Veritas Storage Foundation™ and High Availability 6.0.1 Release Notes - Linux

Symantec™

Veritas Storage Foundation and High Availability 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 7

Legal Notice

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization

- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information

- Upgrade assurance that delivers software upgrades

- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis

- Premium service offerings that include Account Management Services

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Veritas Storage Foundation and High Availability Release Notes

This document includes the following topics:

- About this document
- Component product release notes
- About Veritas Storage Foundation High Availability
- About Symantec Operations Readiness Tools
- Important release information
- Changes introduced in SFHA 6.0.1
- No longer supported
- System requirements
- SFHA: Issues fixed in 6.0.1
- Known issues
- Software limitations
- Documentation

About this document

This document provides important information about Veritas Storage Foundation and High Availability (SFHA) version 6.0.1 for Linux. Review this entire document before you install or upgrade SFHA.

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

This is "Document version: 6.0.1 Rev 7" of the *Veritas Storage Foundation and High Availability Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

https://sort.symantec.com/documents

Component product release notes

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

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

`/docs/product_name`

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

This release includes the following component product release notes:

- *Veritas Storage Foundation Release Notes* (6.0.1)
- *Veritas Cluster Server Release Notes* (6.0.1)

About Veritas Storage Foundation High Availability

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

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

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

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

About Symantec Operations Readiness Tools

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

SORT can help you do the following:

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

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

To access SORT, go to:

https://sort.symantec.com

Important release information

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

- For the latest patches available for this release, go to:
  https://sort.symantec.com/

- The hardware compatibility list contains information about supported hardware
  and is updated regularly. For the latest information on supported hardware visit
  the following URL:
  http://www.symantec.com/docs/TECH170013

  Before installing or upgrading Storage Foundation and High Availability Solutions
  products, review the current compatibility list to confirm the compatibility of your
  hardware and software.

Changes introduced in SFHA 6.0.1

This section lists the changes in Veritas Storage Foundation and High Availability
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.

SFHA now detects VCS and SF licenses for SFHA installation and configuration

In 6.0.1, you can install and configure SFHA if you install a pair of valid VCS and
SF keys. The VCS and SF keys may not show when you run `vxkeyless display`.
However, you can still install and use SFHA if you install the SF and VCS keys.

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.

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).

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.

Installer support to configure Coordination Point servers

You can now use the `-configcps` option in the installer to configure CP servers. This functionality to configure CP servers is now integrated with the installer. The `configure_cps.pl` script used earlier to configure CP servers is now deprecated.

You can also configure CP servers by generating response files. You can use the `-responsefile '`*/tmp/sample1.res*`'` option in the installer to configure CP servers.

See the *Installation Guide* for more details.

Changes related to Veritas Storage Foundation and High Availability (SFHA)

Storage Foundation and High Availability (SFHA) 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 *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 *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.

Enhanced support for Fusion-io iodrive and iodrive2 on Linux

In this release, support for Fusion-io iodrive and iodrive2 is enhanced to support TRIM operations. TRIM operations erase unused blocks of data to improve SSD performance. The future writes on the Fusion-io cards are more efficient.

The SFHA components, Veritas File System (VxFS) and Veritas Volume Manager (VxVM), use the TRIM operations to free up the blocks that do not contain valid data. The TRIM cabaility is simliar to thin reclamation , and is performed with the same commands.

SFHA supports the TRIM operation for Fusion-io devices for Red Hat Linux 6.0 (RHEL6) and SUSE Linux Enterprise Server 11 (SLES11).

Changes related to Veritas File System

Veritas File System includes the following changes in 6.0.1:

Support for Thin Reclamation on a Thin Reclaimable LUN and TRIMs for an SSD on Linux

The `fsadm -R` command and the `vxfs_ts_reclaim()` call can perform Thin Reclamation on a Thin Reclaimable LUN and TRIMs for an SSD. In a volume set, the action taken is as per the type of device.

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

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.

The vxfsstat command can display pinned memory counters information

You can use the `vxfsstat -m` command to display pinned memory counters information.

For more information, see the `vxfsstat`(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*.

Changes related to SFDB tools

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

Support for FlashSnap and Database Storage Checkpoint for DB2

In this release, the SFDB tools support FlashSnap (Third-mirror break-off snapshots) and Database Storage Checkpoint operations for DB2 databases.

You can capture an online image of actively changing data at a given instant: a point-in-time copy. You can perform system backup, upgrade and other maintenance tasks on point-in-time copies while providing continuous availability of your critical data. You can also offload processing of the point-in-time copies onto another host.

Database FlashSnap lets you make backup copies of your volumes online and with minimal interruption to users.

Database Storage Checkpoint quickly creates a persistent image of a file system at an exact point in time. It reduces I/O overhead by identifying and maintaining only the file system blocks that have changed since the last Storage Checkpoint or backup via a copy-on-write technique.

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:

Protection of the VFR target file system from accidental writes (on Linux)

The `protected=off|on` option of the `mount_vxfs` command protects the target file system from accidental writes. Modifications to the target file system by anything other than the file replication job may cause replication to fail. The new `protected=off|on` option mounts the file system at the target system as read-write, and only allows the replication daemon to apply updates, thus preventing accidental writes that could cause replication to fail.

Finer granularity of replication at the file and directory level using VFR consistency groups (on Linux)

VxFS supports replication at the file and directory level to a single target using a consistency group. A consistency group consists of an include list and an exclude list, which specify a combination of files and directories within a file system that needs to be replicated as a single consistent unit, and files and directories that should not be replicated, respectively. Both include and exclude lists are optional, and if no path is specified in the include list, the entire file system is replicated. The consistency group must be configured on both the source and the target systems.

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.

