation

About this document

This document provides important information about Veritas Storage Foundation (SF) version 6.0.1 for HP-UX 11i v3. Review this entire document before you install or upgrade SF.

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

This is "Document version: 6.0.1 Rev 0" of the *Veritas Storage Foundation 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.

About Veritas Storage Foundation

Veritas Storage Foundation by Symantec includes Veritas File System (VxFS) and Veritas Volume Manager (VxVM.)

Veritas File System is a high performance journaling file system that provides easy management and quick-recovery for applications. Veritas File System delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity.

Veritas Volume Manager removes the physical limitations of disk storage. You can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Veritas Volume Manager also provides easy-to-use, online storage management tools to reduce downtime.

VxFS and VxVM are included in all Veritas Storage Foundation products. If you have purchased a Veritas Storage Foundation product, VxFS and VxVM are

installed and updated as part of that product. Do not install or update them as individual components.

Veritas Storage Foundation includes the dynamic multi-pathing functionality.

The Veritas Replicator option, which replicates data to remote locations over an IP network, can also be licensed with this product.

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

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

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 introduced in 6.0.1

This section lists the changes in Veritas Storage Foundation 6.0.1.

New versioning process for SFHA Solutions products

Symantec made some changes to simplify the versioning process to ensure that customers have a unified experience when it comes to deploying our different products across Storage, Availability, Backup, Archiving and Enterprise Security products. With this change, all the products will have a 3 digit version. In complying with this approach, the current SFHA Solutions release is available as version 6.0.1.

New directory location for the documentation on the software media

The PDF files of the product documentation are now located in the `/docs` directory on the software media. Within the `/docs` directory are subdirectories for each of the bundled products, which contain the documentation specific to that product. The `sfha_solutions` directory contains documentation that applies to all products.

Changes related to installation and upgrades

The product installer includes the following changes in 6.0.1.

Locally-installed installation and uninstallation scripts now include the release version

When you run local scripts (/opt/VRTS/install) to configure Veritas products, the names of the installed scripts now include the release version.

Note: If you install your Veritas product from the install media, continue to run the `installsf` command without including the release version.

To run the script from the installed binaries, run the `installsf<version>` command.

Where *<version>* is the current release version with no periods or spaces.

For example, to configure the 6.0.1 version of your product, run this command:

```
# /opt/VRTS/install/installsf601 -configure
```

VxVM private region backup pre-checks for disk groups prior to upgrade

The installer verifies that recent backups of configuration files of all the disk groups in VxVM private region have been saved in the `/etc/vx/cbr/bk` directory prior to doing an upgrade. If not, a warning message is displayed.

Warning: Backup `/etc/vx/cbr/bk` directory.

Support for tunables file templates

You can use the installer to create a tunables file template. If you start the installer with the `-tunables` option, you see a list of all supported tunables, and the location of the tunables file template.

Additional installation postcheck options

The `postcheck` option has been enhanced to include additional checks.

You can use the installer's post-check option to perform the following checks:

■ General checks for all products.

- Checks for Volume Manager (VM).

- Checks for File System (FS).

- Checks for Cluster File System (CFS).

Changes related to Veritas Storage Foundation

Veritas Storage Foundation (SF) includes the following changes in 6.0.1:

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 6.0.1:

Enhancements to vxassist for controlling storage allocations and managing volume intents

In this release, the `vxassist` command has been enhanced to provide more flexibility and control in volume allocations and intent management.

The following list describes the enhancements:

- A rich set of new predefined disk classes.
  The new disk classes cover comprehensive characteristics of the available storage. These disk properties are automatically discovered. You can use these disk classes to select the required type of storage for allocations.

- Ability to define alias names for predefined disk classes.
  For administrative convenience, you can customize alias names that are shorter or more user-friendly.

- Ability to change the precedence order for the predefined disk classes that are supported for mirror or stripe separation and confinement.
  You can now customize the precedence order for the predefined disk classes that are supported for mirror or stripe separation and confinement. The mirror or stripe operation honors the higher priority disk class specified in the custom precedence order.

- Ability to define new disk classes.
  You can associate user-defined properties to disks that satisfy a particular criterion. This functionality enables you to customize device classification or grouping. You can use these custom disk classes to specify storage selections.

- New clauses for precise disk selection.
  The new `use` and `require` clauses enable you to select storage from well-defined sets of intended disk properties. The `require` type of clauses select disks from an intersection set where all specified properties are met. The `use` type of clauses select disks from a union set where at least one of the specified

properties is met. The `use` and `require` constraints are made persistent by default, for disk group version 180 and onwards.

■ Management commands for the volume intents.
Use the volume intent management commands to manage the `use` and `require` type of persistent intents. You can set, clear, update, and list the `use` and `require` intents for the volume, after the volume is created.

For more information about `vxassist` and these enhancements, see the *Veritas Storage Foundation Administrator's Guide* and the `vxassist`(1M) manual page.

Upgrade for instant snap Data Change Objects (DCOs)

Instant snap Data Change Objects (DCOs), formerly known as version 20 DCOs, support the creation of instant snapshots for VxVM volumes. Starting with release 6.0, the internal format for instant DCOs changed. Upgrade the instant snap DCOS and DCO volumes to ensure compatability with the latest version of VxVM. The upgrade operation can be performed while the volumes are online.

The upgrade operation does not support upgrade from version 0 DCOs.

