Veritas Storage Foundation™ and High Availability Release Notes

Linux

6.0

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

Document version: 6.0.5

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

This document provides important information about Veritas Storage Foundation and High Availability (SFHA) version 6.0 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: 6.0.5 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:

/product_name/docs

Symantec recommends copying the docs directory on the software media that contains the product guides to the /opt/VRTS directory on your system.

This release includes the following component product release notes:

■ *Veritas Storage Foundation Release Notes* (6.0)

■ *Veritas Cluster Server Release Notes* (6.0)

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:
  http://sort.symantec.com/

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

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

Changes introduced in 6.0

This section lists the changes in Veritas Storage Foundation and High Availability
6.0.

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

Storage Foundation and High Availability (SFHA) includes the new features and
changes introduced in 6.0 of the underlying products.

Changes related to Veritas File System

Veritas File System includes the following changes in 6.0:

Data deduplication

You can run post-process periodic deduplication in a file system, which eliminates
duplicate data without any continuous cost. This feature requires an Enterprise
license.

See the *Administrator's Guide*.

File compression

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

See the *Administrator's Guide*.

Default disk layout Version is now 9

In this release, disk layout Version 9 is now the default version, which enables support for the following features:

■ File compression

■ Data deduplication

■ File replication

See the *Administrator's Guide*.

Multi-threaded Thin Reclamation

You can perform multi-threaded Thin Reclamation operations for improved performance.

See the `fsadm_vxfs`(1M) and `vxfs_ts_reclaim`(3) manual pages.

Storage Checkpoints

The following changes were made to Storage Checkpoints:

■ You can tune Veritas File System (VxFS) file systems to create removable Storage Checkpoints by default.
See the `vxtunefs`(1M) manual page.

■ VxFS now attempts to remove removable Storage Checkpoints if the file system does not have enough space instead of failing the operation.

■ Storage Checkpoints have improved visibility to the file system. With the `ckptautomnt` mount option, all Storage Checkpoints are made accessible automatically through a directory in the root directory of the file system that has the special name `.checkpoint`, which does not appear in directory listings. Inside this directory is a directory for each Storage Checkpoint in the file system. Each of these directories behave as a mount of the corresponding Storage Checkpoint with some exceptions.
See the *Veritas Storage Foundation and High Availability Administrator's Guide*.

Partitioned directories

Normally, a large volume of parallel threads performing access and updates on a directory that commonly exist in an file system suffers from exponentially longer wait times for the threads. This feature creates partitioned directories to improve the directory performance of file systems. When any directory crosses the tunable threshold, this feature takes an exclusive lock on the directory inode and redistributes the entries into various respective hash directories. These hash directories are not visible in the name-space view of the user or operating system. For every new create, delete, or lookup thread, this feature performs a lookup for

the respective hashed directory (depending on the target name) and performs the operation in that directory. This leaves the parent directory inode and its other hash directories unobstructed for access, which vastly improves file system performance.

See the *Administrator's Guide*.

Delayed allocation for extending writes

Performance of extending writes on local mounts is improved using the delayed allocation feature, which is turned on by default for all applicable writes.

See the *Administrator's Guide*.

Parallel Direct I/O on Linux

On VxFS, each iovec is performed synchronously for the readv(2) call and writev(2) call. For both readv(2) and writev(2), the Single Unix Specification states, "The readv/writev() function shall always fill an area completely before proceeding to the next." However, for direct I/O, Linux ignores this requirement and submits a number of iovecs in parallel before waiting for completion. In this release, VxFS now performs parallel direct I/O for both reads and writes, which improves VxFS performance. This support for parallel direct I/O can be enabled by setting the VxFS module load tunable vx_parallel_dio.

See the *Administrator's Guide*.

vxfsconvert can upgrade additional Veritas File System disk layout versions

The vxfsconvert command can upgrade the VxFS disk layout Version 4.

FileSnap creation over Network File System

You can create a FileSnap over Network File System (NFS) by creating a hard link from an existing file to a new file with the extension ":snap:vxfs:".

See the *Administrator's Guide*.

Free space defragmentation

You can now specify the -C option with the fsadm command to minimize file system free space fragmentation. This attempts to generate bigger chunks of free space in the specified device.

Improved fsadm fragmentation reporting

The fsadm command's fragmentation reporting now includes the file fragmentation index and the free space fragmentation index. Both of these indices range between 0 and 100, and give an idea about the level of file fragmentation and free space

fragmentation, respectively, which gives you a good indication of when you should defragment a file system.

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 6.0:

Recovery for synchronization tasks

In this release, VxVM tracks the plex synchronization for the following commands: vxplex att, vxassist mirror, vxsnap addmir, vxsnap reattach, and vxsnap restore. If the system crashes or the vxconfigd daemon fails, VxVM provides automatic recovery for the synchronization task. When the system is recovered,VxVM restarts the synchronization from the point where it failed. The synchronization occurs in the background, so the volume is available without delay.

Secure deletion of Veritas Volume Manager disks

When you decommission a disk that contained sensitive data, you may need to destroy any remaining data on the disk. In this release, VxVM provides the ability to shred the data on the disk to minimize the chance that the data is recoverable. When you specify the disk shred operation, VxVM shreds the entire disk, including any existing disk labels. After the shred operation, VxVM writes a new empty label on the disk to prevent the disk from going to the error state. The VxVM shred operation overwrites all of the addressable blocks with a digital pattern in one, three, or seven passes.

Caution: All data in the volume will be lost when you shred it. Make sure that the information has been backed up onto another storage medium and verified, or that it is no longer needed.

For more information on shredding disks, see the *Veritas Storage Foundation Administrator's Guide*.

Creating a volume of maximum size

In previous releases, Veritas Volume Manager provided a two-step approach to creating a volume of the maximum size. You had to run the vxassist maxsize command to find the maximum size of the volume to be created with the given constraints. Then, you had to run the vxassist make command and specify the volume size as the maximum determined by the vxassist maxsize command.

In this release, you can create a maximum sized volume with a single command. Specify the vxassist make command with the maxsize keyword. The vxassist

command creates the maximum sized volume possible, taking into consideration any other allocation attributes that you specify.

Veritas Volume Manager co-existence with Oracle Automatic Storage Management (ASM) disks

ASM disks are the disks used by Oracle Automatic Storage Management software. Veritas Volume Manager (VxVM) co-exists with Oracle ASM disks, by recognizing the disks as the type Oracle ASM. VxVM protects ASM disks from any operations that may overwrite the disk. VxVM classifies and displays the ASM disks as ASM format disks. You cannot initialize an ASM disk, or perform any VxVM operations that may overwrite the disk.

For more information about VxVM co-existence with ASM disks, see the *Veritas Storage Foundation Administrator's Guide*.

Changing VxVM tunables

The `vxtune` command is used to display or modify the values of Veritas Volume Manager tunable parameters. In this release, the `vxtune` command is extended and enhanced. The `vxtune` command has the following new functionality:

- manages an extended list of Veritas Volume Manager tunable parameters, including Veritas Volume Replicator and Cluster Volume Manager tunable parameters.

- provides a template format for tuning parameters. The template feature enables you to export the list of tunable parameters into a file, modify the values as necessary, then reload the tunables with an import command.

- enhanced command output. The output now displays the current value, the default value, and whether a reboot is required for the new value to take effect. Optionally, the output displays a description of the tunable parameters.

- makes the tunable values persistent across reboots.

- categorizes the tunable parameters by VxVM component. Specify the component to list or export the tunable parameters in that category. The components are the following:

    - basevm
      Basic core VxVM functionality.

    - fmr
      FlashSnap functionality.

    - cvm
      Cluster Volume Manager.

    - vvr

Veritas Volume Replicator.

Changes to the instant snapshot (version 20) data change object (DCO) volume layout

In this release, the volume layout of the data change object (DCO) has been changed to improve the I/O performance and scalability of instant snapshots. The change in layout does not alter how you administer instant snaphosts. The only visible effect is in improved I/O performance and in some cases, increased size of DCO volume. As with previous releases, you create DCOs for instant snapshots using "vxsnap prepare" or by specifying "logtype=dco dcoversion=20" while creating volume with "vxassist make".

The instant snapshot DCO (previously known as a version 20 DCO) now uses dynamic creation of maps on the preallocated storage.

Veritas Volume Manager throttling of administrative I/O

In this release, Veritas Volume Manager (VxVM) provides throttling of administrative I/O. During heavy I/O loads, VxVM throttles I/O that it creates to do administrative operations. This behavior ensures that the administrative I/Os do not affect the application I/O performance. When the application I/O load is lighter, VxVM increases the bandwidth usage for administrative I/O operations.

VxVM automatically manages the I/O throttling for administrative tasks, based on its perceived load on the storage. Currently, I/O throttling is supported for the copy operations which use ATOMIC_COPY and involve one destination mirror. The I/O throttling is transparent, and does not change the command usage or output. The following commands are supported:

- `vxassist mirror`

- `vxassist snapcreate`

- `vxevac`

- `vxplex att`

- `vxplex cp`

- `vxplex mv`

- `vxprint`

- `vxsnap addmir`

- `vxsnap reattach`

- `vxsd mv`

- `vxtune`

The administrative I/O operations allocate memory for I/O from a separate memory pool. You can tune the maximum size of this pool with the tunable parameter, `vol_max_adminio_poolsz`.

See the *Veritas Storage Foundation Administrator's Guide* for information about tuning the `vol_max_adminio_poolsz` parameter.

Command completion for Veritas commands

Veritas Storage Foundation and High Availability now supports command completion for Veritas Volume Manager (VxVM) commands and Dynamic Multi-Pathing (DMP) commands. In this release, command completion is supported only on the bash shell. The shell must be bash version 2.4 or later.

To use this feature, press **Tab** while entering a supported VxVM or DMP command. The command is completed as far as possible. When there is a choice, the command completion displays the next valid options for the command. Enter one of the displayed values. A value in brackets indicates a user-specified value.

Note: Platform-specific options are not supported with command completion in this release.

The following commands support command completion:

- vxassist
- vxdisk
- vxplex
- vxprint
- vxsnap
- vxstat
- vxtune
- vxcache
- vxconfigd
- vxtask
- vxreattach
- vxdmpadm
- vxddladm
- vxvol

- vxcdsconvert
- vxresize
- vxdctl
- vxsd
- vxdisksetup
- vxdiskunsetup
- vxrecover
- vxedit
- vxdg
- vxclustadm

vxdisk -o thin list command now shows the disk space used by a VxFS file system

The `vxdisk -o thin list` command now shows the disk space used by a VxFS file system.

Changes related to Veritas Dynamic Multi-Pathing (DMP)

The following sections describe changes in this release related to Veritas Dynamic Multi-Pathing (DMP).