Changes in supported Linux virtualization technologies

Veritas Storage Foundation and High Availability (SFHA) Solutions 6.0.1 products support the following virtualization technologies in Linux environments:

- Kernel-based Virtual Machine (KVM) technology for Red Hat Enterprise Linux (RHEL)
- Kernel-based Virtual Machine (KVM) technology for SUSE Linux Enterprise Server (SLES)

SFHA Solutions products provide the following functionality for KVM guest virtual machines:

- Storage visibility
- Storage management

- High availability
- Cluster failover
- Replication support

Table 1-1   SFHA Solutions supported configurations in guest and host for KVM technologies

Objective	Recommended SFHA Solutions product configuration	KVM technology
Storage visibility for KVM guest virtual machines	Dynamic Multi-Pathing (DMP) in the KVM guest virtual machines	RHEL SLES
Storage visibility for KVM hosts	DMP in the KVM hosts	RHEL SLES
Storage management features and replication support for KVM guest virtual machines	Storage Foundation (SF) in the KVM guest virtual machines	RHEL SLES
Advanced storage management features and replication support for KVM hosts	Storage Foundation Cluster File System (SFCFSHA) in the KVM hosts	RHEL SLES
End-to-end storage visibility in KVM hosts and guest virtual machines	DMP in the KVM host and guest virtual machines	RHEL SLES
Storage management features and replication support in the KVM guest virtual machines and storage visibility in in the KVM host	DMP in the KVM host and SF in the KVM guest virtual machines	RHEL SLES
Application monitoring and availability for KVM guest virtual machines	Symantec ApplicationHA in the KVM guest virtual machines	RHEL
Virtual machine monitoring and failover for KVM hosts	Veritas Cluster Server (VCS) in the KVM hosts	RHEL SLES
Application failover for KVM guest virtual machines	VCS in the KVM guest virtual machines	RHEL SLES

Table 1-1   SFHA Solutions supported configurations in guest and host for KVM technologies *(continued)*

Objective	Recommended SFHA Solutions product configuration	KVM technology
Application availability and virtual machine availability	Symantec Application HA in the KVM guest virtual machines and VCS in the KVM host	RHEL
Application failover across KVM guest virtual machines and physical hosts	VCS in KVM guest virtual machines and KVM physical host machines	RHEL SLES

VCS provides virtual to virtual (in-guest) clustering support for the following Linux virtualization environments:

- Red Hat Enterprise Virtualization (RHEV)

- Microsoft Hyper-V

- Oracle Virtual Machine (OVM)

For implementation details:

See the *Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide for Linux*.

For VMware support, see *Veritas Storage Foundation in a VMware ESX Environment*.

Changes to LLT

This release includes the following change to LLT:

Setting the value of peerinact in the `/etc/llttab` file

Symantec recommends not to set the value of peerinact to 0. To achieve the infinite timeout functionality for peerinact, you must set peerinact to a large value. The supported range of value is between 1 through 2147483647.

Changes to I/O fencing

This section covers the new features and changes related to I/O fencing in this release.

Enhancement to the CoordPoint agent

The CoordPoint agent monitors changes to the Coordinator Disk Group constitution, such as when a disk is deleted from the Coordinator Disk Group due to accidental execution of a VxVM administrative command or if the VxVM private region of a disk is corrupted.

The agent performs detailed monitoring on the CoordPoint resource and reports faults. You can tune the frequency of the detailed monitoring by setting the LevelTwoMonitorFreq attribute introduced in this release. For example, if you set this attribute to 5, the agent monitors the Coordinator Disk Group constitution in every fifth monitor cycle.

For more information on the CoordPoint agent, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

For information on configuring the CoordPoint agent using script-based installer and manually configuring the CoordPoint agent to monitor coordinator disks, see the *Veritas Cluster Server Installation Guide*.

For more information on replacing I/O fencing coordinator disks or coordinator diskgroup when the cluster is online, see the *Veritas Cluster Server Administrator's Guide*.

No longer supported

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

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

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

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

- FlashSnap reverse resync

- Checkpoint policy and Checkpoint quotas

- Interactive modes in clone and rollback

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems

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

Table 1-2 shows the supported operating systems for this release.

Table 1-2   Supported operating systems

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 6	Update 1, 2, 3	2.6.32-131.0.15.el6 2.6.32-220.el6 2.6.32-279.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only
Red Hat Enterprise Linux 5	Update 5, 6, 7, 8, 9	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5 2.6.18-308.el5 2.6.18-348.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 11	SP1, SP2	2.6.32.12-0.7.1 3.0.13-0.27.1	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 10	SP4	2.6.16.60-0.85.1	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Linux 6	Update 1, 2, 3	2.6.32-131.0.15.el6 2.6.32-220.el6 2.6.32-279.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Linux 5	Update 5, 6, 7, 8, 9	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5 2.6.18-308.el5 2.6.18-348.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only

* Extended Memory Technology

Note: Only 64-bit operating systems are supported.

Note: For SLES11 SP2 kernel versions later than February, 2012, you need to apply the following Veritas patch: sf-sles11_x86_64-6.0.1.100. This patch is available on the patch download page at https://sort.symantec.com/patch/detail/6732.

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

Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

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

Mandatory patch required for Oracle Bug 4130116

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

For more information, refer to the following TechNote:

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

Storage Foundation for Databases features supported in database environments

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

Table 1-3   SFDB features supported in database environments

Veritas Storage Foundations feature	DB2	Oracle	Oracle RAC	Sybase	Sybase ASE CE
Oracle Disk Manager	No	Yes	Yes	No	No
Cached Oracle Disk Manager	No	Yes	No	No	No
Concurrent I/O	Yes	Yes	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes	Yes	Yes

Table 1-3   SFDB features supported in database environments *(continued)*

Veritas Storage Foundations feature	DB2	Oracle	Oracle RAC	Sybase	Sybase ASE CE
SmartTier	Yes	Yes	Yes	Yes	Yes
Database Storage Checkpoints **Note:** Requires Enterprise license	Yes	Yes	Yes	No	No
Database Flashsnap **Note:** Requires Enterprise license	Yes	Yes	Yes	No	No
SmartTier for Oracle **Note:** Requires Enterprise license	No	Yes	Yes	No	No

Notes:

- SmartTier is an expanded and renamed version of Dynamic Storage Tiering (DST).

- Storage Foundation for Databases (SFDB) tools Database Checkpoints, Database Flashsnap, and SmartTier for Oracle are supported with an Enterprise product license.

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

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

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

Veritas Storage Foundation memory requirements

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

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

SFHA: Issues fixed in 6.0.1

This section covers the incidents that are fixed in SFHA 6.0.1.

Installation and upgrades: issues fixed in 6.0.1

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

Table 1-4   Fixed issues related to installation and upgrades

Incident	Description
2329580	Unable to stop some SFCFSHA processes.
2873102	Perl module error on completion of SFHA installation
2627076	Incorrect server names sometimes display if there is a clock synchronization issue.
2622987	sfmh discovery issue when you upgrade your Veritas product to 6.0.1
2585899	On RHEL, unable to create storage for OCR and Vote disk when using FQDN instead of using only the node name.
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.

Installation and upgrades: Issues fixed in 6.0 RP1

There are no new fixed incidents for installation and upgrades in 6.0 RP1.

Veritas Storage Foundation and High Availability fixed issues

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

See "Veritas File System: issues fixed in 6.0.1" on page 25.