See the *Veritas Storage Foundation Administrator's Guide* and the `vxsnap`(1M) manual page.

Dynamic Reconfiguration tool

Dynamic Multi-Pathing provides a Dynamic Reconfiguration tool. The Dynamic Reconfiguration tool is an interactive tool to automate dynamic reconfiguration of LUNs or HBAs. Dynamic reconfiguration includes addition, removal or replacement of LUNs, and replacement of certain HBAs, without requiring a reboot. The Dynamic Reconfiguration tool simplifies the process, so that you do not need a complex set of DMP and operating system related commands.

Changes related to Veritas File System

Veritas File System includes the following changes in 6.0.1:

The glmstat command can display GLM cache memory usage information

You can use the `glmstat -M` command to display GLM cache memory usage information.

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

SmartTier can compress or uncompress files

SmartTier can compress or uncompress files during relocation, or can perform in-place compression or uncompression of an entire tier.

See the *Administrator's Guide*.

File compression

You can compress files to reduce the space used, while retaining the accessibility of the files and having the compression be transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files.

See the *Administrator's Guide*.

Changes related to SFDB tools

The following sections describe the changes related to Storage Foundation for Databases (SFDB) tools in 6.0.1.

See "Support for creation of Golden Image snapshots using FlashSnap for Oracle" on page 14.

See "Support for Flashsnap at the VVR Secondary site for Oracle" on page 14.

See "Introduction of the Compression Advisor tool for Oracle" on page 15.

Support for creation of Golden Image snapshots using FlashSnap for Oracle

In this release, the SFDB tools support the creation of Golden Image snapshots using FlashSnap for Oracle databases.

Online mode, third-mirror-break-off type snapshot i.e. online FlashSnap snapshot of a database instance contains all the information needed to create a clone of the database instance. It can act as a template for creating clone database instances. You can thus allocate a FlashSnap snapshot that can be used as a master copy for creating one or more clone instances. The clone instances created from a FlashSnap image, termed as the 'golden image', are incremental copies of the master or the golden image. These depend on the FlashSnap image for their operations.

Support for Flashsnap at the VVR Secondary site for Oracle

In this release, the SFDB tools support Flashsnap operation at the VVR Secondary site for Oracle databases.

Online mode snapshots (i.e. traditional, third-mirror-break-off snapshots) are supported in VVR replication environment. Also, support for more than one secondary site is added. For online mode snapshots in VVR environment, IBC (In-Band Control) messages are used to synchronize activities on the Primary and Secondary sites. Snapshot is initiated from VVR Secondary site.

Introduction of the Compression Advisor tool for Oracle

In this release, the SFDB tools provide the Compression Advisor tool for Oracle databases.

Veritas File System (VxFS) provides the `vxcompress` utility that can be used to compress individual files transparent to the underlying applications. An application reading a compressed file automatically receives the uncompressed data that is uncompressed in memory only; the on-disk part of the data remains compressed. If an application writes to a compressed file, parts of the file are uncompressed on disk.

Compression Advisor provides extended compression functionality for Oracle database files in Oracle single instance and Oracle RAC environments. The Compression Advisor command `sfae_comp_adm` resides in the `/opt/VRTS/bin` directory, and it must be run by the DBA user.

Changes related to replication

Veritas Storage Foundation and High Availability Solutions includes the following changes related to replication in 6.0.1:

VVR CPU utilization improvements with fine granular locking and optimizations

CPU usage is reduced due to VVR lock and code optimization. I/O throughput is improved due to faster I/O processing.

CPU utilization improvements and memory optimizations in VVR compression engine

CPU usage is reduced while compression is enabled. The reduced CPU footprint is achieved by memory pre-allocation optimizations, and changing the compression window size and memory levels to provide optimum compression performance.

VVR replication performance improvements in TCP protocol

Overall improvement of replication throughput due to introducing the following:

- An I/O throttling implementation at the VVR layer to improve network bandwidth usage for TCP. (Not applicable to UDP protocol).

- Per RVG read-back memory pool to avoid contention of memory between the RVGs in the SRL read-back.

- A separate read-back thread to read the data from the SRL. This is disabled by default.

Improved resiliency in case of VVR data volume failure in clustered storage environments using CVM I/O shipping framework

In the event of a data volume failure, there may be some writes to the SRL that do not also write to the data volume due to an I/O failure. To make the data consistent, the writes are flushed to the data volume. In previous releases, there was no mechanism to flush the writes from the node with storage connectivity; to avoid data inconsistency, the data volume was detached cluster wide. Using the I/O shipping framework, in flight I/Os (where the I/O finishes on the SRL but does not write to the data volume) are now shipped to the node with storage connectivity and written to the data volume. As a result, the data volume remains consistent and is available on all nodes that have storage connectivity.

No longer supported

The following features are not supported in this release of SF 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 HP-UX 11i v3 operating systems

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

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

Table 1-1   Supported operating systems

Operating system	Operating system version	Architecture
HP-UX 11i Version 3 March 2011 Operating Environments Update Release or later	HP-UX B.11.31.1103 HP-UX B.11.31.1109 HP-UX B.11.31.1203	PA-RISC Itanium

Storage Foundation for Databases features supported in database environments

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

Table 1-2   SFDB features supported in database environments

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

Fixed issues

This section covers the incidents that are fixed in this release.