Tuning Dynamic Multi-Pathing with templates

Veritas Dynamic Multi-Pathing (DMP) has multiple tunable parameters and attributes that you can configure for optimal performance. In this release, DMP introduces a template method to update several tunable parameters and attributes with a single operation. The template represents a full or partial DMP configuration, showing the values of the parameters and attributes of the host.

To view and work with the tunable parameters, you can dump the configuration values of the DMP tunable parameters to a file. Edit the parameters and attributes, if required. Then, load the template file to a host to update all of the values in a single operation.

For more information about tuning DMP with templates, see the *Veritas Dynamic Multi-Pathing Administrator's Guide*.

You can manage the DMP configuration file with the `vxdmpadm config` commands.

See the `vxdmpadm`(1M) man page.

Changes to DMP support for ALUA arrays

In this release, DMP has improved support for ALUA arrays. DMP now efficiently handles most implementations of the ALUA standard. The enhancements include the following:

- DMP now detects whether an ALUA array is A/A-A, A/A or A/P-F.

- DMP handles the array state correctly, when a node is taken out of the cluster. The enclosure level attribute failoverpolicy is now set internally.

- DMP handles Standby and unavailable LUN states for ALUA arrays.

- DMP monitors LUN ownership changes. DMP can shift the I/O load depending on the current state of the LUN.

DMP detects "persist through power loss" storage device server capability

In this release, DMP detects when a storage device server has the capability "persist through power loss". Certain arrays, such as Oracle's Sun Storage 7310, use this capability to preserve the persistent reservation and registrations across power cycles, controller reboots, and other similar operations.

If DMP detects that the device supports this capability, then DMP sets the APTPL (Activate Persist Through Power Loss) bit to 1 in the PERSISTENT RESERVE OUT parameter data sent with a REGISTER, REGISTER AND IGNORE EXISTING KEY service action, according to SPC-3 specifications.

When APTPL is set to 1, the persistent reservation (PR) keys are preserved during array controller takeover or failback operations.

Dynamic Multi-Pathing (DMP) detects and reports extended attributes from Veritas Operations Manager

If you have Veritas Operations Manager (VOM), and you have configured a central Management Server, the Device Discovery layer (DDL) of DMP can obtain extended attributes for managed hosts. DDL obtains these additional attributes out of band from the VOM database. DMP displays these attributes as output of the `vxdisk -p list` command.

See the *Administrator's Guide*.

DMP enhancements

The following DMP enhancements have been made in this release:

- The `vxdmpadm enable` command and the `vxdmpadm disable` command now accept multiple controllers on the command line.

- In addition, you can now enable or disable paths between a given controller and a port-id pair. If you specify both an HBA controller and an array port, DMP disables I/O on the specific portion of the Storage Area Network (SAN).

- The `vxdmpadm stat errord` command and the `vxdmpadm stat restored` command are deprecated.

  To see status for the restore tasks, use the `vxdmpadm gettune` command.

- Excluding or including paths from DMP is deprecated.

  Excluding paths from DMP but not from VxVM can lead to unsupported configurations. The command operations to exclude or include paths from DMP are now deprecated. You can exclude or include paths from VxVM. The deprecated commands are as follows:

  `vxdmpadm exclude dmp`

  `vxdmpadm include dmp`

  `vxdiskadm`: DMP options under `Suppressing or including devices for VxVM`

- `vxddladm list devices` command now displays the name of the ASL even if the device is skipped.

- `vxddladm status eventsource` is added to show the status of the `vxesd` daemon

- `vxscsiinq` diagnostic utility is enhanced to take hexadecimal page numbers as arguments.

Changes related to replication

Veritas Storage Foundation and High Availability Solutions includes the following changes related to replication in 6.0:

File Level Replication on Linux

Veritas File Replicator (VFR) supports file-level replication of application data, tracks all updates to the File System and periodically replicates these updates at the end of a configured time interval. VFR leverages Veritas File System (VxFS) data deduplication and will not replicate data that is already on the destination. VFR also supports VxFS compression and compressed files will be replicated as such. In addition, VFR also supports reversible data transfers. VFR is available as an option to Storage Foundation, included in the Veritas Replicator (new name for Veritas Volume Replicator) license and also in Symantec VirtualStore 6.0.

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

vvrcheck configuration utility

There is now a configuration utility, `/etc/vx/diag.d/vvrcheck`, that displays current replication status, detects and reports configuration anomalies, and creates statistics files that can be used by display tools. The `vvrcheck` also runs diagnostic checks for missing daemons, valid licenses, and checks on the remote hosts on the network. For more information, see the `vvrcheck(1M)` man page.

SmartMove for VVR

The initial sync between the Primary and Secondary is performed using the autosync option. The autosync to sync the volume now uses the SmartMove API from VxFS and provides the data only sync between the Primary and Secondary. This increases the initial autosync performance, which is dependent on the file system usage in the volume. This feature also helps thin provision LUNs configured on the Secondary site to use storage space only for data.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Thin provisioning and reclamation support for VVR

Storage reclamation is now possible on VVR volumes with VxFS file system on it. The storage corresponding to the volumes on the Secondary RVG is automatically reclaimed when the Primary volumes are reclaimed. The existing `vxdisk reclaim` or `fsadm -R` commands function for reclaiming VVR objects as well. For storage reclamation to work, the volumes on the Primary RVG must be mounted.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Enable compression with VVR

VVR compression lets you send data over the network in a compressed format from a Primary to one or more Secondary hosts. Compression reduces network bandwidth consumption and is useful in scenarios where there is low available bandwidth or where the bandwidth is shared among several applications. The compression option can be enabled on a per system or per Secondary basis using the CLI.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Replication performance improvement

Replication performance is improved by introducing Secondary logging (logging the I/O on the Secondary SRL before writing to the data volume). The primary requirement for this feature to work is to have the same size SRL on both the

Secondary and Primary. The Secondary SRL is used for staging the I/O from the Primary, and parallelize the data volume write. This improves the replication performance both in VVR and CVR. By default, this feature is enabled in 6.0.

There are other replication-specific tunables that may be increased to obtain the maximum replication performance.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide.*

Support for 8-node cluster applications

In a shared disk group environment, VVR supports replication of 8-node cluster applications. In previous releases, support was limited to 4-node cluster applications.

The following improvements enable scalability to 8-node support:

■ Improved message processing allows the logowner to process more messages per second, resulting in improved application throughput

■ Secondary logging feature improves replication performance

■ Improved CPU usage provides more CPU cycles to the logowner to process requests from other nodes in a cluster

■ Increased limit on max outstanding I/Os with VVR

See the *Veritas Storage Foundation and High Availability Solutions Replication 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.

Support for space-optimized snapshots for database cloning

You can use Storage Foundation for Databases (SFDB) tools to take space-optimized snapshots of your Oracle database and then create database clones by using those snapshots. SFDB tools use the underlying features of Storage Foundation for this operation.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases.*

SmartTier integration with OEM

You can now view the following SmartTier related information in the Oracle Enterprise Manager (OEM) reports:

- Storage allocation and free space in each tier
- Space occupied by a data file in each tier
  This is useful when a part of a data file is moved from tier to tier when database objects such as table or index are moved.

Enhancements to Cached ODM Advisor (dbed_codm_adm)

You can use the Cached ODM Advisor command `dbed_codm_adm` to generate a variety of reports that help you determine which data files are suitable for enabling Cached ODM. The reports generated by Cached ODM Advisor are enhanced to use the historical data from Oracle Automatic Workload Repository (AWR).

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Support for space-optimized snapshots on DR site for database cloning

You can use Storage Foundation for Databases (SFDB) tools in a replicated environment to take space-optimized snapshots on a disaster recovery (DR) site. This functionality lets you create clones of your Oracle database on the DR site in a setup where the database on the primary site is being replicated on the DR site.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Single CLI for different point-in-time copy operations

You can use the new SFDB command `vxsfadm` to perform various point-in-time copy operations on your Oracle database. `vxsfadm` provides the following benefits:

- Uniform command line for multiple operations
- Use case based functionality
- Enhanced error handling

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Support for file-level snapshots for database cloning

You can use Storage Foundation for Databases (SFDB) tools to take file-level snapshots of your Oracle database and then create one or more clones based on those snapshots. SFDB tools use the underlying features of Storage Foundation for this operation.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Enhanced authentication support

The authentication support for Storage Foundation for Databases (SFDB) tools is enhanced in this release. You can use the `sfae_auth_op` to set up and configure authentication for SFDB tools.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Changes related to Veritas Cluster Server (VCS)

The following sections contain changes related to VCS kernel components such as LLT, GAB, and I/O fencing, and clusters in secure mode.

For more information on changes related to VCS, see the *Veritas Cluster Server Release Notes*.

Changes to SFHA clusters in secure mode

In this release, the installation and configuration experience of secure cluster is considerably simplified. You can easily convert the cluster into secure cluster with this simplified secure cluster configuration model.

The new architecture is based on embedded VxAT, where the security components are installed as a part of the SFHA package. The root broker is no longer a single-point-of-failure in the new architecture. There is no dependency on a separate VRTSat package. Non-root users who are already logged on SFHA hosts are now not prompted for password. Additionally, a cluster-level user feature is introduced to simplify user administration in secure clusters.

See the *Installation Guide* and *Administrator's Guide* for more information.

Changes to LLT

This release includes the following new features and changes to LLT:

- LLT now supports VLAN tagging (IEEE 802.1Q).

- The `lltconfig` command includes the following new options:

  - -N
    You can use this option to list all the used cluster IDs.

  - -M
    You can use this option to display the currently loaded LLT module version information.

See the `lltconfig` manual page for more information.
See the `llttab` manual page for more information.

- Link utilization statistics are enhanced that help in the root cause analysis of performance related issues.

- Periodic flushing of ARP cache is disabled.

- When MAC address of a NIC changes, LLT immediately relearns the new MAC address and also updates the peer nodes about the change.

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

Changes to GAB

This section covers the new features and changes related to GAB in this release.

Better GAB and I/O fencing integration to ensure application availability

In the event of a split-brain situation before VxFEN module implements the decision, sometimes GAB proceeds with attempting to resolve the join after the split-brain. GAB removes all but one joining subcluster. This behavior can cause the entire cluster to shut down. To avoid this scenario, GAB now gives priority to the fencing module.

With the GAB and I/O fencing integration in this release, if the I/O fencing module's decision is still pending before GAB initiates a join of the subcluster, GAB delays the iofence message. GAB wait depends on the value of the VxFEN tunable parameter *panic_timeout_offst* based on which VxFEN computes the delay value and passes to GAB.

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

GAB can now recognize clients with names in addition to ports

When kernel clients initialize GAB API, they can now define a client name string. GAB now adds a client name which enables GAB to track the client even before GAB port is registered. GAB also passes the client name information to LLT when registering the LLT port. The `lltstat -p` command also displays the GAB client names when providing the status details of the ports in use.