See "Veritas Volume Manager: issues fixed in 6.0.1" on page 28.

Storage Foundation for Databases (SFDB) tools: issues fixed in 6.0.1

Table 1-5 describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-5   SFDB tools fixed issues

Incident	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 File System: issues fixed in 6.0.1

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

Table 1-6   Veritas File System fixed issues

Incident	Description
2764861	Uncompress by vxcompress ignores quota limitation.
2753944	The file creation threads can hang.

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

Incident	Description
2735912	The performance of tier relocation using fsppadm enforce is poor when moving a large amount of files.
2712392	Threads hung in VxFS.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2682550	Access a VxFS file system via NFS could cause system panic on Linux while unmount is in progress.
2674639	The cp(1) command with the –p option may fail on a file system whose File Change Log (FCL) feature is enabled. The following error messages are displayed: cp: setting permissions for 'file_name': Input/output error cp: preserving permissions for 'file_name': No data available.
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.
2651922	ls -l command on local VxFS file system is running slow and high CPU usage is seen.
2597347	fsck should not coredump when only one of the device record has been corrupted and the replica is intact.
2584531	vxfs hangs on ls, du and find.
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.
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.
2086902	Spinlock held too long on vxfs spinlock, and there is high contention for it.
1529708	Formatting issue with the output of vxrepquota.

Veritas File System: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas File System (VxFS) in 6.0 RP1.

Table 1-7   Veritas File System 6.0 RP1 fixed issues

Fixed issues	Description
2679361	Network Customization screen doesn't show any NICs in I18N-level0 environment.
2678096	The fiostat command dumps core when the count value is 0.
2672201	Certain commands get blocked by kernel, causing EACCES(ERRNO = 13).
2672148	vxdelestat (1M) when invoked with -v option goes into infinite loop.
2663750	Abrupt messages are seen in engine log after complete storage failure in cvm resiliency scenario.
2655786	Shared' extents are not transferred as 'shared' by the replication process.
2655754	Deadlock because of wrong spin lock interrupt level at which delayed allocation list lock is taken.
2653845	When the fsckptadm(1M) command with the '-r' and '-R' option is executed, two mutually exclusive options gets executed simultaneously.
2649367	Kernel crashes in vx_fopen because of NULL pointer dereference.
2646936	The replication process dumps core when shared extents are present in the source file system.
2646930	Permission denied errors(EACCES) seen while doing I/O's on nfs shared filesystem.
2645435	The following error message is displayed during the execution of the fsmap(1M) command:'UX:vxfs fsmap: ERROR: V-3-27313'.
2645112	write operation on a regular file mapping to shared compressed extent results in corruption.
2645109	In certain rare cases after a successful execution of vxfilesnap command, if the source file gets deleted in a very short span of time after the filesnap operation, then the destination file can get corrupted and this could also lead to setting of VX_FULLFSCK flag in the super block.
2645108	In certain cases write on a regular file which has shared extent as the last allocated extent can fail with EIO error.
2634483	On RHEL6U1 writing to VxFS /proc hidden interface fails with EINVAL.
2630954	The fsck(1M) command exits during an internal CFS stress reconfiguration testing.

Table 1-7   Veritas File System 6.0 RP1 fixed issues *(continued)*

Fixed issues	Description
2613884	Metadata corruption may be seen after recovery.
2609002	The De-duplication session does not complete.
2599590	Expanding or shrinking a DLV5 file system using the fsadm(1M)command causes a system panic.
2583197	Upgrade of a file system from version 8 to 9 fails in the presence of partition directories and clones.
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.
2389318	Enabling delayed allocation on a small file system sometimes disables the file system.

Veritas Volume Manager: issues fixed in 6.0.1

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

Table 1-8   Veritas Volume Manager fixed issues

Incident	Description
2838059	VVR Secondary panic in vol_rv_update_expected_pos.
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.
2812355	CVS rolling upgrade : vxconfigd hung in join when tried to join upgraded slave node to cluster during upgrade from 5.1sp1rp2 to 6.0sp1 on "sles11sp1-Issue 2".
2794625	Unable to configure ASM to use DMP native block device path.

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

Incident	Description
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.
2751278	The vxconfigd daemon hung on all cluster nodes during `vxsnap` operation.
2751102	Random panics seen in vx_worklist_thr on SLES11 and VxFS.
2747032	Write is taking long time to complete when read/write happen simultaneously.
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.
2726148	System unbootable after `/usr/lib/vxvm/bin/vxupdatelvm` script updates filter in `lvm.conf` file.
2721807	Root disk encapsulation: On SLES11 SP2, machine went to maintenance mode during final reboot after encap.
2711312	Missing symbolic link is created after pulling FC cable on RHEL6.
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.
2709743	Inplace upgrade is not working from 6.0.
2703858	Site failure (storage and all nodes including master node) led to 'configuration daemon not accessible' error on all the sites.

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

Incident	Description
2701654	Phantom DMP disk partition causes panic.
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.
2680482	Empty `vx.*` directories are left in the `/tmp` directory.
2680343	Node panic during cur pri path update in cluster while running I/O shipping.
2679917	Corrupt space optimized snapshot after a refresh with CVM master switching.
2675538	The `vxdisk resize` command may cause data corruption.
2664825	Disk group import fails when disk contains no valid UDID tag on config copy and config copy is disabled.
2660151	The `vxconfigd` daemon is generating a series of LVM header messages for devices (CLONES/replicated devices). Secondary EMC MirrorView LUNS in an error state.
2656803	Race between `vxnetd start` and `stop` operations causes panic.
2652485	Inactive snapshot LUNs cause trespassing.
2648176	Performance difference on Master versus Slave during recovery with Data Change Object (DCO).
2645196	Campus Cluster + Hot Relocation: When a disk failure is detected, the associated disks for that site are detached and ALL disks as marked as RLOC.
2644248	The `vxunroot` command fails as root partition "logvol" mounted on `/var/log`.