Installation and upgrades fixed issues

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

Table 1-3   Fixed issues related to installation and upgrades

Incident	Description
2628469	Base501 packages doesn't get removed on fresh installation with 1109 Fusion.
2626311	Installer checks for VRTSfsadv if you specify -version.
2532432	CPI fails to remove older VRTSperl package after the upgrade.
2526709	DMP-OSN tunable value not get persistence after upgrade from 5.1SP1 to 6.0.
2088827	During product migration the installer overestimates disk space use.

Veritas Storage Foundation fixed issues

Issues fixed for Veritas Storage Foundation (SF) includes issues fixed for Veritas File System and Veritas Volume Manager.

Veritas File System fixed issues

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

Table 1-4   Veritas File System fixed issues

Incident	Description
2781322	VxFS returning error 61493 (VX_EFCLNOSPC) on CFS.
2764861	Uncompress by vxcompress ignores quota limitation.
2753944	The file creation threads can hang.
2735912	The performance of tier relocation using fsppadm enforce is poor when moving a large amount of files.
2715186	System panic spinlock: locker forgot to unlock.
2712392	Threads hung in VxFS.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2703747	CFS failover takes up to 20 minutes due to slow log replay.
2696067	When a getaccess() command is issued on a file which inherits the default Access Control List (ACL) entries from the parent, it shows incorrrect group object permissions.
2684573	The performance of the cfsumount(1M) command for the VRTScavf package is slow when some checkpoints are deleted.
2670022	Duplicate file names can be seen in a directory.
2655788	Using cross-platform data sharing to convert a file system that has more than 32k nlinks does not update the vx_maxlink and maxlink_enable tunables.

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

Incident	Description
2651922	ls -l command on local VxFS file system is running slow and high CPU usage is seen.
2599590	Expanding or shrinking a DLV5 file system using the fsadm(1M)command causes a system panic.
2597347	fsck should not coredump when only one of the device record has been corrupted and the replica is intact.
2583197	Upgrading from disk layout Version 8 to 9 on a file system with partitioned directories and Storage Checkpoints can return with a read-only file system error message.
2566875	The write(2) operation exceeding the quota limit fails with an EDQUOT error (Disc quota exceeded) before the user quota limit is reached.
2559450	Command fsck_vxfs(1m) may core-dump with SEGV_ACCERR error.
2555198	sendfile() does not create DMAPI events for Hierarchical Storage Management(HSM) on VxFS.
2536130	fscdsconv fails to convert FS between specific platforms if FCL is enabled.
2272072	GAB panics the box because VCS engine HAD did not respond. The lobolt wraps around.
2183320	VxFS mmap performance degradation on HP-UX 11.31.
2086902	Spinlock held too long on vxfs spinlock, and there is high contention for it.
1529708	Formatting issue with the output of vxrepquota.

Veritas Volume Manager fixed issues

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

Table 1-5   Veritas Volume Manager fixed issues

Incident	Description
2838059	VVR Secondary panic in vol_rv_update_expected_pos.

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

Incident	Description
2832784	ESX panicked after applying a template file from GUI.
2826958	The pwwn number is not displayed in the output of command `vxdmpadm list dmpnode dmpnodename=`*`dmpnode name`*.
2818840	Enhance the `vxdmpraw` utility to support permission and "root:non-system" ownership to be set and make it persistent.
2815517	The `vxdg adddisk` command should not allow mixing clone & non-clone disks in a disk group.
2794625	Unable to configure ASM to use DMP native block device path.
2792242	I/O hang after performing zone remove/add operations.
2774406	The svol_flush_srl_to_dv_start fails to start.
2771452	IO hung because of hung port deletion.
2763206	The `vxdisk rm` command core dumps when list of disknames is very long.
2756059	Panic in voldco_or_drl_to_pvm when volume started at boot.
2754819	Live deadlock seen during disk group rebuild when the disk group contains cache object.
2743926	DMP `restored` daemon fails to restart during system boot.
2741240	The `vxdg join` transaction failed and did not rollback to the sourcedg.
2739709	Disk group rebuild related issues.
2739601	VVR: repstatus output occasionally reports abnormal timestamp.
2737420	The `vxconfigd` daemon dumps core while onlining of the disk.
2729501	Exclude path not working properly and can cause system hang while coming up after enabling native support.

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

Incident	Description
2710579	Do not write backup labels for CDS disk - irrespective of disk size.
2710147	Node panics in dmp_pr_do_reg during key registration with fencing enabled.
2700792	SEGV in vxconfigd daemon during CVM startup.
2700486	The vradmind daemon coredumps when Primary and Secondary have the same hostname and an active Stats session exists on Primary.
2700086	EMC BCV (NR) established devices are resulting in multiple DMP events messages (paths being disabled/enabled).
2698860	The vxassist mirror command failed for thin LUN because statvfs failed.
2689845	After upgrade, some VxVM disks changed to error status and the disk group import failed.
2688747	Logowner local sequential I/Os starved with heavy I/O load on logclient.
2688308	Do not disable other disk groups when a re-import of a disk group fails during master take-over.
2680604	The vxconfigbackupd daemon does not work correctly with NUM_BK in bk_config.
2680482	Empty vx.* directories are left in the /tmp directory.
2679917	Corrupt space optimized snapshot after a refresh with CVM master switching.
2677016	The vxesd daemon dumps core with SIGILL.
2675538	The vxdisk resize command may cause data corruption.
2674465	Data corruption while adding/removing LUNs.
2666163	A small portion of possible memory leak introduced.
2664825	Disk group import fails when disk contains no valid UDID tag on config copy and config copy is disabled.