This feature is applicable only to GAB kernel clients, and not applicable for user-land GAB clients such as HAD.

The gabconfig command has new -C option

The -C option of the `gabconfig` command lists the names of the GAB clients that have registered with GAB. The `-C` option when used with `-a` option lists the client names along with the port membership details.

Changes to I/O fencing

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

Support for racer node re-election during I/O fencing race

At the time of a network partition, the VxFEN module elects the lowest node in each sub-cluster as the racer node to race for the coordination points on behalf of the sub-cluster. The other spectator nodes wait on the racer node to do the fencing.

In the previous releases, the I/O fencing race was entirely dependent on the single racer node as follows:

■ If the racer node is not able to reach a majority of coordination points, then the VxFEN module on the racer node sends a LOST_RACE message and all nodes in the subcluster also panic when they receive the LOST_RACE message.

■ If the racer node panics during the arbitration, then the spectator nodes in the sub-cluster assume that the racer node lost the race and the spectator nodes also panic.

With the new racer node re-election feature, the VxFEN module re-elects the node with the next lowest node id in the sub-cluster as the racer node. This feature optimizes the chances for the sub-cluster to continue with the race for coordination points.

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

With fencing enabled, GAB can now automatically seed the cluster when some cluster nodes are unavailable

In the earlier releases, if some of the nodes are not up and running in a cluster, then GAB port does not come up to avoid any risks of preexisting split-brain. In such cases, you can manually seed GAB using the command `gabconfig -x` to bring the GAB port up. However, if you have enabled I/O fencing in the cluster, then I/O fencing can handle any preexisting split-brain in the cluster.

In this release, I/O fencing has extended this functionality to be able to automatically seed GAB as follows:

■ If a number of nodes in a cluster are not up, GAB port (port a) still comes up in all the member-nodes in the cluster.

■ If the coordination points do not have keys from any non-member nodes, I/O fencing (GAB port b) also comes up.

This new functionality is disabled by default. You must manually enable this automatic seeding feature of GAB in clusters where I/O fencing is configured in enabled mode.

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

You can still use the `gabconfig -x` command to manually seed the cluster.

Installer support to migrate between fencing configurations in an online cluster

You can now use the installer to migrate between disk-based and server-based fencing configurations. You can also replace the coordination points for any I/O fencing configuration in an online cluster using the same installer option. The installer uses the `vxfenswap` script internally.

You can also use response files to perform these I/O fencing reconfiguration operations.

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

Support for multiple virtual IP addresses in CP servers

You can now configure multiple network paths (virtual IP addresses) to access a CP server. CP server listens on multiple virtual IP addresses. If a network path fails, CP server does not require a restart and continues to listen on one of the other available virtual IP addresses.

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

Support for Quorum agent in CP servers

With the support for multiple virtual IP addresses, you can now use the Quorum agent to configure CP server service group failover policies. You can specify the minimum number of IP resources that must be online for the Quorum resource to remain online.

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

Graceful shutdown of a node no longer triggers I/O fencing race condition on peer nodes

In the earlier releases, a gracefully leaving node clears its I/O fencing keys from coordination points. But the remaining sub-cluster races against the gracefully

leaving node to remove its registrations from the data disks. During this operation, if the sub-cluster loses access to the coordination points, the entire cluster may panic if the racer loses the race for coordination points.

In this release, this behavior has changed. When a node leaves gracefully, the CVM or other clients on that node are stopped before the VxFEN module is unconfigured. Hence, data disks are already clear of its keys. The remaining sub-cluster tries to clear the gracefully leaving node's keys from the coordination points but does not panic if it is not able to clear the keys.

Entering and displaying values in human-friendly units

Storage Foundation now supports reporting and inputting values in human-friendly units.

The following commands were modified to display human-friendly units:

- `diskusg`

- `ff`

- `fsadm`

- `fsckptadm`

- `fsvoladm`

- `vxdg free`

- `vxdisk list`

- `vxdisk -o thin list`

- `vxdisk -o thin,fssize list`

- `vxdmpadm iostat show`

- `vxedquota`

- `vxmemstat`

- `vxprint`

- `vxquot`

- `vxquota`

- `vxrepquota`

- `vxstat`

- `vxtune`

See the manual pages for more information.

Displaying SFHA information with vxlist

The `vxlist` command is a new display command that provides a consolidated view of the SFHA configuration. The `vxlist` command consolidates information from Veritas Volume Manager (VxVM) and Veritas File System (VxFS). The `vxlist` command provides various options to display information. For example, use the following form of the command to display file system information including information about the volume, disk group, and so on. In previous releases, you needed to run at least two commands to retrieve the following information.

```
# /opt/VRTSsfmh/bin/vxlist fs
TY FS   FSTYPE SIZE    FREE    %USED DEVICE_PATH          MOUNT_POINT
fs /    ext3   65.20g 51.70g 17%    /dev/sda1                 /
fs mnt  vxfs   19.84g  9.96g 49%    /dev/vx/dsk/bardg/vol1  /mnt
```

For help on the `vxlist` command, enter the following command:

```
# vxlist -H
```

See the `vxlist(1m)` manual page.

Support for Kernel-based Virtual Machines (KVM) on Linux

Storage Foundation High and Availability Solutions provide configurations to enhance the Kernel-based Virtual Machine (KVM) environment. Storage Foundation High and Availability Solutions 6.0 products are supported on the Red Hat Enterprise Linux (RHEL) 6.1 distribution.

Storage Foundation and High Availability Solutions products provide the following functionality for KVM guest virtual machines:

- Storage visibility

- Storage management

- High availability

- Cluster failover

- Replication support

For implementation information:

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

Licensing changes in the SFHA Solutions 6.0 release

Storage Foundation and High Availability Solutions 6.0 introduces the following licensing changes:

- The Cluster File System license is deprecated. CFS customers are entitled to the Storage Foundation Cluster File System High Availability (SFCFS HA) functionality.

- The VVR Option is renamed as Veritas Replicator Option. This option includes VVR (volume-based replication) and the new file-based replication solution.

- The VVR Enterprise license is deprecated; you can use Storage Foundation Enterprise and add Veritas Replicator Option to get this functionality. VVR Enterprise customers are entitled to Storage Foundation Enterprise with Replicator Option.

- The VCS license enables full cluster functionality as well as the limited start/stop functionality.

- Storage Foundation Enterprise CFS for Oracle RAC (Linux/x64) customers are entitled to Storage Foundation Enterprise for Oracle RAC (Linux/x64.)

The following functionality is included in the Standard and Enterprise licenses:

- The Compression feature is available with the Standard license.

- The SmartTier feature is now available with the Standard license.

- The Deduplication feature is available with the Enterprise license.

The following products are included in this release:

- Dynamic Multi-Pathing

- VirtualStore

- Storage Foundation Basic

- Storage Foundation Standard

- Storage Foundation Enterprise

- Veritas Cluster Server

- Veritas Cluster Server HA/DR

- Storage Foundation Standard HA: Storage Foundation Standard plus Veritas Cluster Server

- Storage Foundation Enterprise HA: Storage Foundation Enterprise plus Veritas Cluster Server

- Storage Foundation Enterprise HA/DR

- Storage Foundation Enterprise Cluster File System HA

- Storage Foundation Enterprise Cluster File System HA/DR

- Storage Foundation Enterprise for Oracle RAC

- Storage Foundation Enterprise HA/DR for Oracle RAC

- Storage Foundation Enterprise for Sybase ASE CE

- Storage Foundation Enterprise HA/DR for Sybase CE

HA: High Availability

HA/DR: High Availability and Disaster Recovery

Veritas Replicator Option can be added to all Storage Foundation and High Availability products, except Dynamic Multi-Pathing and Veritas Cluster Server.

Note that products, features, and options may differ by operating system and platform. Please see the product documentation for information on supported platforms.

Changes related to installation and upgrades

The product installer includes the following changes in 6.0.

The installsfha script and the uninstallsfha script are now available

The installsfha script and the uninstallsfha script scripts are now available in the storage_foundation_high_availability directory to directly install, uninstall, or configure the Storage Foundation and High Availability product.

See the *Veritas Storage Foundation and High Availability Installation Guide* for more details.

Unencapsulation not required for some upgrade paths

Unencapsulation is no longer required for certain upgrade paths.

See the *Veritas Storage Foundation and High Availability Installation Guide* for more details.

Creating a backup boot disk group when the boot disk is encapsulated and mirrored during upgrades

When you upgrade from a 5.1 Service Pack (SP) 1 or later release, the installer can split a mirrored boot disk group to create a backup disk group. You can use this backup in case of an upgrade failure.

Support for product installation using yum on Linux

You can now install any of the Veritas products with yum. Yum installation is supported for Red Hat Enterprise Linux 5 and 6.

See the *Installation Guide* for more information.

The installer can now detect duplicate VCS cluster IDs and can automatically generate cluster IDs

The installer can now detect duplicate VCS cluster IDs and prompt you to select an unused one. It can also generate an unused ID during installation.

The installer can check product versions and hotfixes

You can check the existing product versions using the installer command with the `-version` option before or after you install. After you have installed the current version of the product, you can use the `showversion` script in the /opt/VRTS/install directory to find version information.

You can discover the following information with these commands:

- The installed version of all released Strorage Foundation and High Availability Suite of products

- The missing required RPMs or patches as applicable for platform

- The available updates (including patches or hotfixes) from SORT for the installed products

Depending on the product, the script can identify versions from 4.0 onward.

Using the installer's postcheck option

You can use the installer's postcheck option to diagnose installation-related problems and to provide troubleshooting information.

Rolling upgrade improvements

The rolling upgrade procedure has been streamlined and simplified.

Allow Response files to change tuning parameters

You can set non-default product and system tunable parameters using a tunables template file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing during or after the installation procedure.

See the *Installation Guide* for more information.

Packaging updates

The following lists the package changes in this release.

- New `VRTSsfcpi60` RPM for product installer scripts
  The `VRTSsfcpi60` RPM is introduced in this release. The VRTSsfcpi60 RPM contains the installer scripts and libraries that the installer uses to install, configure and upgrade Veritas products.

- New `VRTSfsadv` RPM for product data deduplication
  The `VRTSfsadv` RPM is introduced in this release. The `VRTSfsadv` RPM contains the libraries for the data deduplication feature.

For more information, see the *Installation Guide*.

Enhancements to collecting a VxExplorer troubleshooting archive

The Symantec Operations Readiness Tools (SORT) data collector contains functionality to collect and submit a VxExplorer archive. You can send this archive to Symantec Technical Support for problem diagnosis and troubleshooting. VxExplorer does not collect customer data.