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

Incident	Description
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.
2626199	The `vxdmpadm list dmpnode` printing incorrect path type.
2623182	vxvm-boot not cleaning up `/tmp/vx.*` directories whenever system reboot is done for Linux environment.
2620555	I0 hang due to SRL overflow & CVM reconfig.
2612301	Upgrading kernel on encapsulated boot disk does not work on Red Hat Enterprise Linux (RHEL) 5, 6, and SUSE Linux Enterprise Server (SLES) 10.
2607706	Encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path names are not the same.
2580393	Removal of SAN storage cable on any node brings Oracle Application Groups down on all nodes.
2566174	Null pointer dereference in `volcvm_msg_rel_gslock()`.
2564092	Automate the LUN provisioning (addition) / removal steps using `vxdiskadm`.
2553729	Status of the EMC Clariion disk changed to "online clone_disk" after upgrade.
2486301	"VXFS" RPM installation failed.
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.
1675482	The `vxdg list` *dgname* command gives error 'state=new failed'.
1533134	DMP: depreciated SCSI `ioctl` use sg_io type of error.
1190117	vxdisk -f init can overwrite some of the public region contents.
2698035	Tunable values do not change as per the values applied through vxtune.

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

Incident	Description
2682491	The vxvmconvert command displays an error message while converting an LVM volume.

Veritas Volume Manager: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.0 RP1.

Table 1-9   Veritas Volume Manager 6.0 RP1 fixed issues

Fixed issues	Description
2674465	Data Corruption while adding/removing LUNs.
2666163	A small portion of possible memory leak introduced due to addition of enhanced messages.
2660151	vxconfigd is generating a series of LVM header messages for devices (CLONES/replicated devices)Secondary EMC MirrorView LUNS in an error state.
2657797	Starting 32TB RAID5 volume fails with unexpected kernel error in configuration update.
2649958	vxdmpadm dumps core due to null pointer reference.
2647795	Intermittent data corruption after a vxassist move.
2629429	vxunroot does not set original menu.lst and fstab files, SUSE 10.0 NETAPP FAS3000 ALUA SANBOOT.
2627056	vxmake -g <DGNAME> -d <desc-file> fails with very large configuration due to memory leaks.
2626741	Using vxassist -o ordered and mediatype:hdd options together do not work as expected.
2621465	When detached disk after connectivity restoration is tried to reattach gives 'Tagid conflict' error.
2620556	I/O hung after SRL overflow.
2620555	I/O hang due to SRL overflow and CVM reconfig.

Table 1-9   Veritas Volume Manager 6.0 RP1 fixed issues *(continued)*

Fixed issues	Description
2610877	In cluster configuration dg activation can hang due to improper handling of error codes.
2610764	In cluster configuration i/o can hang on master node after storage is removed.
2608849	Logowner local I/O starved with heavy I/O load from Logclient.
2607519	Secondary master panics in case of reconfig during autosync.
2607293	Primary master panic'ed when user deleted frozen RVG.
2605702	Bail out initialising disk with large sector size and its foreign device.
2600863	vxtune doesn't accept tunables correctly in human readable format.
2591321	while upgrading dg version if rlink is not up-to-date the vxrvg command shows error but dg version gets updated.
2590183	write fails on volume on slave node after join which earlier had disks in "lfailed" state.
2576602	vxdg listtag should give error message and display correct usage when executed with wrong syntax.
2575581	vxtune -r option is printing wrong tunable value.
2574752	Support utility vxfmrmap (deprecating vxfmrshowmap) to display DCO map contents and verification against possible state corruptions.
2565569	read/seek i/o errors during init/define of nopriv slice.
2562416	vxconfigbackup throws script errors due to improper handling of arguments.
2556467	disabling all paths and rebooting host causes /etc/vx/.vxdmprawdev record loss.
2535142	New crash was detected on RHEL6.1 during upgrade due to mod unload, possibly of vxspec.
2530698	after "vxdg destroy" hung (for shared DG), all vxcommands hang on master.
2527289	Both sites become detached after data/dco plex failue at each site, leading to i/o cluster wide outage.
2526498	Memory leaks seen in some I/O code path.
2516584	startup scripts use 'quit' instead of 'exit', causing empty directories in /tmp.

Table 1-9   Veritas Volume Manager 6.0 RP1 fixed issues *(continued)*

Fixed issues	Description
2402774	Install Upgrade : After upgrade to 6.0 encapsulated root disk is marked as 'clone_disk'.
2348180	Failure during validating mirror name interface for linked mirror volume.
2176084	Intermittent failure to start ESD on a node.
1765916	Correcting World Writable and unapproved file permissions.

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

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

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

Incident	Description
2708619	If you set the scsi3_disk_policy attribute to dmp, you cannot enable the Veritas fencing module (VxFEN). The VxFEN source code is updated to pick up the dmp device path that contains the full disk name instead of a partition or slice.
2845244	`vxfen` startup script gives error `grep: can't open /etc/vxfen.d/data/cp_uid_db`. The error comes because vxfen startup script tries to read a file that might not be present. This error is typically seen when starting vxfen for the very first time after installation.
2554167	Setting peerinact value to 0 in the `/etc/llttab` file floods the system log file with large number of log messages.

Known issues

This section covers the known issues in this release.

Installation known issues

This section describes the known issues during installation and upgrade.

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

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

Workaround:

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

To unfreeze the service groups manually

1   List all the frozen service groups

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

2   Unfreeze all the frozen service groups:

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

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

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

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

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

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

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

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

Error messages in syslog (1630188)

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

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

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

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

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

```
The partioning on disk /dev/sda is not readable by
The partioning tool parted, which is used to change the
partition table.
You can use the partitions on disk /dev/sda as they are.
You can format them and assign mount points to them, but you
cannot add, edit, resize, or remove partitions from that
disk with this tool.
```

Or

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

```
reiserfs aedd (xennet xenblk)

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

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

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

After a locale change restart the vxconfig daemon (2417547)

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

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

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

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

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

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

1   Encapsulate the root disk.

2   Reinstall the nash utility.

3   Upgrade to the SF 6.0.1 release.

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

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

Workaround:

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

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

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

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

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

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

Workaround: Do one of the following:

■   Kill the start.pl process.

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

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.

After upgrade from VxVM version 6.0, 6.0.1, or 6.0.3 with an encapsulated boot disk, the system fails to boot (2750782)

On Red Hat Enterprise Linux 6 (RHEL6), during the Veritas Volume Manager (VxVM) upgrade from 6.0, 6.0.1, or 6.0.3 to higher versions, the RPM runs the installation scripts of the VxVM higher version first.Then, the RPM runs the uninstallation scripts of the VxVM 6.0, 6.0.1, or 6.0.3 version. Due to a defect in the 6.0, 6.0.1, and 6.0.3 uninstallation scripts, it corrupts the file installed by the higher version. This leads to the boot failure.

Workaround:

1 Unroot the encapsulated root disk.

2 Uninstall `VRTSvxvm` (6.0, 6.0.1, or 6.0.3) RPM.

3 Install `VRTSvxvm` of higher version (above 6.0.3).

The uninstaller does not remove all scripts (2696033)

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

Workaround:

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

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

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

Veritas File System modules fail to unload during uninstall or upgrade if a break-off snapshot volume is created or reattached (2851403)

If a break-off snapshot volume is created or reattached on the system, the Veritas
File System modules, `vxportal` and `vxfs`, may fail to unload during uninstall or
upgrade. The situation occurs if the SmartMove feature is enabled, which is the
default setting. When you use the installer to uninstall or upgrade, you may see a
message similar to the following:

```
Veritas Storage Foundation Shutdown did not complete successfully

vxportal failed to stop on dblxx64-21-v1
vxfs failed to stop on dblxx64-21-v1
```

Workaround:

1   Open a new session and manually unload the modules that failed to unload. Use commands similar to the following:

    # **/sbin/modprobe -r vxportal**
    # **/sbin/modprobe -r vxfs**

2   Because some processes failed to stop, the installer recommends a reboot and asks you if you want to continue.

    Press y to continue to the next phase. You can ignore the reboot requirement.

LLT known issues

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

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

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

Workaround: Do not specify the MAC address of a NIC in the llttab file while configuring LLT if you want to configure a vlan later. If you have already specified the MAC address of a NIC, then delete the MAC address from the llttab file, and update the file before you restart LLT.

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

With each received packet, LLT increments the following variables:

- recvcnt (increment by one for every packet)

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

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

This does not impact the LLT functionality.

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

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This

occurs in some corner cases even though a port is not even registered on the peer node.

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

When LLT is configured over a bonded NIC and that bonded NIC is DOWN with the `ifconfig` command, LLT marks the corresponding link down. When the bonded NIC is UP again using the `ifconfig` command, LLT fails to detect this change and it doesn't mark the link UP.

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

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

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

Workaround: None

GAB known issues

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

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

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

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

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

Workaround: There is no workaround for this issue.

Cluster panics during reconfiguration (2590413)

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

Workaround: There is no workaround for this issue.

I/O fencing known issues

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

After you run the vxfenswap utility the CoordPoint agent may fault (3462738)

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

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

CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Workaround: Start VxFEN again after some time.

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

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

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

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

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

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

Unable to customize the 30-second duration (2551621)

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

Workaround: There is no workaround for this issue.

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

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

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

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

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

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

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

To resolve this issue

1   Rename `cpsadm` to `cpsadmbin`:

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

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

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

3   Change the permissions of the new file to 775:

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

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

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

Workaround:

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

Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later (2824472)

The issue exists because the 5.1SP1 release version does not support separate directories for truststores. But, release version 6.0 and later support separate directories for truststores. So, because of this mismatch in support for truststores, you cannot set up trust between client systems and servers.

Workaround: Set up trust manually between the coordination point server and client systems using the `cpsat` or `vcsat` command. Now, the servers and client systems can communicate in secure mode.

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

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

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

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

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

Workaround: To resolve this issue:

1   Activate the storage domain in RHEV-M.

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

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

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

Workaround:

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

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

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

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

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

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

Workaround: There is no workaround for this issue.

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

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

Workaround:

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

The vxfentsthdw utility fails to launch before you install the VRTSvxfen package (2858190)

Before you install the VRTSvxfen package, the file of `/etc/vxfen.d/script/vxfen_scriptlib.sh` where stores the vxfentsthdw utility does not exist. In this case, the utility bails out.

Workaround:

Besides installing the VRTSvxfen package, run the vxfentsthdw utility directly from the installation DVD.

Veritas Storage Foundation and High Availability known issues

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

NFS issues with VxFS Storage Checkpoints (2027492)

NFS clients mounting VxFS Storage Checkpoints that are NFS-exported by SFHA cluster nodes using a Virtual IP may receive the following error message upon Virtual IP failover:

```
Stale NFS file handle
```

This is a result of major numbers of VxFS Storage Checkpoints not necessarily being the same on all SFHA cluster nodes.

Workaround: The following procedure resolves this issue.

To resolve this issue

◆ You can modify the Options attribute of the Share resource corresponding to the VxFS Storage Checkpoint and add the `fsid` share option to force the `fsid` of an NFS-exported VxFS Storage Checkpoint to remain the same on all cluster nodes.

See the `exports`(5) manual page for more information.

For example:

If `share1` is the VCS Share resource corresponding to an NFS-exported VxFS Storage Checkpoint, then run the following commands:

```
# haconf -makerw
# hares -modify share1 Options "fsid=num"
# haconf -dump -makero
```

where *num* is any 32-bit number that is unique amongst all the exported file systems.

Installer exits upgrade to 5.1 RP1 with Rolling Upgrade error message (1951825, 1997914)

Installer exits upgrade to 5.1 RP1 with Rolling Upgrade error message, if protocol version entries are present in `/etc/gabtab` and `/etc/vxfenmode` files. Installer program may exit with either one of the following error messages during upgrade from 5.1 to 5.1 RP1:

```
SF51 is installed. Rolling upgrade is only supported from 5.1 to
higher version for the products
```

Or

```
To do rolling upgrade, VCS must be running on <node>.
```

Workaround: If the protocol version entries are present in /etc/gabtab and /etc/vxfenmode files, then installer detects it as Rolling Upgrade (RU). If you are not attempting RU, and doing full upgrade to 5.1 RP1, remove the protocol version entries from these two files for installer to proceed with regular upgrade.

In an IPv6 environment, db2icrt and db2idrop commands return a segmentation fault error during instance creation and instance removal (1602444)

When using IBM DB2 db2icrt command to create a DB2 database instance on a pure IPv6 environment, the db2icrt command returns segmentation fault error message. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2icrt -a server -u db2fen1 db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 4700: 26182 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -addfcm -i ${INSTNAME?}
```

The db2idrop command also returns segmentation fault, but the instance is removed successfully after the db2idrop command is issued. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2idrop db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 3599:  7350 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -remove -s DB2_${INSTNAME?} 2> /dev/null

DBI1070I  Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

This issue has been identified as an IBM issue. Once IBM has fixed this issue, then IBM will provide a hotfix for this segmentation problem.

At this time, you can communicate in a dual-stack to avoid the segmentation fault error message until IBM provides a hotfix.

To communicate in a dual-stack environment

◆ Add an IPv6 hostname as an IPv4 loopback address to the /etc/hosts file. For example:

    127.0.0.1 swlx20-v6

    Or

    127.0.0.1 swlx20-v6.punipv6.com

    127.0.0.1 is the IPv4 loopback address.

    swlx20-v6 and swlx20-v6.punipv6.com are the IPv6 hostnames.