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

Incident	Description
2657797	Starting 32TB RAID5 volume fails with V-5-1-10128 Unexpected kernel error in configuration update.
2656803	Race between `vxnetd start` and `stop` operations causes panic.
2652485	Inactive snapshot LUNs cause trespassing.
2647795	Intermittent data corruption after a `vxassist move` operation .
2643634	Message enhancement for a mixed (non-cloned and cloned) disk group import.
2627126	Lots of I/Os and paths are stuck in dmp_delayq and dmp_path_delayq respectively. DMP daemon did not wake up to process them.
2626741	Using `vxassist -o ordered` and `mediatype:hdd` options together do not work as expected.
2626199	The `vxdmpadm list dmpnode` printing incorrect path type.
2621465	When detached disk after connectivity restoration is tried to reattach gives 'Tagid conflict' error.
2620556	IO hung after SRL overflow.
2617277	Need man pages for the `vxautoanalysis` and `vxautoconvert` commands.
2608849	VVR Logowner local I/O starved with heavy I/O load from Logclient.
2580393	Removal of SAN storage cable on any node brings Oracle Application Groups down on all nodes.
2576602	The `vxdg listtag` command should give error message and display correct usage when executed with wrong syntax.
2566174	Null pointer dereference in `volcvm_msg_rel_gslock()`.
2564092	Automate the LUN provisioning (addition) / removal steps using `vxdiskadm`.

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

Incident	Description
2556467	DMP-ASM: disable all paths and reboot host cause `/etc/vx/.vxdmprawdev` records to be lost.
2553729	Status of the EMC Clariion disk changed to "online clone_disk" after upgrade.
2516584	Startup scripts use 'quit' instead of 'exit', causing empty directories in `/tmp`.
2495346	The `vxvmconvert` utility is broken to convert LVM to VxVM:hpdisk for larger configurations.
2495338	Disks with hpdisk format can't be initialized with private region offset other than 128.
2441283	The `vxsnap addmir` command sometimes fails under heavy I/O load.
2427894	Opaque disk support for VIS appliance.
2249445	Develop a tool to get the disk-related attributes like geometry, label, media capacity, partition info etc.
2240056	The `vxdg move` transaction not completing and backups fail.
2227678	The second rlink gets detached and does not connect back when overflowed in a multiple-secondaries environment.
2149922	When importing a disk group, record the event in the `/var/adm/messages` (syslog) file.
1675482	The `vxdg list` *dgname* command gives error 'state=new failed'.

Storage Foundation for Databases (SFDB) tools fixed issues

Table 1-6 describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-6   SFDB tools fixed issues

Fixed issues	Description
2585643	If you provide an incorrect host name with the `-r` option of `vxsfadm`, the command fails with an error message similar to one of the following: `FSM Error: Can't use string ("") as a HASH ref while "strict refs" in use at /opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 776. SFDB vxsfadm ERROR V-81-0609 Repository location is invalid.` The error messages are unclear.
2703881 (2534422)	The FlashSnap validation operation fails with the following error if the mirrors for data volumes and archive log volumes share the same set of disks: `SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.`
2582694 (2580318)	After you have done FlashSnap cloning using a snapplan, any further attempts to create a clone from the same snapplan using the `dbed_vmclonedb` continue to use the original clone SID, rather than the new SID specified using the *new_sid* parameter. This issue is also observed when you resynchronize the snapplan, take a snapshot again without specifying the new clone SID, and then try to clone with the new SID.
2579929	The `sfae_auth_op -o auth_user` command, used for authorizing users, fails with the following error message: `SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username>` The authentication setup might have been run with a strict umask value, which results in the required files and directories being inaccessible to the non-root users.

Veritas Storage Foundation fixed issues in 6.0 RP1

Table 1-7 lists the Veritas Volume Manager (VxVM) issues fixed in this release.

Table 1-7   Veritas Volume Manager fixed issues

Incident	Description
2689104	Data corruption occurs while adding/removing LUNs.

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

Incident	Description

2682534	Starting a 32TB RAID5 volume fails with the following error:
	
	`V-5-1-10128`
	`Unexpected kernel error in configuration update.`

2680605	`vxconfigbackupd` does not work correctly with NUM_BK in bk_config.
2666174	A small memory leak may be seen in vxconfigd, the VxVM configuration daemon, when Serial Split Brain (SSB) error is detected in the import process.
2660157	The `vxtune -r` option prints wrong tunable value.
2647086	The `vxvmconvert` utility does not work while converting LVM to VxVM:hpdisk for larger configurations.
2644185	`vxdmpadm` dumps core in display_dmpnodes_of_redundancy.
2643159	In a mirrored volume, when read error occurs, unwanted messages are displayed on console.
2643156	DG activation may hang due to a bug in activation code path, when memory allocation fails in the kernel.
2643155	After deleting RVG, primary master panics in rv_ibc_freeze_timeout.
2643154	vxtune does not accept tunables correctly in human readable format.
2643151	Disks with the hpdisk format cannot be initialized with private region offset other than 128.
2643142	`vxmake -g <DGNAME> -d <desc-file>` fails with very large configuration due to memory leaks.
2643139	I/O hangs on the primary after SRL flush and while reconnecting the RLINK.
2643138	During CVM reconfiguration, VVR waits for the object iocout to go to zero to start the recovery.
2643137	VxVM displays read I/O error messages when a VxVM `nopriv` disk is defined on a partition slice other then slice 2.
2643134	Failure occurs while validating mirror name interface for linked mirror volume.