The legacy `VxExplorer` script now works differently. When you run the script, it launches the SORT data collector on the specified local host with the `-vxexplorer` option.

To learn more about using the data collector to collect a VxExplorer archive, see:

www.symantec.com/docs/HOWTO32575

Changes related to product documentation

The Storage Foundation and High Availability Solutions 6.0 release includes the following changes to the product documentation.

Table 1-1 lists the documents introduced in this release.

Table 1-1   New documents

New documents	Notes
Veritas Storage Foundation Installation Guide	Installation and upgrade information for Storage Veritas Foundation.
Veritas Storage Foundation Administrator's Guide	Administration information for Veritas Storage Foundation.

Table 1-1 New documents *(continued)*

New documents	Notes
Veritas Storage Foundation and High Availability Release Notes	Release-specific information for Veritas Storage Foundation and High Availability users.
Veritas Storage Foundation and High Availability Solutions Solutions Guide	Solutions and use cases for Veritas Storage Foundation and High Availability Solutions.
Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide	Troubleshooting information for Veritas Storage Foundation and High Availability Solutions.
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	Virtualization-related information for Veritas Storage Foundation and High Availability Solutions.
Symantec VirtualStore Release Notes	Release-specific information Symantec VirtualStore.
Veritas Storage Foundation for Sybase ASE CE Release Notes	Release-specific information for Veritas Storage Foundation for Sybase ASE CE.
Veritas Storage Foundation for Sybase ASE CE Installation Guide	Installation information for Veritas Storage Foundation for Sybase ASE CE.
Veritas Storage Foundation for Sybase ASE CE Administrator's Guide	Administration information for Veritas Storage Foundation for Sybase ASE CE.
Virtual Business Services–Availability User's Guide	Information about Virtual Business Services. This document is available online.

Table 1-2 lists the documents that are deprecated in this release.

Table 1-2 Deprecated documents

Deprecated documents	Notes
Veritas File System Administrator's Guide	Content now appears in the *Veritas Storage Foundation Administrator's Guide* and in the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

Table 1-2   Deprecated documents *(continued)*

Deprecated documents	Notes
Veritas Volume Manager Administrator's Guide	Content now appears in the *Veritas Storage Foundation Administrator's Guide* and in the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide.*
Veritas Storage Foundation Advanced Features Administrator's Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Solutions Guide.*
Veritas Volume Manager Troubleshooting Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide.*
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	Content now appears in the *Veritas Cluster Server Bundled Agents Reference Guide.*
Veritas Volume Replicator Planning and Tuning Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide.*
Veritas Volume Replicator Advisor User's Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide.*

Table 1-3 lists documents that are no longer bundled with the binaries. These documents are now available online.

Table 1-3   Online documents

Document
Veritas Cluster Server Agent Developer's Guide
Veritas File System Programmer's Reference Guide

No longer supported

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

■ Several documents are deprecated in this release.
   See "Changes related to product documentation" on page 32.

- Disk layout Version 4 is no longer supported. You cannot create nor mount a file system with disk layout Version 4. You can use the `vxfsconvert` utility to upgrade the disk layout to Version 7 or later after installing this release.
  See the `vxfsconvert`(1M) manual page.

- Disk layout Version 6 is deprecated. You can only local mount a file system with disk layout Version 6, and the only operation that you can perform is to upgrade the disk layout to a supported version by using the `vxupgrade` utility. Symantec recommends that you upgrade from Version 6 to the latest default disk layout version. You cannot create new file systems with disk layout Version 6. If you upgrade a file system from disk layout Version 6 to a later version, once the upgrade operation finishes, you must unmount the file system cleanly, then re-mount the file system.
  See the `vxupgrade`(1M) manual page.

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

The following topics describe the system requirements for this release:

Supported Linux operating systems

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

Table 1-4 shows the supported Linux operating systems for this release.

Table 1-4   Supported Linux operating systems

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 6	Update 1, 2	2.6.32-131.0.15.el6 2.6.32-220.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only

Table 1-4   Supported Linux operating systems *(continued)*

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 5	Update 5, 6, 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 11	SP1	2.6.32.12-0.7	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	**6.1	2.6.32-131.0.15.el6	64-bit x86, EMT*/Opteron
Oracle Linux 5	**Update 5, 6, 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron

* Extended Memory Technology

** RHEL-compatible mode only.

Note: Only 64-bit operating systems are supported.

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.

For DMP, SF, SFHA, SFCFSHA, SFRAC, VCS, and VirtualStore, Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

On Linux, 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

Hardware compatibility list (HCL)

The hardware compatibility list contains information about supported hardware and is updated regularly. 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.

For the latest information on supported hardware, visit the following URL:

http://www.symantec.com/docs/TECH170013

For information on specific High Availability setup requirements, see the *Veritas Cluster Server Installation Guide*.

VMware Environment

For information about the use of this product in a VMware Environment, refer to http://entsupport.symantec.com/docs/289033

Veritas Storage Foundation for Database features supported in database environments

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

Table 1-5   SFDB features supported in database environments

SFDB feature	DB2	Oracle	Sybase
Oracle Disk Manager, Cached Oracle Disk Manager	No	Yes	No
Concurrent I/O	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes
Database Storage Checkpoints	No	Yes	No
Database Flashsnap	No	Yes	No
SmartTier for Oracle	No	Yes	No

For the most current information on SFHA and single instance Oracle versions supported, see:

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

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

Veritas Storage Foundation memory requirements

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

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

Fixed issues

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

See the corresponding Release Notes for a complete list of fixed incidents related to that product.

See "Documentation" on page 115.

Installation and upgrades: Issues fixed in 6.0

Table 1-6   Fixed issues related to installation and upgrades

Incident	Description
1952659	If a system check fails on one node, the CPI allows you to proceed with the installation on the remaining systems.
2070448	Adding a node to a running cluster no longer fails in secure mode.
2173459	The installer no longer fails after starting GAB on new node if the cluster uses secure CPS.
2185707	The installer no longer hangs upon starting vxfen, when a remote node cannot connect by ssh.
2313718	The installer now provides the option to synchronize clocks if there is a clock skew of more than 5 seconds between nodes.

Table 1-6   Fixed issues related to installation and upgrades *(continued)*

Incident	Description
2370156	The `-version` option can now detect and show the versions of the packages when there is a mix of different versions, including RU upgrades.
2371882	The installer now provides a script `/opt/VRTS/install/showversion` for easier version checker calling.
2624441	SLES10 SP4, IBM DS APF, CPI, "VXFS" package installation failed.

Veritas Storage Foundation for Databases (SFDB) tools: Issues fixed in 6.0

Table 1-7   SFDB tools fixed issues

Fixed issues	Description
1840672	In a multiple disk group environment, if the snapshot operation fails then `dbed_vmsnap` fails to reattach all the volumes.
1469310	If the database fails over during FlashSnap operations, various error mesages appear.

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

Table 1-8 describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in this release.

Table 1-8   Storage Foundation for Databases fixed issues

Incident	Description
2203917	Process table has been changed to use per-hash-bucket locks, and the number of buckets has been increased from 32 to 256.
2237709	The `dbdst_preset_policy` command no longer aborts when you specify the volume class as MEDIUM.

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

There are no SFDB fixed issues in 5.1 SP1 RP2.

Veritas File System: Issues fixed in 6.0

Table 1-9   Veritas File System fixed issues

Incident	Description
2565400	Poor read performance with DSMC (TSM) backup on CFS file systems.
2528888	CFS mount fails after recovery from I/O path failure.
2528819	VxFS thread creates warning messages.
2527578	Panic in vx_bhash_rele().
2526174	Wrong offset calculation affects replication functionality.
2515459	mount command still hung even with the fix of e1466351.
2515380	ff_vxfs ERROR: V-3-24347: program limit of 30701385 exceeded.
2492304	File entry is displayed twice if the find or ls command is run immediately after creation.
2486589	Threads blocked behind vx_ireuse_steal().
2429566	Memory leak in internal buffercache after 497 days (lbolt wrap-over).
2412488	Do not disable read ahead on files with shared extents if page cache optimization is not enabled.
2386331	vxfs panic in iput() from vx_softcnt_flush() ,after filesystem full fsck,and run reboot.
2379139	newline in vx_osdep.c: snprintf(cmp->cm_name, sizeof(cmp->cm_name), "vxclonefs-%d\n" breaks native LVM (pvs).
2371710	User quota information corruption.
2343792	vxedquota slow on some systems.
2325449	VxFS unmount performance on large memory systems.

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

Incident	Description
2320611	There was discrepancy between vxi_bcache_maxkbyte and vx_bc_bufhwm.
2316051	DB2 9.5 onwards can cause contention of the mmap_sem.
2307933	Support online resize of RCQ similar to intent log.
2294285	WARNING message with fsmigadm start command.
2289522	Time and size issues in fsppadm query command output.
2246127	Mount should read IAUs multiple blocks and in parallal.
2242213	vx_sched' is hogging CPU resources.
2242211	Unable to grow filesystem on SVM Volume.
2226294	VXFS 5.1GA on CentOS 5.4 causing system panic in _list_add() corruption in vx_ftenter codepath when using named data streams.
2203917	Peformance problem with Oracle 10g using ODM.
2200631	Use fixed extent size for clone pushes; also increase max push per transaction from 64k to 256k.
2180476	System panic in vx_iupdat_clustblks().
2172485	Metadata was not updated correctly after write() with O_SYNC flag.
2152337	/dev/odm/* 666 permission.
2074806	dm_punch_hole() request does not invalidate pages.

Veritas File System: Issues fixed in 5.1 SP1 RP1

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

Table 1-10   Veritas File System fixed issues

Fixed issues	Description
1296491	Fixed issues seen during a force unmount of a parent cluster file system while a child was being mounted or umounted.

Table 1-10   Veritas File System fixed issues *(continued)*

Fixed issues	Description
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.
2030119	fsppadm core dumps when analysing a badly formatted XML file, is resolved
2032525	Fixed the cause of an NFS stale file handle.
2061554	Sequential extents are now collated.
2111921	Improved the performance of VxFS file systems with concurrent I/O or direct I/O enabled.
2149659	Fixed the cause of an error that resulted during the truncate operation of a file with a shared extent in the presence of a Storage Checkpoint containing FileSnaps.
2162822	During online migration from ufs to vxfs, df command returns a non-zero return value.
2163084	The listxattr() call now uses rwlock.
2169273	During online migration, nfs export of the migrating file system leads to system panic
2177253	A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for vxupgrade purposes
2178147	Linking a IFSOC file now properly calls vx_dotdot_op(), which fixes the cause of a corrupted inode.
2181833	The vxfilesnap command no longer gives an incorrect error message on a Storage Checkpoint file system.
2184528	fsck no longer fails to repair corrupt directory blocks that have duplicate directory entries.
2178147	Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system
2198553	A forced unmount now properly clears the bd_super structure member.
2221623	Fixed a performance loss due to a delxwri_ilist spin lock with the default values for vx_idelxwri_timelag.