Process start-up may hang during configuration using the installer (1678116)

After you have installed a Storage Foundation product, some Veritas Volume Manager processes may hang during the configuration phase.

Workaround: Kill the installation program, and rerun the configuration.

Oracle 11gR1 may not work on pure IPv6 environment [1819585]

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

Tools like `dbca` may hang during database creation.

Workaround: There is no workaround for this, as Oracle 11gR1 does not fully support pure IPv6 environment.

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
```

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

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

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

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

Workaround:

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

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

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

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

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

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

To resolve this issue

1 In the Web GUI, go to **Settings** > **Deployment**.

2 Select **HF020008500-06 hotfix**.

3 Click **Install**.

4 Check the **force** option while reinstalling the hotfix.

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

Server panic after losing connectivity to the voting disk (2787766)

This issue occurs on A/P arrays. If the voting disk loses connectivty to the primary paths, DMP takes some time to analyze the error and fail over the paths. During this time, the `cssd` reports a timeout and panics. When using Oracle ASM over DMP devices, set the `disktimeout` parameter to an appropriate value. This parameter indicates the maximum time allowed for a voting file I/O to complete. If this time is exceeded, the voting disk is marked as offline.

The default of `disktimeout` is 200. If the value of the tunable is less that this value, reset the value to the default value.

Workaround:

To set the `disktimeout` to 200:

```
$CRS_HOME/bin/crsctl set css disktimeout 200 [-force] test
```

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

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

Cascaded failure of nodes with ioship enabled may cause the vxconfigd daemon to hang (2865771)

In a shared disk group environment with ioship enabled, the `vxconfigd` daemon may hang in certain cases. When the I/O is initiated from the slave node that has lost connectivity to the disks locally, the I/O is shipped to other nodes. If the node processing the shipped I/O also leaves the cluster shortly after the first node, and tries to rejoin the cluster as a slave, the cascaded failures may cause the `vxconfigd` daemon to hang.

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

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

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

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

vxrestored daemon fails to restore disabled paths (1663167)

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

Workaround:

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

To enable the mpt_disable_hotplug_remove tunable

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

    options mptsas mpt_disable_hotplug_remove=0

2   Rebuild the initrd image:

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

3   Reboot the system.

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

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

Workaround:

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

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

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

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

Workaround:

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

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

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

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

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

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

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

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

Workaround:

To recover from this situation

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

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

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

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

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

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

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

See the vxdmpadm(1M) manual page.

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

VxVM starts before OS device scan is done (1635274)

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

Workaround:

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

I/O fails on some paths after array connectivity is restored, due to high restore daemon interval (2091619)

If a path loses connectivity to the array, the path is marked as suspected to fail and hence is not used for I/O. After the connectivity is restored, the restore daemon detects that the path is restored when the restore daemon probes the paths. The restore daemon makes the path available for I/O. The restore daemon probes the paths at the interval set with the tunable parameter dmp_restore_interval. If you set the dmp_restore_interval parameter to a high value, the paths are not available for I/O until the next interval.

Changes in enclosure attributes are not persistent after an upgrade to VxVM 6.0.1 (2082414)

The Veritas Volume Manager (VxVM) 6.0.1 includes several array names that differ from the array names in releases prior to release 5.1SP1. Therefore, if you upgrade from a previous release to VxVM 6.0.1, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0.1. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-11 shows the Hitachi arrays that have new array names.

Table 1-11 Hitachi arrays with new array names

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

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0.1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

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

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

Workaround:

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

DMP marks the subpaths as DISABLED while these subpaths are accessible from OS level (2037222)

For iSCSI devices on SLES 10 SP3, the DMP tunable parameter dmp_fast_recovery needs to be turned off.

```
# vxdmpadm settune dmp_fast_recovery=off
```

DMP panics if a DDL device discovery is initiated immediately after loss of connectivity to the storage (2040929)

When using EMC Powerpath with VxVM 5.1SP1 on SLES11, set the fast_io_fail_tmo on the HBA port to any non-zero value that is less than the dev_loss_tmo value so as to avoid a panic in case a DDL device discovery is initiated by the vxdisk scandisks command or the vxdctl enable command immediately after loss of connectivity to the storage.

Failback to primary paths does not occur if the node that initiated the failover leaves the cluster (1856723)

When CVM is configured on non-A/A storage, if a node loses access to the storage through all the primary paths, then all the nodes in the cluster switches to the secondary paths. If the node which raised the protocol leaves the cluster and if all the rest of the nodes in the cluster are seeing the primary paths as healthy, then failback to primary paths never happens.

Issues if the storage connectivity to data disks is lost on a CVM slave node while vxconfigd was not running on the node (2562889)

If storage connectivity to data disks is lost on a CVM slave node while vxconfigd was not running on the node, this may result in following issues when vxconfigd comes up on this node:

■ The shared disk groups on the disconnected storage are marked as dgdisabled on the slave node only.

■ The shared disk groups are available to rest of the cluster nodes but no transactions, such as VxVM configuration changes, are possible on any shared disk group.

- Attempts to deport such shared disk groups will fail.

Work-arounds:

Use one of the following work-arounds:

- Remove the faulty slave node out of CVM cluster, restore storage connectivity, and rejoin the node to the cluster.
- Restart `vxconfigd` on the CVM master node.

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

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

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

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

Workaround:

To work around this issue

1   Restore storage connectivity.

2   Deport the disk group.

3   Import the disk group.

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

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

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

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

Workaround:

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

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

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

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

Unable to upgrade the kernel on an encapsulated boot disk on SLES 11 (2612301)

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

Workaround: Perform the following procedure on the system with the encapsulated root disk to upgrade the kernel.

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

1   Unroot the encapsulated root disk:

    # **/etc/vx/bin/vxunroot**

2   Upgrade the kernel:

    # **rpm -Uvh Kernel-*upgrade_version***

3   Reboot the system.

4   Re-encapsulated the root disk:

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

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 and High Availability 5.x to 6.0.1 may fail for IBM XIV Series arrays (2715119)

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

Workaround:

After the upgrade, run `vxddladm assign names`.

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

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

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

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

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

Workaround:

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

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

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

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
```

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

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

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

Workaround:

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

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

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

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

Workaround:

No workaround available.

A master node is not capable of doing recovery if it cannot access the disks belonging to any of the plexes of a volume (2764153)

A master node with missing disks is not capable of doing recovery, as it does not have access to the disks belonging to any of the plexes of a volume.