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

Incident	Description
2637183	Intermittent data corruption occurs after a `vxassist` move.
2633978	On a thin LUN, the creation of VxVM rootdisk mirror using `vxrootmir`(1M) command fails.
2630111	`vxcdsconvert` fails while converting from the hpdisk format to the cdsdisk format.
2630074	`vxdg destroy` hangs for shared DG, following which all vxcommands hang on master.
2628978	Startup scripts use `quit` instead of `exit`, causing empty directories in `/tmp`.
2626994	`vxdg listtag` should give error message and display correct usage when executed with wrong syntax.
2626746	`vxassist -o ordered` and `mediatype:hdd` options together do not work as expected.
2625766	I/O hang occurs on the master node after storage is removed.
2625762	During initial sync from VVR primary site to VVR secondary site, if there is a cluster reconfiguration, the CVM Master on the VVR secondary site may panic.
2625743	While upgrading diskgroup version, if RLINK is not up-to-date, vxrvg shows error but diskgroup version gets updated.
2625724	`vxvmconvert` fails to convert LVM volume on DMP raw devices to vxvm volume with native support set to on.

2625718	The `vxconfigbackup` script throws the following error:
	
	 `c1062-ucs-bl18-vm8:/usr/lib/vxvm/bin`
	 `# mv: missing destination file operand`

2625709	`vxvmconvert` does not support migration of data for large number of LVM/VG configuration since it always creates private region at static offsets - 128th block.
2624574	Under a heavy I/O load on the logclient node, write I/Os on VVR Primary logowner take a very long time to complete.
2615288	Both sites become detached after data/DCO plex failure at each site, leading to I/O cluster wide outage.

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

Incident	Description
2607793	Disabling all paths and rebooting host, causes all `/etc/vx/.vxdmprawdev` records to be lost.
2605706	The write operation after join fails on volume on slave node which earlier had disks in the lfailed state.
2598525	Memory leaks are reported during nightly test runs.
2589962	Enhance the `vxfmrmap` utility, while deprecating `vxfmrshowmap`, to display DCO map contents and verify against possible state corruptions.

Table 1-8 lists the Veritas File System (VxFS) issues fixed in this release.

Table 1-8   Veritas File System fixed issues

Incident	Description
2655754	Deadlock occurs because of wrong spin lock interrupt level at which delayed allocation list lock is taken.
2683409	Binaries linking the `libvxfssnap.sl` library dump core while reading the inodes.
2670022	Duplicate file names can be seen in a directory.
2566875	A `write`(2) operation exceeding the quota limit fails with an EDQUOT error.
2678096	The `fiostat` command dumps core when the count value is 0.
2655788	The nlink count of PD needs to be calculated while validating directories in the fscdsadm(1M) command.
2653845	When the `fsckptadm`(1M) command with the `-r` and `-R` option is executed, two mutually exclusive options gets executed simultaneously.
2645109	The vxfilesnap ioctl does not flush any dirty buffer that is associated with the source inode.
2645112	A write operation on a regular file that maps to a shared compressed extent results in corruption.
2583197	Upgrade of a file system from version 8 to 9 fails in the presence of partition directories and clones.

Table 1-8   Veritas File System fixed issues *(continued)*

Incident	Description
2613884	Metadata corruption may be seen after recovery.
2624459	Listing of a partitioned directory using the DMAPI does not list all the entries.
2646936	The replication process dumps core when shared extents are present in the source file system.
2552095	The system may panic while re-organizing the file system using the `fsadm`(1M) command.
2536130	The fscdsconv(1M) command which is used to convert corrupted or non-VxFS file systems generates core.
2645435	The following error message is displayed during the execution of the `fsmap`(1M) command: `UX:vxfs fsmap: ERROR: V-3-27313`
2599590	Expanding or shrinking a DLV5 file system using the `fsadm`(1M)command causes a system panic.

Table 1-9 lists the Veritas Storage Foundation Installer issue fixed in this release.

Table 1-9   Veritas Storage Foundation Installer fixed issue

Incident	Description
2628469	Base501 packages do not get removed when performing fresh installation with HP-UX 11i v3 September 2011.

Table 1-10 lists the Veritas Enterprise Administrator (VEA) issues fixed in this release.

Table 1-10   Veritas Enterprise Administrator fixed issues

Incident	Description
2672039	The vxsvc process crashes and dumps core.
2677191	File placement policy operations are not available in the VEA GUI.

Known issues

This section covers the known issues in this release.

Installation known issues

This section describes the known issues during installation and upgrade.

Warning messages may be seen during script-based installation (2615500)

When you install SF using the script-based installer, you may see the following warning message:

```
interpreter "/opt/VRTSperl/bin/perl" not found
```

Workaround: You must install perl to resolve the issue.