Table 1-10   Veritas File System fixed issues *(continued)*

Fixed issues	Description
2226257	Fixed the cause of a system panic in the `in _list_add`() call, which led to corruption in the `vx_ftenter`() codepath when using named data streams.

Veritas File System: Issues fixed in 5.1 SP1 RP2

Table 1-11 describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP2.

Table 1-11   Veritas File System fixed issues

Fixed issues	Description
2340953	cfs.stress.enterprise hit an assert f:vx_iget:1a.
2481984	file system will hang if a customer creates 400 shares
2247387	LM stress.S3 test hit an assert "vx_ino_update:2"
2486589	threads blocked behind vx_ireuse_steal
2440584	node panic in vx_sync() during shutdown
2424240	Dedup ioctl sharing extents incorrectly under certain scenarios
2431674	panic in vx_common_msgprint() via vx_inactive()
2480935	V-3-26626: File Change Log IOTEMP and ACCESSTEMP index creation failure for /vx/fsvm with message Argument list too long
1892045	Improve the memory allocation for per-cpu data.
2413172	There is a priority 1 issue reported by AXA Rosenburg for Filestore replication and issue seems related to VxFS
2399228	TRuncate up size updates can be missed
2412604	It does not work when set homedir user softlimit numspace quota after generate data
2242630	Remove limits on inode and buffer cache sizes
2422574	Reboot one node and the node can't mount file system , after turn on the homedir quota on

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

Fixed issues	Description
2403126	cfs recovery didn't finished timely in the primary node after one slave left.
2283893	Add functionality of free space defragmentation through fsadm.
2372093	new fsadm -C hung
2387609	User quota corruption
2371710	user quota information corrupts on 5.1SP1
2371903	newline in vx_osdep.c: snprintf(cmp->cm_name, sizeof(cmp->cm_name), "vxclonefs-%d" breaks native LVM(pvs)
2384831	vxfs panic in iput() from vx_softcnt_flush() ,after filesystem full fsck,and run reboot
2399178	fsck : pass2c needs performance enhancements
2374887	Accessing FS hung. FS marked full fsck after reboot of node.
2374887	Accessing FS hung. FS marked full fsck after reboot of node.
2283315	cfs-stress_S5 hits assert of "f:vx_reorg_emap:10 via vx_extmap_reorg"
2368737	RCQ processing code should set FULLFSCK flag if it finds a corrupt indirect block.
1956458	fsckpt_fbmap for changed blocks failed with ENXIO due to inode mapped to hole in ILIST of down stream checkpoint
2337470	In the process of shrink fs, the fs out of inodes, fs version is 5.0MP4HF*
2332460	vxedquota slow on some systems
2300682	Question about IOTemp on fsppadm query
2316793	After removing files df command takes 10 seconds to complete
2302426	Unaligned Reference Fault in vx_copy_getemap_structs
2272072	Threads stuck in vx_rwsleep_rec_lock_em
2290800	investigation on ilist HOLE
2192895	Panic while set/get acls - possible race condition
2059611	Panic in vx_unlockmap() due to NULL ml_tranp

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

Fixed issues	Description
2282201	vxdump core dumped whilst backing up layout 7 local VxFS file system
2337737	killing IOs to a CFS and ls command to the same CFS is hanging.
2316094	There was discrepancy between vxi_bcache_maxbyte and vx_bc_bufhwm.
2253938	EAU delegation timeouts
2419991	ncheck: no way to limit output to specific filesets, as with limiting output to specific inodes
2419989	ncheck -i does not limit output to the specified inodes when using -o device/block/sector
2074806	dm_punch_hole request does not invalidate pages
2296107	Operation not applicable appear on fsppadm query result
2246579	Panic at getblk() when growing a full filesystem with fsadm
2061177	fsadm -de' command erroring with 'bad file number' on filesystem(s) on 5.0MP3RP1
1475345	write() system call hangs for over 10 seconds on VxFS 3.5 on 11.23
2251223	df -h after removing files takes 10 seconds
2253617	LM stress aborted due to "run_fsck : Failed to full fsck cleanly".
2220300	vx_sched' is hogging CPU resources.
2161379	repeated hangs in vx_event_wait()
1949445	hang due to large number of files in a directory
1296491	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2239412	system panics while writing to cfs share exported as NFS to ESX server4.1.
2169324	Test LM-stress_S5 hits an assert of "f:vx_idelxwri_off:5a vai vx_trunc_tran"

Veritas Volume Manager: Issues fixed in 6.0

Table 1-12   Veritas Volume Manager fixed issues

Incident	Description
2595557	Multiple execution of "sysctl -a" caused OS panic.
2578336	Failed disk due to cdsdisk format.
2561012	VxVM operations in failover clusters causes inconsistency in the public region disk_offset.
2560843	I/O hang in slave nodes after one of slave is rebooted in a 4-node setup.
2536667	System panics after xmfree in volcvmdg_delete_msg_receive and voldiodone.
2527289	Both sites gets detached after data/dco plex failue at each site, leading to I/O cluster wide outage.
2524936	Diskgroup disabled after vxconfigd found the process file table is full.
2513101	User data corrupted with disk label information.
2510523	I/O hangs on both master and slave after doing "vxclustadm setmaster".
2495332	vxcdsconvert is broken if the private region of the disk to be converted is less than 1 MB.
2495186	I/O throttling with TCP protocol due to memory flow control.
2489350	volkmsg_cb_t,vol_fsvm_info_t leaked in VVR Primary node.
2484685	Race between two vol_subdisk sios while done processing which causes one thread to free sio_fsvm_priv before other thread accesses it.
2484334	Panic in dmp_stats_is_matching_group.
2483053	Master node out of memory.
2445066	Panic in vol_rv_service_message_start on primary.
2441937	vxconfigrestore precommit fails with awk errors.
2440349	DCO volume may grow into any 'site' even when 'alloc=site:xxxx' is specified by a list of 'site' to be limited.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2438426	EFI flag is added to a path in ddi_path_list even though it is non-EFI.
2432006	Pending read count with kio cache is not decremented when read object is locked in transaction.
2431470	vxpfto uses DM name when calling vxdisk, but vxdisk will match DA name first and thus cause corruption.
2428875	I/O on both nodes (wait for the DCM flush started), and crash the slave node, lead to the master reconfiguration hang.
2428631	Allow same fence key to be used for all Disk groups.
2425722	vxsd move operation failed for disk size >= 2TB.
2425551	I/O hangs for 6 mintues or more when rebooting the slave node if there is I/O on both master and slave.
2425259	vxdg join fails with VE_DDL_PROPERTY: Property not found in the list.
2421067	vxconfigd hung in both nodes of primary.
2419803	Pinnacle secondary master panics at nmcom_send_tcp during autosync.
2419486	Data corruption when changing naming scheme.
2419348	DMP panic; race condition between DMP reconfig and DMP pass through ioctl.
2413904	Multiple issues are seen while performing dynamic LUN reconfiguration.
2411874	A failing disk does not detached across the cluster if the disk access failure is seen in a slave.
2411698	I/O hangs on both master and slave.
2410845	Lots of 'reservation conflict' messages seen in clustered environment with XIV arrays.
2407699	vxassist core dump if the /etc/default/vxassist file contains wantmirror=ctlr.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2407192	Application I/O hangs because of a race condition between CVM reconfiguration and log-owner change protocol.
2406292	Panic in vol_subdisksio_delete.
2400654	Stale array.info file can cause vxdmpadm commands to hang.
2396293	I/Os loaded, sanboot failed with a vxconfigd core dump.
2390431	VVR vxio panic at the end of autosync, when transitioning from DCM to SRL logging mode.
2389554	vxdg listssbinfo output is not correct.
2388725	Panic in dmp_get_dmpsymbols when attempting to load an APM.
2387993	Including/excluding libvxpp.so vxconfigd goes into disabled mode.
2386120	Enhancement request to add diagnostic logging to help triage a CVM master takeover failure situation.
2379029	Changing of enclosure name is not working for all devices in enclosure.
2367564	Long boot times observed due to vxvm-udev.sh since upgrading to 5.1SP1.
2365951	Growto failing with error V-5-1-10128 Unexpected kernel error in configuration update.
2364253	VVR: Kernel memory is leaked on VVR secondary while using SO snapshots.
2361295	CVM reconfiguration hang at vxconfigd level join when vxconfigd on MASTER node is either dead or restarted and a SLAVE node is trying to join the cluster and has to send connection request to MASTER vxconfigd.
2359814	vxconfigbackup doesn't handle errors well.
2358321	Remove usage of __invalidate_device() from VxVM. Symbol is no longer in kABI whitelist.
2357798	CVR:Memory leak due to unfreed vol_ru_update structure.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2357507	In presence of large number of NR (Not-Ready) devices, server panics due to NMI triggered and when DMP continuously generates large no of path disable/enable events.
2356744	VxVM script daemons should not allow its duplication instance in itself.
2356293	Hung in the DMP stack vxdmpread -> uphysio.
2355706	IO hang when cache object was full.
2349352	During LUN provisioning in single path IO mode environment a data corruption is observed.
2346470	Excluding and including a LUN in a loop triggers a huge memory leak.
2344186	CCT: Volume recovery is not clearing the needsync flag from volumes with DCO in BADLOG state causing node join to fail.
2337353	vxdmpadm include vxvm dmpnodename=<emcpower#> includes all excluded dmpnodes along with the requested one.
2337233	vxdmpadm exclude vxvm dmpnodename=<emcpower#> does not suppress TPD device.
2334757	memory consumption for the vxconfigd grows because of a lot of DMP_IDLE, DMP_UNIDLE events.
2334544	In CVM environment, vxconfigd level join is hung when Master returns error "VE_NO_JOINERS" to parallel slave joiners.
2334534	In CVM environment, vxconfigd level join is hung when Master returns error "VE_NO_JOINERS" to a joining node and cluster nidmap is changed in new reconfiguration.
2334485	vxconfigd restart on master node while CVM reconfiguration is on-going/pending leads to disabling of shared diskgroups.
2324507	The manpage for vxrelayout(1M) command is incorrect.
2323925	If rootdisk is encapsulated and if install-db is present, clear warning should be displayed on system boot.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2322752	Duplicate DA records seen for NR devices upon restart of vxconfigd.
2320917	vxconfigd core dump and lost diskgroup config after removing volume and disk on thin reclaim LUN.
2317703	Vxesd/Vxconfigd leaks file descriptors.
2317540	System panic due to kernel heap corruption while DMP device driver unload.
2316297	Error message "Device is in use" appears during boot time.
2313021	Sun cluster: CVM slave failed to join after reboot.
2299670	Disk Groups created on EFI LUNs do not auto import at boot time using VxVM version 5.1SP1 and later.
2291226	Skip writing backup label for CDS disks > 1TB to avoid block level corruption.
2286559	kernel heap corruption detected panic after array controller reboot.
2280624	Need to set site consistent only on mirrored-volumes.
2276542	HFs created for Linux contain additional APM key files for vscsi leading to errors at apm load time.
2268408	suppressing a powerpath disk's path using vxdiskadm 17-2 causes the disk to go in error state.
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated.
2253970	Support per-disk maxiosize for private region I/Os.
2253552	Leak in vxsfdefault_parse.y at function vxsf_getdefault (*val).
2252680	vxtask abort does not cleanup tasks properly.
2248730	vxdg import command hangs as vxrecover daemon (spawned by vxdg) doesn't close standard error stream.
2245121	Rlinks do not connect for NAT configurations.
2240056	vxdg move' transaction not completing and backups fail.