Workaround:

If other nodes have access to the storage, they can do the recovery. Switch the master role to some other node with better storage connectivity.

CVM fails to start if the first node joining the cluster has no connectivity to the storage (2787713)

If the first node joining the cluster has no connectivity to disks, the import of shared disk groups fails. Other nodes that join the cluster later assume that the auto-import of disk groups is already done as part of the existing cluster processing.

Workaround:

Perform a master switch to the node that has connectivity to the disks. Then import the disk groups manually.

Issue with a configuration with large number of disks when the joining node is missing disks (2869514)

In a configuration with large number of disks (more than 500) where the joining node is missing a few disks (for example. 100 disks), the node join time takes a long time. The joining node attempts to online all the disks as it searches for the missing disks on the node. When the disks are not found the REMOTE LMISSING disks are created on the joining node to complete the join process. This process is found to take time and in such cases the VCS resource online process can timeout.

Workaround:

- Connect the missing disks on the joining node.
- If the intention is to join with missing disks, the VCS timeout needs to be increased.

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.

CVMVolDg agent may fail to deport CVM disk group

The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the `CVMDeportOnOffline` attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.

Workaround: If multiple CVMVolDg resources are configured for a shared disk group, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

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

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

Workaround: There is no workaround for this issue.

Veritas File System known issues

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

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

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

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

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

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

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

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

Workaround:

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

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

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

Workaround:

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

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

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

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

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

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

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

Deduplication can fail with error 110 (2591473)

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

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

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

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

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

Workaround:

Make more space available on the file system.

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

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

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

Workaround:

Rerun the shrink operation after stopping the I/Os.

Possible assertion failure in vx_freeze_block_threads_all() (2244932)

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

Workaround:

There is no workaround for this issue.

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.

fsppadm operations issued on multi-volume file system fail if there are other mounted file systems with a disk layout Version less than 6 (2909206, 2909203)

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

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

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

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

Workaround:

There is no workaround for this issue.

Replication known issues

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

vradmin syncvol command compatibility with IPv6 addresses (2075307)

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

Workaround:

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

RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2054804)

The RVGPrimary agent initiated operation to start replication between the original Primary and the bunker fails during failback – when migrating back to the original Primary after disaster recovery – with the error message:

```
VxVM VVR vxrlink ERROR V-5-1-5282 Error getting information from
remote host. Internal Error.
```

The issue applies to global clustering with a bunker configuration, where the bunker replication is configured using storage protocol. It occurs when the Primary comes back even before the bunker disk group is imported on the bunker host to initialize the bunker replay by the RVGPrimary agent in the Secondary cluster.

Workaround:

To resolve this issue

1   Before failback, make sure that bunker replay is either completed or aborted.

2   After failback, deport and import the bunker disk group on the original Primary.

3   Try the start replication operation from outside of VCS control.

Bunker replay did not occur when the Application Service Group was configured on some of the systems in the Primary cluster, and ClusterFailoverPolicy is set to "AUTO" (2047724)

The time that it takes for a global cluster to fail over an application service group can sometimes be smaller than the time that it takes for VVR to detect the configuration change associated with the primary fault. This can occur in a bunkered, globally clustered configuration when the value of the `ClusterFailoverPolicy` attribute is `Auto` and the `AppGroup` is configured on a subset of nodes of the primary cluster.

This causes the RVGPrimary online at the failover site to fail. The following messages appear in the VCS engine log:

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

Workaround:

To resolve this issue

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

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

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

Workaround:

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround:

The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

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

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

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

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

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

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

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

Workaround:

The following procedure resolves this issue.

To resolve this issue

1   Edit the `/etc/hosts` file to specify the correct IPv6 address.

2   Restart the `vradmind` daemon on all VVR hosts:

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

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

Issue: After upgrading VVR to an IPv6-only environment in 6.0 release, vradmin commands may not work when a colon is specified in the RVG, data volume(s) and/or SRL name. It is also possible that after upgrading VVR to an IPv6-only environment, vradmin createpri may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

Workaround:

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

While vradmin commands are running, vradmind may temporarily lose heart beats (2071568, 2275444)

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

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

Workaround:

To resolve this issue

1 Depending on the application I/O workload and network environment, uncomment and increase the value of the `IPM_HEARTBEAT_TIMEOUT` variable in the `/etc/vx/vras/vras_env` on all the hosts of the RDS to a higher value. The following example increases the timeout value to 120 seconds.

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

2 Restart `vradmind` on all the hosts of the RDS to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following on all the hosts of the RDS:

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

vxassist relayout removes the DCM (145413)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

1 Pause or stop the applications.

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

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

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

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

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

Workaround:

There are two workarounds for this issue.

- Upgrade the hosts running Storage Foundation 5.1 to Storage Foundation 5.1SP1 or later and re-run the `vradmin verifydata` command.

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

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

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

Workaround:

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

1 Pause or stop the applications.

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

    # **vxrlink -g** *diskgroup* **status** *rlink*

3 Stop the affected RVG. Enter the following:

    # **vxrvg -g** *diskgroup* **stop** *rvg*

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

    # **vxvol -g** *diskgroup* **dis** *vol*

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

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

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

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

7 Start the RVG. Enter the following:

    # **vxrvg -g** *diskgroup* **start** *rvg*

8 Resume or start the applications.

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

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

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
```

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

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

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

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

Workaround:

Manually edit the following files to update the script:

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

In each file, modify the following line:

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

to

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

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

SFDB commands do not work in IPV6 environment (2619958)

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

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

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

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

Workaround:

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

Attempt to use SmartTier commands fails (2332973)

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

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

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

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

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

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

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

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

- BALANCE

- CHECKPOINT

- METADATA

Workaround:

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

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

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

Workaround:

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

- For FlashSnap, resync the snapshot and try the clone operation again.

- For FileSnap and Database Storage Checkpoints, destroy the clone and create the clone again.

- For space-optimized snapshots, destroy the snapshot and create a new snapshot.

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

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

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

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

Workaround:

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

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0.1 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.0.1.

When upgrading from SFHA version 5.0 to SFHA 6.0.1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

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

Workaround

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

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

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

Workaround

There is no workaround for this issue.

Clone command errors in a Data Guard environment using the MEMORY_TARGET feature for Oracle 11g (1824713)

The `dbed_vmclonedb` command displays errors when attempting to take a clone on a STANDBY database in a dataguard environment when you are using the MEMORY_TARGET feature for Oracle 11g.

When you attempt to take a clone of a STANDBY database, the `dbed_vmclonedb` displays the following error messages:

```
Retrieving snapshot information ...                          Done
Importing snapshot diskgroups ...                           Done
Mounting snapshot volumes ...                               Done
Preparing parameter file for clone database ...            Done
Mounting clone database ...
ORA-00845: MEMORY_TARGET not supported on this system


SFDB vxsfadm ERROR V-81-0612 Script
/opt/VRTSdbed/applications/oracle/flashsnap/pre_preclone.pl failed.
```

This is Oracle 11g-specific issue known regarding the MEMORY_TARGET feature, and the issue has existed since the Oracle 11gr1 release. The MEMORY_TARGET feature requires the `/dev/shm` file system to be mounted and to have at least 1,660,944,384 bytes of available space. The issue occurs if the `/dev/shm` file system is not mounted or if the file system is mounted but has available space that is less than the required minimum size.

Workaround:

To avoid the issue, remount the `/dev/shm` file system with sufficient available space.

To remount the /dev/shm file system with sufficient available space

1   Shut down the database.

2   Unmount the `/dev/shm` file system:

    # **umount /dev/shm**

3   Mount the `/dev/shm` file system with the following options:

    # **mount -t tmpfs shmfs -o size=4096m /dev/shm**

4   Start the database.

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 [1985055]

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

Workaround: Set MonitorOption attribute for Oracle resource to 0.

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

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

There is no workaround for this issue.

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

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

You might encounter an Oracle error such as the following:

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

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

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

- Bug 14102418: DATABASE DOESNT START DUE TO ORA-1513

- Bug 14036835: SEEING ORA-01513 INTERMITTENTLY

Workaround:

Retry the cloning operation until it succeeds.

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.

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

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

Workaround:

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

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

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

Workaround:

There is no workaround at this point of time.

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

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

```
# ./sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
```

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

Workaround:

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

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

For more information, see TECH64812.

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 92.

Veritas Volume Manager software limitations

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

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

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

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

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

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

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

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

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

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

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

DMP settings for NetApp storage attached environment

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

Table 1-12 describes the DMP tunable parameters and the new values.

Table 1-12   DMP settings for NetApp storage attached environment

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

The change is persistent across reboots.

To change the tunable parameters

1   Issue the following commands:

    # **vxdmpadm settune dmp_restore_interval=60**

    # **vxdmpadm settune dmp_path_age=120**

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

    # **vxdmpadm gettune dmp_restore_interval**

    # **vxdmpadm gettune dmp_path_age**

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

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

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

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

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

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

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

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

Veritas File System software limitations

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

Linux I/O Scheduler for Database Workloads

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

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

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

Configuration File	Architecture and Distribution
/boot/grub/menu.lst	RHEL5 x86_64, RHEL6 x86_64, SLES10 x86_64, and SLES11 x86_64

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

For example, for RHEL5, change:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2
    initrd /boot/initrd-2.6.18-128.el5.img
```

To:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2 \
```

```
    elevator=deadline
    initrd /boot/initrd-2.6.18-128.el5.img
```

For RHEL6, change:

```
title RHEL6
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.32-71.el6 ro root=/dev/sdb2
    initrd /boot/initrd-2.6.32-71.el6.img
```

To:

```
title RHEL6
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.32-71.el6 ro root=/dev/sdb2 \
    elevator=deadline
    initrd /boot/initrd-2.6.32-71.el6.img
```

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

Reboot the system once the appropriate file has been modified.

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

Recommended limit of number of files in a directory

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

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.

Limitations with delayed allocation for extending writes feature

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

■ In the cases where the file data must be written to disk immediately, delayed allocation is disabled on that file. Examples of such cases include Direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.

- Delayed allocation is not supported on memory mapped files.

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

- Delayed allocation is not supported for shared mounts in a cluster file system.

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

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

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

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

Replication software limitations

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

VVR Replication in a shared environment

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

VVR IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.

- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.

- A replication configuration in which an IPv6 address is specified for the `local_host` attribute of a primary RLINK and an IPv4 address is specified for the `remote_host` attribute of the same RLINK.

- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.

- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

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

Softlink access and modification times are not replicated on RHEL5 and SLES10 for VFR jobs

When running a file replication job on RHEL5 and SLES10, softlink access and modification times are not replicated.

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

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

Stopping systems in clusters with I/O fencing configured

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

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

cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordination points and data disks to prevent possible difficulties with subsequent cluster startup.

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

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

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

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

Veritas Storage Foundation for Databases (SFDB) tools software limitations

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

Oracle Data Guard in an Oracle RAC environment

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

Upgrading to Oracle 10.2.0.5 is required if using SFDB tools

If you are running Oracle version 10.2.0.4 and upgrading a Storage Foundation product with SFDB tools to 6.0.1, you must upgrade the Oracle binaries and database to version 10.2.0.5, before upgrading to 6.0.1.

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.

Limitations related to installation

This section covers installation-related limitations.

Limitations related to rolling upgrade

Rolling upgrade with responsefile to 6.0.1 is not supported.

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 1-13 lists the documentation for Veritas Storage Foundation and High Availability.

Table 1-13   Veritas Storage Foundation and High Availability documentation

Document title	File name
Veritas Storage Foundation and High Availability Release Notes	sfha_notes_601_lin.pdf
Veritas Storage Foundation and High Availability Installation and Configuration Guide	sfha_install_601_lin.pdf

Table 1-14 lists the documents for Veritas Cluster Server.

Table 1-14   Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_604_lin.pdf
Veritas Cluster Server Release Notes	vcs_notes_604_lin.pdf
Veritas Cluster Server Administrator's Guide	vcs_admin_601_lin.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_601_lin.pdf
Veritas Cluster Server Agent Developer's Guide (This document is available online, only.)	vcs_agent_dev_601_unix.pdf
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_601_lin.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_601_lin.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_601_lin.pdf

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

Table 1-15   Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_601_lin.pdf
Veritas Storage Foundation Installation Guide	sf_install_601_lin.pdf
Veritas Storage Foundation Administrator's Guide	sf_admin_601_lin.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_lin.pdf

Table 1-16 lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-16   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_lin.pdf
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	sfhas_virtualization_601_lin.pdf
Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide	sfhas_replication_admin_601_lin.pdf

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

http://sort.symantec.com/documents

Manual pages

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

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

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

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

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

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

See the man(1) manual page.

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

To edit the man(1) configuration file

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

    ```
    export LC_ALL=C
    ```

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

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

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

    where other man paths are specified in the configuration file.

3   Add new section numbers. Change the line:

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

    to

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

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

https://sort.symantec.com/documents