To install perl

1   Exit the installer.

2   Install the VRTSperl depot from the product media manually:

```
# cd /dvd_path/depot
# /usr/sbin/swinstall -x enforce_dependencies=false
 -x autoreboot=false -s `pwd` VRTSperl
```

3   Start the installer again.

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 depots 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` depots after the upgrade process completes.

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

If you install or configure SF 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.

Error message seen in swagent.log after removing the 6.0.1 VRTS packages (2324553)

After removing the 6.0.1 VRTS packages and before rebooting the system, you sometimes see the following message in the `swagent.log` file:

```
vxfs mount: V-3-21272: mount option(s) incompatible with file system
/dev/vg00/lvol1
```

This message appears because the VRTS packages are removed and the kernel is not yet loaded.

Workaround: Reboot the system.

Installer installs VRTSfsadv if you specify certain options (2626333)

On the HP-UX Precision Architecture (PA) platform, if you run the installer certain options, such as `-minpkgs`, `-recpkgs`, `-allpkgs`, `-pkginfo`, `-pkgtable`, the installer

installs the `VRTSfsadv` depot in addition to the required Veritas File System (VxFS) depots. This depot is not required by VxFS.

Workaround: There is no workaround for this issue. This issue is harmless.

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.

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.

Some unused packages are not removed after upgrade SFORA from 5.0.1 to SFHA 6.0.1 (2821560)

If you upgrade from 5.0.1 or a previous release, the `VRTSobc33`, `VRTSpbx`, and `VRTSicsco` packages are not uninstalled even if no other package depends on them. You can safely ignore these packages. When you uninstall the product, the installer uninstalls these packages.

Workaround: There is no workaround for this issue. This issue is harmless.

Veritas File System modules may fail to unload if SmartMove is enabled and a break-off snapshot volume has been reattached (2851403)

The Veritas File System modules, vxportal and vxfs, may fail to unload if SmartMove is enabled and a break-off snapshot volume is reattached. Reattaching the snapshot causes an extra reference count to the vxportal module, which causes the module unload to fail.

Workaround:

Manually unload the Veritas Volume Manager modules (vxspec, vxio, vxdmp) before unloading the vxportal module. This decrements the vxportal module's reference count.

Perl module error on completion of SF installation (2879417)

When you install, configure, or uninstall SF, the installer prompts you to optionally upload installation logs to the Symantec Web site. If the installer encounters connectivity problems, you may see an error similar to the following:

```
Status read failed: Connection reset by peer at
<midia_path>/../perl/lib/5.14.2/Net/HTTP/Methods.pm line 269.
```

Workaround:

Ignore this error. It is harmless.

Veritas Storage Foundation known issues

There are no new known issues in this release of Veritas Storage Foundation (SF).

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

After upgrading SF stack from 5.0MP3RP2 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the SFM GUI.

Workaround:

To resolve this known issue

◆ On each manage host where VRTSsfmh 2.1 is installed, run:

   # **/opt/VRTSsfmh/adm/dclisetup.sh -U**

A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880081)

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

Workaround: To see the placement class tags in the VEA GUI, you must tag the volumes prior to constructing the volume set. If you already constructed the

volume set before tagging the volumes, restart `vxsvc` to make the tags visible in the GUI.

Veritas Volume Manager known issues

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

The vxrecover command fails with 'Cannot execute /etc/vx/type/static/vxassist: No such file or directory' (2857827)

In some cases, the `vxrecover` command fails with the following error:

```
 Cannot execute /etc/vx/type/static/vxassist: No such file or directory.
```

This case typically happens if linked volume grow or shrink recovery is triggered as part of the `vxrecover` operation.

Workaround:

To resolve this issue, copy the file from `/usr/sbin/vxassist` to `/etc/vx/type/static/vxassist` .

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

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

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

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

Workaround:

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

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

After you initialize a disk that is under the operating system's native LVM control and not under Veritas Volume Manager (VxVM) control by using the `pvcreate` `path_to_physical_disk` command, the first time that you run the `vxdisk list` `disk_name` command results in a VxVM error message related to get_contents, and the `flags` field is incorrectly populated. However, in the next instantiation of the same command, VxVM does not produce an error and the flags are correctly populated with the LVM tag.

Workaround:

Issue the `vxdisk list` `disk_name` command a second time.

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

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

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

Workaround:

There is no workaround for this issue.

Known Issue related to EFI disk initialization (2585433)

For disks initialized with EFI format using `idisk`, DA record becomes invisible from "vxdisk list" output after executing "vxdisk scandisks".

Workaround:

For devices to be correctly seen with slices in "vxdisk list" output, VxVM needs to flush the cached open and reopen the disk device. Further, VxVM needs to search for this new EFI format on the disk and generate new DA record.

To recover from this issue

◆ To achieve this functionality run following VxVM commands:

```
# vxdisk rm <DANAME>
# vxdctl cacheflush
# vxdisk scandisks
```

After disconnecting and reconnecting the Fibre Channel, DMP is unable to present the device tree (2509636)

On some HP-UX 11i version 3 systems, after disconnecting and reconnecting the Fibre Channel, DMP is unable to present the device tree.

Workaround:

Restart the `vxconfigd` daemon with the following command:

```
# vxconfigd -k
```

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.

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 SF 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.

Recovery and rollback to original configuration may not succeed if the system reboots while the online migration setup is in partial state (2611423)

During online migration from LVM to VxVM volumes, if there is a system reboot when the migration setup is in partial state, that is, the start operation has not completed successfully, then the recover and abort operations might not be able to recover and rollback the configuration.