Table 1-12 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2234821	DMP can't detect the re-enabled os device status" does not work on RHEL5.
2233611	HDS wants the ASL for the USP-V & VSP (R700) to check page 00 to see if E3 is supported, if E3 is supported then issue inquiry on E3 and wants the R700 array name set as Hitachi VSP.
2232789	supporting NetApp Metro Cluster.
2230377	Differences based sync fails for volumes/RVG sizes greater than 1TB.
2228531	cvm master vxconfigd process hung in vol_klog_lock.
2227923	renaming of enclosure name is not persistent.
2226813	VVR: rlinks remain disconnected with UDP protocol if data ports are specified.
2226771	Unable to configure disk in VM using vxdiskadd or vxdiskadm.
2220926	vxprivutil -D set <attr>' command leads to permanent vxprivutil command hang.
2218470	Some of the VxVM init scripts need to be compliant to the Linux Standard Base.
2212784	Enhance VM device suppression and disk reconfiguration handling.
2205108	vxconfigd clubbing all luns in a single dmpnode.
2202710	VVR:During SRL to DCM flush, commands should not hang and come out with proper error.
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault.
2201056	DCO creation does not take into account the region size specified in "default" file.
2200670	vxattachd does not recover disks if disk group is not imported.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2199496	Data Corruption seen with "site mirror" Campus Cluster feature.
2197254	While creating volumes on thinrclm disks, the option "logtype=none" does not work with vxassist command.
2196918	Snapshot creation with cachesize fails, as it doesn't take into account diskgroup alignment.
2196480	The disk initialization failed due to wrong number of cylinders reported in devintf_disk_geom_raw gotten from raw geometry.
2194492	VxVM-ASM co-existence enablement.
2193429	IO policy not getting preserved when vxconfigd is restarted and migration from one devlist to other is taking place.
2192612	XP ASL is claiming EVA lun.
2191693	vxdmpadm native list' command is not displaying any output nor error.
2190020	dmp_deamon applying 1m continuous memory paging which is too large.
2188590	An ilock acquired by a slave node for a read on a DCL object can lead to I/O hang when the node becomes master before reading is done.
2183984	System panics due to race condition while updating DMP I/O statistics.
2181631	Striped-mirror volume cannot be grown across sites with -oallowspansites with DRL.
2176601	SRDF-R2 devices are seen in error state when devices are in write-protected mode.
2168720	Removal of stale ASL's.
2165394	Diskgroup imported by selecting wrong disks. After destroying original diskgroup, import without useclonedev option imports diskgroup with original disks rather than clone disks.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2165141	VxVM resets b_clock_ticks to zero if I/O hints are passed by VxFS.
2160199	Master takeover fails as the upcoming master could not import shared diskgroup.
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core.
2154287	Improve handling of Not-Ready (NR) devices that are triggering "VxVM vxdmp V-5-3-1062 dmp_restore_node: Unstable path" messages.
2152830	In a multilevel clone disks environment, a regular DG import should be handled properly and in case of DG import failure, it should report correct error message.
2148851	vxdisk resize failed to resize a disk that is expanded physically from the array console.
2144775	Failoverpolicy "local" is not getting preserved after VxVM upgrade.
2139179	SSB check invalid with a copy of a LUN.
2136046	Need to log additional information and handle hang case in CVM scripts.
2133503	Renaming enclosure results in dmpevents.log reports 'Mode for Enclosure has changed from Private to Private'.
2105547	tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations.
2104887	vxdg error messaging improvement required for cloned devices, report conflicting dgid and suggest running "-o updateid".
2102929	Deprecate vxdiskadm option 17:4/5/6/7 and 18:4/5/6/7 in configurations and modify 17:8 and 18:8 to drop support for exclude/include from vxdmp contol.
2100865	Memory leaks in vxconfigd.
2092921	Enhancements in vxrecover and if possible DCO plex attach implementation.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2088007	Possibility of reviving only secondary paths in DMP.
2082450	In case of failure, vxdisk resize should display more meaningful error message.
2081043	vxconfigd core dump in clist_next while running cvmtc.
2080730	VxVM/vxdmp exclude file contents after being updated should be consistent via vxdiskadm and vxdmpadm.
2070561	Improve diskgroup import error messaging in regards to cloned devices.
2038928	Creation of older version diskgroup fails.
2033909	In SFRAC configuration, I/O hung after disable secondary path of A/PG array Fujitsu ETERNUS3000.
2015467	Performance improvement in VxVM mapping provider.
2008721	DMAPI dm_handle_to_path() libxdsm.a call could be interrupted with kill -9 leaving session and access right.
2000661	Diskgroup rename during import with enhanced noreonline fails as slave uses the cached private region header info for diskgroup rather than info sent by master.
1959513	Propagate -o noreonline option of disk group import to slave nodes.
1940052	vxconfigd hung on master after removing the hba alias from zone and node leave followed by join.
1869002	Introduction of circular buffer at vold level for master-slave communication.
1829285	vxconfigd coredumps while assigning unique native name to a disk.
1675599	Memory leaks in DDL and ASLs.
1468885	The vxbrk_rootmir script does not complete and is hanging after invoking vxprivutil.
1431223	vradmin syncvol and syncrvg does not work if the remote diskgroup and vset name are specified when synchronizing vsets.

Table 1-12   Veritas Volume Manager fixed issues *(continued)*

Incident	Description
1426480	VOLCVM_CLEAR_PR() ioctl does not propogate the error returned by DMP to the caller.
1291519	After vradmin -s migrate is used twice, vrstat stops printing stats.
1192166	vxdg -n [newdg] deport [origdg] causes memory leak.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Volume Manager in this release.

Table 1-13   Veritas Volume Manager 5.1 SP1 RP1 fixed issues

Fixed issues	Description
1426480	The volcvm_clear_pr() ioctl now propogates the error returned by DMP to the caller.
1829285	vxconfigd no longer dumps core while assigning a unique native name to a disk.
1869002	Introduced a Circular buffer at the vold level for master-slave communication.
1940052	vxconfigd no longer hangs on the master after removing the HBA alias from the zone and node leave followed by join
1959513	The -o noreonline option of a diskgroup import is now propogated to slave nodes.
1970560	vxconfigd no longer dumps core on the master node when vxconfigd on a passive slave dies and command shipping is in progress.
2015467	Improved performance for NetBackup 6.5.5 on Veritas Storage Foundatoin 5.1 VxVM mapping provider.
2038928	Added support for creating and using pre-5.1 SP1 release diskgroups on CDS-initialized disks.
2062190	vxrootadm : split/join operation fails when there is a rvg present in the rootdg/backupdg

Table 1-13   Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2080730	The vxvm exclude file and vxdmp exclude file contents are now consistent after updating the files using the `vxdiskadm` command and `vxdmpadm` command.
2082450	vxdisk resize should output more meaningful error message
2088007	possibility of reviving only secondary paths in dmp_revive_paths()
2105547	tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations
2125306	Fixed a few issues related to loading the HBA API library and the `vxinstall` script.
2129477	`vxdisk reclaim` no longer fails after resize.
2129989	EVA ASL should report an error message if pref_bit is not set for a LUN
2133503	Renaming enclosure results in dmpevents.log reporting 'mode for Enclosure has changed from Private to Private'
2148682	while shipping a command node hangs in master selection on slave nodes and master update on master node
2152830	In a multi-level clone disks environment, a regular diskgroup import is now handled properly, and in the case of a diskgroup import failure, the correct error message is now displayed.
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core.
2160199	An upcoming master can now import a shared diskgroup, which allows the master takeover to succeed.
2166682	checks needed to make sure that a plex is active before reading from it during fsvm mirror read interface
2172488	FMR2 restore doesn't sync the existing snapshot mirrors
2179479	The flags on a disk are no longer incorrectly set as "error" even after running the `vxdisk scandisks` command after creating a PV and volume group.
2181631	striped-mirror volume cannot be grown across sites with -oallowspansites w/ DRL
2183984	system panic in dmp_update_stats() routine

Table 1-13   Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
2188590	an ilock acquired by a SLAVE node for a read on a DCL object can lead to IO hang when the node becomes MASTER before reading is done
2191693	`vxdmpadm native list` command now displays output and error messages.
2194492	VxVM-ASM co-existence enablement
2199496	Fixed a data corruption issue with the "site mirror" Campus Cluster feature.
2199836	The system with the root volume group and DMP native support enabled now successfully boots and mounts.
2200670	The `vxattachd` command can now recover disks if even if the disk group is not imported.
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault
2215216	`vxkprint` now reports TP-related values.
2220926	The `vxprivutil -D set` *attr* command no longer causes the `vxprivutil` command to hang.
2226813	Rlinks no longer remain disconnected with the UDP protocol if data ports are specified.
2227923	Renaming an enclosure is now persistent.
2234844	An `asm2vxfs` conversion with Linux partitions no longer fails.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP2

Table 1-14 describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP2.

Table 1-14   Veritas Volume Manager 5.1 SP1 RP2 fixed issues

Fixed issues	Description
2484685	Race between two vol_subdisk sios while doing done processing which causes one thread to free sio_fsvm_priv before other thread accesses it

Table 1-14   Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2480600	I/O permanent hung on master node when IO size larger than 512K, and 32+ threads write in parallel
2440349	DCO volume may grow into any 'site' even when 'alloc=site:xxxx' is specified by a list of 'site' to be limited
2431470	vxpfto uses DM name when calling vxdisk, but vxdisk will match DA name first and thus cause corruption
2431423	CVR: Panic in vol_mv_commit_check after I/O error on DCM
2428875	I/O on both nodes (wait for the DCM flush started), and crash the slave node, lead to the master reconfiguration hang
2428631	Allow same fence key to be used for all Disk groups
2425722	vxsd move operation failed for disk size greater than or equal to 2TB
2425551	IO hung for 6 mintues when reboot the slave node, if there is I/O on both master and slave
2424833	while autosync_deport#2 primary logowner hits ted assert nmcom_send_msg_tcp
2422058	The VxVM diskgroup can NOT import with I/O fencing enabled of both dmp and raw mode
2421067	Vxconfigd hung in both nodes of primary
2419348	DMP panic: race between dmp reconfig and dmp pass through ioctl
2413904	Multiple issues are seen while performing Dynamic LUN reconfiguration
2411698	VVR:iohang: On I/O to both master and slave
2410845	Lots of 'reservation conflict' messages seen on 51SP1RP1P1 clusters with XIV arrays
2408771	vxconfigd does not scan and discover all the storage device; some storage devices are skipped
2407192	Application I/O hangs because of race between CVM reconfiguration and Log-owner change protocol
2406292	Panic in vol_subdisksio_delete()
2400654	Stale arrayinfo file can cause vxdmpadm commands to hang

Table 1-14   Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2400076	vxconfigd produced kernel panic when you run "vxinstall" command
2396293	I/Os loaded, sanboot failed with vxconfigd core dump
2388725	Panic in dmp_get_dmpsymbols when attempting to load an APM
2387993	While testing including/excluding libvxppso vxconfigd goes into disabled mode
2386120	Enhancement request to add diagnostic logging to help triage a CVM master takeover failure situation
2385694	IO hung if the slave node rebooted
2385680	vol_rv_async_childdone+1147
2383158	VVR: vxio panic in vol_rv_mdship_srv_done+680
2379029	Changing of enclosure name is not working for all devices in enclosure
2369786	VVR:A deadloop about NM_ERR_HEADR_IO
2369177	DDL: do_diskio function should be able to handle offset greater than 2 TB
2365951	Growto failing with error V-5-1-10128 Unexpected kernel error in configuration update
2364253	VVR: Kernel memory is leaked on VVR secondary while using SO snapshots
2359814	vxconfigbackup doesn't handle errors well
2358321	Remove usage of __invalidate_device() from VxVM Symbol is no longer in kABI whitelist
2357798	CVR:Memory leak due to unfreed vol_ru_update structure
2357507	In presence of large number of NR (Not-Ready) devices, server panics due to NMI triggered and when DMP continuously generates large no of path disable/enable events
2356744	VxVM script daemons should not allow its duplication instance in itself
2349352	During LUN provisioning in single path IO mode environment a data corruption is observed
2346470	Excluding and including a LUN in a loop triggers a huge memory leak

Table 1-14   Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2337694	TP "vxdisk -o thin list" showing size 0 for over 2TB LUNs on RHEL5
2337353	vxdmpadm include vxvm dmpnodename=*emcpower#* includes all excluded dmpnodes along with the requested one
2334534	In CVM environment, vxconfigd level join is hung when Master returns error "VE_NO_JOINERS" to a joining node and cluster nidmap is changed in new reconfiguration
2323925	If rootdisk is encapsulated and if install-db is present, clear warning should be displayed on system boot
2322752	Duplicate DA records seen for NR devices upon restart of vxconfigd
2320917	vxconfigd core dump and lost dg config after removing volume and disk on thin reclaim LUN
2317703	Vxesd/Vxconfigd leaks file descriptors
2316297	After applying 51SP1RP1 error message "Device is in use" appears during boot time
2299670	Disk Groups created on EFI LUNs do not auto import at boot time using VxVM version 51SP1 and later
2286559	kernel heap corruption detected panic after array controller reboot
2263317	CLONE: Diskgroup import with dgid needs to be clearly documented in manual for the case in which original dg was destroyed and cloned disks are present
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated
2255182	Handling misconfiguration of CLARiiON array reporting one failovermode value through one HBA and different from other HBA
2253970	Support per-disk maxiosize for private region I/Os
2253552	Leak in vxsfdefault_parsey at function vxsf_getdefault (*val)
2249113	vol_ru_recover_primlog_done return the same start address to be read from SRL, if the dummy update is greater than MAX_WRITE
2248730	vxdg import command hangs as vxrecover daemon (spawned by vxdg) doesn't close standard error stream

Table 1-14   Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2242268	panic in voldrl_unlog
2240056	vxdg move' transaction not completing and backups fail
2237089	vxrecover might start the recovery of data volumes before the recovery of the associated cache volume is recovered
2234821	etrack 1946267 - DMP can't detect the re-enabled os device status does not work on RHEL5
2232411	supporting NetApp Metro Cluster
2228531	cvm master vxconfigd process hung in vol_klog_lock()
2218470	Some of the VxVM init scripts need to be compliant to the Linux Standard Base
2205108	SVS 51SP1: vxconfigd clubbing all luns in a single dmpnode
2204752	Multiple VM commands succeed but throw "GPT entries checksum mismatch" error message for hpdisk format
2200670	vxattachd does not recover disks if disk group is not imported
2197254	While creating volumes on thinrclm disks, the option "logtype=none" does not work with vxassist command
2196918	Snapshot creation with cachesize fails, as it doesn't take into account diskgroup alignment
2196480	The disk initialization failed due to wrong number of cylinders reported in devintf_disk_geom_raw() gotten from raw geometry
2194685	vxconfigd daemon core dump during array side switch ports disable and re-enable
2193429	IO policy not getting preserved when vold is restarted and migration from one devlist to other is taking place
2190020	SUSE complains dmp_deamon applying 1m continuous memory paging is too large
2179259	DMP SCSI bypass needs to be enhanced to handle I/O greater than 2 TB

Table 1-14   Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2165394	CLONE: dg imported by selecting wrong disks After destroying original dg, when try to import clone devices without useclonedev option with dgname, then it import dg with original disks
2154287	Improve handling of Not-Ready(NR)devices which are triggering "VxVM vxdmp V-5-3-1062 dmp_restore_node: Unstable path" messages
2152830	In multilevel clone disks environment, regular DG import should be handled properly and in case of DG import failure, it should report correct error message
2144775	Failoverpolicy "local" is not getting preserved after upgrade from 51RP1/Sles10Sp2 to 51Sp1/Sles10Sp3
2139179	SSB check invalid when lun copy
2094672	CVR: vxconfigd on master hangs while reconfig is running in cvr stress with 8 users
2033909	In SF-RAC configuration, IO hung after disable secondary path of A/PG array Fujitsu ETERNUS3000
1791397	VVR:RU thread keeps spinning sending START_UPDATE message repeatedly to the secondary
1675599	Memory leaks in DDL and ASLs

LLT, GAB, and I/O fencing: Issues fixed in 6.0

Table 1-15 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-15   LLT, GAB, and I/O fencing fixed issues

Incident	Description
2515932	[GAB] gabconfig ioctl behaviour changed to return EALREADY if GAB is already configured.
2495020	[Fencing] vxfend does not terminate if you run the `vxfenswap` command to change the fencing mode from 'scsi3' to 'customized', and chooses to rollback when vxfenswap prompts for confirmation.
2442402	[LLT] Reduce lltd CPU consumption by reducing the wakeup calls.

Table 1-15 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
2437022	[Fencing] Fails to run the `vxfenswap` command to the same diskgroup when the disk policy changed.
2426664	[Fencing] vxfend does not terminate when you run the `vxfenswap` command to migrate from the customized mode to the scsi3 mode.
2411652	[GAB] Add a check in GAB for MAX message size of 64KB before enqueuing the message.
2386325	[Fencing] Fencing configuration fails and vxfenadm prints same serial number for all LUNs which have more than 96 bytes of SCSI Inquiry data in page 0x83.
2369742	[Fencing] Once vxfenconfig -c with a particular mode (say customized) has returned EFAULT ("1036 Unable to configure..."), all subsequent runs of vxfenconfig -c with a different mode (say scsi3) fail with error EBADMSG ("1050 Mismatched modes...").
2351011	[Fencing] The vxfenswap utility fails to accurately check for the exit status of the vxfenconfig commands run on the other nodes in the background. This may lead to the vxfenswap utility appearing indefinitely hung if the vxfenconfig process does not succeed for any reason.
2337916	[Fencing] Fencing shutdown script does not retry stopping the fencing module if fencing fails to unconfigure because of clients being registered.
2311361	[Fencing] Fencing details are printed in the engine logs every five minutes if fencing is running and the CoordPoint resource is configured.
2253321	[Fencing] Fencing fails to start if any of the coordination points is unavailable at the startup time.
2252470	[Fencing] Provide options to force the fencing library to obtain serial numbers using standard inquiry or extended inquiry using a variety of ID types.
2218448	[VxCPS] The cpsadm command fails if LLT is not installed or configured on a single-node cluster which hosts the CP server.
2209664	[VxCPS] Configuring fencing is successful with three disks even when single_cp=1 and the formatting of warning messages aer required in vxfend_A.log.
2209144	[VxCPS] There is syntax error while unconfiguring CP server using the configure_cps.pl script.

Table 1-15   LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
2203070	[Fencing] Failed to configure fencing on a 64-node cluster, fencing comes up only on first 33 nodes.
2178126	[GAB] GAB fails to start if it is unable to allocate memory in atomic manner in low memory situations, typically in under-provisioned virtual machine setups.
2161816	[Fencing] Preferred fencing does not work as expected for large clusters in certain cases if you have configured system-based or group-based preferred fencing policy.
2139883	[GAB] On RHEL5 Update 5 and later, messages similar to the following are repeatedly seen on the console: `INFO: task gablogd:22812 blocked for more than 120 seconds.` `"echo 0 > /proc/sys/kernel/hung_task_timeout_secs"` disables this message.
2112742	[VxCPS] Server-based I/O fencing fails to start after configuration on nodes with different locale settings.
2100896	[Fencing] There is failure message even the migration from server-based to disk-based using vxfenswap succeeded.
2085941	[VxCPS] Coordination point server (CP server) binds only to a single virtual IP and listens on the same. Application clusters cannot access the CP server if it fails to establish connection to this virtual IP. Therefore, if the connection fails because of the subnet in which the virtual IP of the CP server exists, you cannot access the CP server even if there is another subnet through which the client can connect to the CP server over a different IP.
2076240	[VxCPS] When you reconfigure an application cluster that uses server-based I/O fencing (customized fencing mode), the installer does not remove the application cluster information from the CP servers before the reconfiguration. As a result, if you reconfigure the application cluster and choose to configure I/O fencing in customized mode using the same CP servers, then reconfiguration of server-based fencing for the application cluster fails.
1973713	[Fencing] The agent XML files are missing for CP server agent.

Known issues

This section covers the known issues in this release.

See the corresponding Release Notes for a complete list of known issues related to that product.

See "Documentation" on page 115.