Workaround: This needs manual intervention for cleanup, depending on the state, to restore the original configuration.

During online migration from LVM to VxVM volumes, LVM sometimes incorrectly reports the remapped LVM device paths as valid LVM volumes

Problem: In a migrated or committed configuration, only the renamed LVM names of the form <lvolname>_vxlv are valid LVM volumes. The original LVM names, in turn, point to target VxVM volumes. However, LVM sometimes incorrectly reports these original LVM device paths pointing to VxVM volumes, as valid LVM volumes.

Do not assume these as LVM volumes or do any operations on them, as it would disrupt the application's access to the target VxVM volumes.

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
    ```

Upgrading from Veritas Storage Foundation 5.x to 6.0.1 may fail for IBM XIV Series arrays (2863512)

Starting in the Veritas Storage Foundation 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 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`.

Continuous trespass loop when a Clariion LUN is mapped to a different host than its snapshot (2761567)

If a Clariion LUN is mapped to a different host than its snapshot, a trespass on one of them could cause a trespass on the other. This behavior could result in a loop for these LUNs, as DMP tries to fail back the LUNs if the primary paths are available.

Workaround

To avoid this issue, turn off the `dmp_monitor_ownership` tunable:

```
# vxdmpadm settune dmp_monitor_ownership=off
```

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
```

In some cases with large LUN setup, the storage disappears after DMP device scan (2828328)

This issue is typically seen on a large LUN setup. In some cases, the storage disappears after the DMP device scan. The DMP device scan is generated with the `vxdisk scandisks` command or the `vxdctl enable` command. Even if the the OS command `ioscan` can discover devices, VxVM/DMP cannot.

Workaround:

Restarting the `vxconfigd` daemon on the affected node may resolve the issue. If that does not work, you must reboot the system.

Diskgroup import of BCV luns using -o updateid and -ouseclonedev 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 -ouseclonedev, 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.

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.

Importing a disk group fails with incorrect error message (2149922)

Importing a disk group using clone disks fails with "wrong usage" or "invalid attribute" error. For example, the following command may show the error.

```
# vxdg -o useclonedev=on import dgname
```

This error message may display if the correct feature licenses are not installed.

Workaround:

Check that the Fast Mirror Resync and Disk Group Split and Join licenses are installed. If not, install the licenses.

Dynamic LUN expansion is not supported for EFI disks in simple or sliced formats (2836798)

Dynamic LUN expansion is not supported for EFI (Extensible Firmware Interface) disks in simple or sliced formats. It may lead to corruption. The recommended format is the Cross-platform Data Sharing (CDS) disk format.

Workaround:

Convert the disk format to CDS using the `vxcdsconvert` utility.

System may not boot from a VxVM root disk on a thin LUN (2753626)

The system may fail to boot from a VxVM root disk on a thin LUN. This is an intermittent issue seen only with thin LUNs. The boot process aborts with the following error:

```
System Console is on the Built-In Serial Interface AF_INET socket/streams output
```

Workaround:

In most cases, rebooting the system resolves the issue.

The vxdmp and other drivers have the incorrect release version. (2878024)

The `vxdmp` and other drivers have the incorrect release version. The version displays as 50.0, as shown in the following output:

```
kcmodule -v vxdmp
Module             vxdmp  (50.0)
Description        VxVM DMP Subsystem
Timestamp          Wed Aug  1 10:17:12 2012 [50195688]
State              static (best state)
State at Next Boot  static (best state)
Capable            static unused
Depends On         interface HPUX_11_31_PERF:1.0
```

Setting the media type for a disk access name fails when the naming scheme is new (2876256)

When the naming scheme is `new`, setting the media type for a disk access name fails:

```
# vxdisk -f -g diskgroup set da_name mediatype=ssd
```

Workaround:

You can use the disk media name for the disk.

```
# vxdisk -f -g diskgroup set dm_name mediatype=ssd
```

Veritas File System known issues

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

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.

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.

Not all partitioned directory entries display after exporting a VxFS file system over an HP-UX NFS server (2623412)

After you export a VxFS file system over an HP-UX NFS server, the file system might not list all of the entries in partitioned directories if accessed by NFS clients. This issue is specific to HP-UX NFS servers and VxFS disk layout Version 8 and later.

Workaround: There is no workaround for this issue.

Performance degradation for buffered writes with delayed allocation turned on (2646933)

With the delayed allocation feature turned on, you might observe a performance degradation for buffered writes.

Workaround: Turn off delayed allocation.

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.

MMAP of sparse files on Veritas File System 5.0 is slower on HP-UX 11.31 than on HP-UX 11.23 (2616622)

The MMAP of sparse files on Veritas File System 5.0 is slower on HP-UX 11.31 than on HP-UX 11.23.

Workaround: There is no workaround for this issue.

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.

A mutex contention in vx_worklist_lk() can use up to 100% of a single CPU (2086902)

A mutex contention in the `vx_worklist_lk`() call can use up to 100% of a single CPU.

Workaround: There is no workaround for this issue.

Deleting a large number of files at the same time drastically increases CPU usage (2129455)

When you delete a large number of files at the same time, the CPU usage drastically increases beyond what you should expect.

Workaround: There is no workaround for this issue.

Replication known issues

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

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

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

For example:

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

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:

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

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.

Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (2478684)

Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (DCM), even if you have enough disk space.

Workaround: Add a LUN to the diskgroup before creating the primary diskgroup.

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.

vradmin repstatus operation may display configuration error after cluster reconfiguration in a CVR environment (2779580)

In a CVR environment, if there is a cluster reconfiguration, the `vradmin repstatus` command may display the following error message:

```
No Primary RVG
```

The `vradmin repstatus` command functions normally on the Primary site.

Workaround: Restart the `vradmind` daemon on both the Primary and Secondary nodes.

I/O hangs on the primary node when running vxrvg snaprestore operation (2762147)

In a CVR environment, if a secondary node is set as the logowner for an RVG, issuing the `vxrvg snaprestore` command on the primary node may result in an I/O hang.

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
```

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 SF. 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.

Clone command fails if PFILE entries have their values spread across multiple lines (1922384)

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.

Frequent occurrence of SFDB remote or privileged command error (2869262)

If you installed a single instance database and try to run SFDB-related commands, then an error similar to the following might occur:

```
$ /opt/VRTSdbed/bin/dbed_update

No repository found for database faildb, creating new one.

SFDB vxsfadm ERROR V-81-0450 A remote or privileged command could not
be executed on host1

Reason: This can be caused by the host being unreachable or the vxdbd
daemon not running on that host.

Action: Verify that the host swpa04 is reachable.  If it is, verify
that the vxdbd daemon is running using the /opt/VRTS/bin/vxdbdctrl
status command, and start it using the /opt/VRTS/bin/vxdbdctrl start
command if it is not running.
```

There is no workaround at this point of time.

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 6.0.1 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.

Swverify error related to VRTSdbed observed after a Phase 2 rolling upgrade of SFRAC 6.0.1 on HP-UX 11.31 (2869263)

Upgrade of the SF or SFRAC stack from 5.x to 6.0.1 could display an swverify warning, as follows:

```
WARNING: Directory "/var/vx/vxdba/locks" should have mode "755" but the
         actual mode is "1755".
WARNING: Directory "/var/vx/vxdba/logs" should have mode "755" but the
         actual mode is "1755".
WARNING: Fileset "VRTSdbed.DBED,l=/,r=6.0.100.000" had file warnings.
```

Workaround: Ignore the warning, or change the directory permissions to 755 for both `/var/vx/vxdba/locks` and `/var/vx/vxbda/logs`.

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.

Relinking ODM after upgrading from 5.0.x

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

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

`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.

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 60.

Veritas Volume Manager software limitations

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

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
```

SF 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.

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).

Veritas File System software limitations

The following are software limitations in the 6.0.1 release of Veritas Storage Foundation.

Recommended limit of number of files in a directory

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

After uninstalling Veritas File System 6.0, a file system with disk layout Version 7 or later cannot be mounted

If you install Veritas File System (VxFS) 5.0 or later, create a file system with disk layout Version 7 or later, and then uninstall VxFS, you are left with the base VxFS release of 4.1. VxFs 4.1 does not recognize disk layout Version 7 or later, and thus you cannot mount the file system.

Workaround: You must reinstall VxFS 5.0 or later to mount a file system that has disk layout Version 7, VxFS 5.1 SP1 or later to mount a file system that has disk layout Version 8, or VxFS 6.0 to mount a file system that has disk layout Version 9.

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.

Data deduplication is not supported on PA architecture

The data deduplication feature is not supported on PA architecture.

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.

Veritas Storage Foundation for Databases (SFDB) tools software limitations

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

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 errata

The following sections cover additions or corrections for Document version: 6.0.1 Rev 0 of the product documentation. These additions or corrections may be included in later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).

See the corresponding Release Notes for documentation errata related to that component or product.

See "Documentation" on page 60.

See "About Symantec Operations Readiness Tools" on page 9.

Veritas Storage Foundation Installation Guide

The following errata applies to the *Veritas Storage Foundation Installation Guide*.

"Disk space requirements" section in the "System requirements" chapter

The -precheck option is only available with the general product installer. It is not available with the installation scripts. Replace the following text:

```
# ./installsf<version> -precheck
```

with this text:

```
#./installer -precheck
```

"Mounting the product disc" section in the "Preparing to install Storage Foundation"

Replace the following text:

```
# mount -F cdfs/dev/dsk/c0t0d0 /dvdrom
```

with this text:

```
# mount -F cdfs /dev/dsk/c0t0d0 /dvdrom
```

There should be space between filesystem and device names.

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 1-11 lists the documentation for Veritas Storage Foundation.

Table 1-11   Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_601_hpux.pdf

Table 1-11   Veritas Storage Foundation documentation *(continued)*

Document title	File name
Veritas Storage Foundation Installation Guide	sf_install_601_hpux.pdf
Veritas Storage Foundation Administrator's Guide	sf_admin_601_hpux.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sfhas_oracle_admin_601_unix.pdf
Veritas File System Programmer's Reference Guide (This document is available online, only.)	vxfs_ref_601_hpux.pdf

Table 1-12 lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-12   Veritas Storage Foundation and High Availability Solutions products documentation

Document title	File name
Veritas Storage Foundation and High Availability Solutions Solutions Guide	sfhas_solutions_601_hpux.pdf
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	sfhas_virtualization_601_hpux.pdf
Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide	sfhas_replication_admin_601_hpux.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.