Issues related to installation

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 (2591399)

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
    ```

EULA changes (2161557)

The locations for all EULAs have changed.

The English EULAs now appear in /*product_dir*/EULA/en/*product_eula.pdf*

The EULAs for Japanese and Chinese now appear in those language in the following locations:

The Japanese EULAs appear in /*product_dir*/EULA/ja/*product_eula.pdf*

The Chinese EULAs appear in /*product_dir*/EULA/zh/*product_eula.pdf*

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

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, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/openv`), then while upgrading to SF 6.0, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure RPMs `VRTSpbx`, `VRTSat`, and `VRTSicsco`. This causes NetBackup to stop working.

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

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

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

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

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

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

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

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

The VRTSacclib RPM is deprecated (2032052)

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

- Fresh installs: Do not install VRTSacclib.

- Upgrade: Ignore VRTSacclib.

- Uninstall: Ignore VRTSacclib.

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:

Refer to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for the section, "vxconfigd daemon recovery."

Upgrading from Veritas Storage Foundation 5.1 Service Pack 1 Rolling Patch 2 to 6.0 with rootability enabled fails (2581313)

Upgrading from Veritas Storage Foundation (SF) 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 to 6.0 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 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 while using an encapsulated root disk

1   Encapsulate the root disk.
2   Reinstall the nash utility.
3   Upgrade to the SF 6.0 release.

During upgrade from 5.1SP1 to 6.0 with an encapsulated root disk, splitting the root mirror fails if the target disk group name is used by a deported disk group (2280560)

During an upgrade from SFHA 5.1 SP1 to SFHA 6.0 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.

Work-around:

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.

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.

sfmh discovery issue when you upgrade your Veritas product to 6.0 (2622987)

If a host is not reporting to any management server but sfmh discovery is running before you upgrade to 6.0, sfmh-discovery may fail to start after the upgrade.

Workaround:

If the host is not reporting to VOM, stop sfmh-discovery manually before upgrading to 6.0 by executing the following command on the managed host:

```
/opt/VRTSsfmh/adm/vxvmdiscovery-ctrl.sh stop
```

Incorrect server names sometimes display if there is a clock synchronization issue (2627076)

When you install a cluster with the Web-based installer, you choose to synchronize your systems with an NTP server due to a clock synchronization issue, you may see the NTP server name in messages instead of your server names.

Workaround:

Ignore the messages. The product is still installed on the correct servers.

When you uninstall CommandCentral Storage Managed Host from a system where Veritas Storage Foundation 6.0 is installed, SF 6.0 reconfiguration or uninstallation fails (2631486)

On a system where Veritas Storage Foundation (SF) 6.0 is installed, if you uninstall CommandCentral Storage (CCS) Managed Host (MH) using the installer script from the CCS media, the installer script removes the contents of `/opt/VRTSperl`. As a result, SF 6.0 reconfiguration or uninstallation using `/opt/VRTS/install/install_sf_product_name` or `/opt/VRTS/install/uninstall_sf_product_name` fails, because the installer script removed the contents of `/opt/VRTSperl`.

Workaround: To uninstall CCS MH from a system where SF 6.0 is installed, before you perform the uninstallation, perform the procedure in the following CCS TechNote:

http://www.symantec.com/business/support/index?page=content&id=HOWTO36496

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.

Issues related to LLT

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 (1788315)

With each received packet, LLT increments the following variables:

■ recvcnt (increment by one for every packet)

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

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

This does not impact the LLT functionality.

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

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

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

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

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

Issues related to GAB

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

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

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

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

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

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.

Issues related to I/O fencing

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

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 nodes' and users' information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster nodes' and users' 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@node1,
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, 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 atleast 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: No workaround.

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 in secure mode (2478502)

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 steps on all the nodes of the CP server:

- Rename cpsadm to cpsadmbin.

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

- 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 "$@"
  ```

- Provide the following permissions to the new file:

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

Veritas Storage Foundation known issues

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

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 VM process may hang during the configuration phase.

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

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. Oracle 11gR2 release may work on a pure IPv6 enviroment, but it has not been tested or released yet.

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 (1880622)

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.

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

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

Workaround

Reboot the system or run the following command.

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

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

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

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

vxrestored daemon fails to restore disabled paths (1663167)

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

Workaround:

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

To enable the mpt_disable_hotplug_remove tunable

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

```
options mptsas mpt_disable_hotplug_remove=0
```

2   Rebuild the initrd image:

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

3   Reboot the system.

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

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

Workaround:

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

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

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

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

Workaround:

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

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

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

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

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

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

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

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

Workaround:

To recover from this situation

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

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

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

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

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

Root disk encapsulation issue (1603309)

Encapsulation of root disk will fail if it has been assigned a customized name with vxdmpadm(1M) command. If you wish 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.

vxdisk -f init can overwrite some of the public region contents (1190117)

If a disk was initialized by a previous VxVM version or defined with a smaller private region than the new default of 32 MB, then the public region data will be overridden.

Workaround:

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

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

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

```
Cannot setup space
```

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

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

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

Workaround: There is no workaround for this issue.

Co-existence check might fail for CDS disks

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

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

In layouts where GPT is used for initialization (typically, when the disk size is currently greater than or had earlier exceeded 1 TB), the AIX co-existence label is placed at sector 55 and VxVM ID block (also known as HP co-existence label) is placed at sector 64. Consequently, AIX utilities would not be able to recognize a cdsdisk initialized using GPT to be a valid VxVM disk. Symantec is working with IBM and third party OEMs to enhance the co-existence check in these utilities.

Workaround: There is no workaround for this issue.

Removing a volume from a thin LUN in an alternate boot disk group triggers disk reclamation (2080609)

If you remove a volume from an alternate boot disk group on a thin LUN, this operation triggers thin reclamation, which may remove information required for the disk to be bootable. This issue does not affect the current boot disk, since VxVM avoids performing a reclaim on disks under the bootdg.

Workaround: If you remove a volume or plex from an alternate boot disk group with the vxedit command, specify the -n option to avoid triggering thin reclamation. For example:

```
# vxedit -g diskgroup -rfn rm volumename
```

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 (2082414)

The Veritas Volume Manager (VxVM) 6.0 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, 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. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-16 shows the Hitachi arrays that have new array names.

Table 1-16   Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC
TagmaStoreUSPV	Hitachi_USP-V

Table 1-16   Hitachi arrays with new array names *(continued)*

Previous name	New name
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. 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

A controller can remain disabled due to udev device removal after loss of connectivity to some paths on RHEL6 and SLES11 (2697321)

The issue may occur with NetApp LUNs in ALUA mode. When a device fails with a `dev_loss_tmo` error, the operating system (OS) device files are removed by `udev`. After this removal, a controller will remain in the disabled state until a reboot is run on the host. To avoid this issue, use the following workaround.

Workaround

To create the new rules file

1 Create the file `/etc/udev/rules.d/40-rport.rules` with the following content line:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports",ACTION=="add",
RUN+="/bin/sh -c'echo 20 > /sys/class/fc_remote_ports/%k/
fast_io_fail_tmo;echo 864000 >/sys/class/fc_remote_ports/%k/
dev_loss_tmo'"
```

2 Reboot the system.

3 If new LUNs are dynamically assigned to the host, run the following command:

```
# udevadm trigger --action=add --subsystem-match=fc_remote_ports
```

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.

vxsnap addmir command sometimes fails under heavy I/O load (2441283)

The `vxsnap addmir` command sometimes fails under heavy I/O load and produces multiple errors.

Workaround: Rerun the `vxsnap addmir` command.

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.

The vxdiskunsetup operation fails the first attempt on EMC powerpath devices (2424845)

Performing `vxdiskunsetup` for the first time on EMC powerpath devices displays an error "Internal Configuration daemon error : disk destroy failed."

Workaround: Retry `vxdiskunsetup` using the same command to resolve the issue.

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 vxassist maxsize option fails to report the maximum size of the volume that can be created with given constraints when the disk group has the siteconsistent flag set (2563195)

The `vxassist maxsize` option fails to report the maximum size of volume that can be created with given constraints when the disk group has the siteconsistent flag set. The following error is reported:

vxassist -g dgname maxsize
VxVM vxassist ERROR V-5-1-752 No volume can be created within the given
constraints

Workaround:

Specify the size explicitly to the `vxassist make` command.

Encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path names are not the same (2607706)

The encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path name are not the same.

For example:

Dmpnode:sdh

Paths: sda sdb

Work-around:

Before running the encapsulation command (`vxencap`), run the following command:

vxddladm assign names

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.

Work-around:

To work around this issue

1 Restore storage connectivity.

2 Deport the disk group.

3 Import the disk group.

After rebooting the array controller for CX4-240-APF array, I/O errors occur on shared file systems (2616315)

For Linux hosts, rebooting the array controller for a CX4-240-APF array may result in I/O errors on shared file systems.

Work-around:

To work around this issue

◆ Set the tunable parameter dmp_lun_retry_timeout to 120 seconds before rebooting the array controller.

```
# vxdmpadm settune dmp_lun_retry_timeout=120
```

On SLES11, sometimes encap fails in partition phase during first reboot (2598859)

On SLES11, sometimes the encapsulation of a root disk fails during the re-partitioning phase, with the following error message:

```
Starting vxvm-reconfig
The Volume Manager is now reconfiguring (partition phase)...
Volume Manager: Partitioning sda as an encapsulated disk.
VxVM vxedpart ERROR V-5-1-13204 partition pre-verification failed : Device
resource busy

VxVM vxedroot ERROR V-5-2-3684 re-partitioning of rootdisk failed!
The partitioning of sda failed.
The following disks failed encapsulation:
sda
All changes the Volume Manager has made during this reconfiguration will be
reversed.
undo sda....
```

There is no work-around for this issue.

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
```

Oracle ASM support with VxVM on SLES 10 requires symlink for raw devices to be created (2556467)

For Oracle ASM support to work with Veritas Volume Manager (VxVM) on SLES 10, a symlink is required. Otherwise, Dynamic Multi-Pathing (DMP) does not create raw devices on reboot.

Workaround:

For Oracle ASM support to work with VxVM on SLES 10, before installing VxVM, create a symlink to `/usr/sbin/raw` under `/bin` using the following command:

```
# ln -s /usr/sbin/raw /bin/raw
```

System booting can take longer than expected (2486301)

System booting can sometimes take longer than expected due to hwpath generation.

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 (2560360)

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
```

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.

Upgrading from disk layout Version 8 to 9 on a file system with partitioned directories and Storage Checkpoints can return with a read-only file system error message (2583201)

Upgrading from disk layout Version 8 to 9 on a file system with partitioned directories and Storage Checkpoints can return with a read-only file system error message. The issue with partitioned directories occurs because disk layout Version 9 has a new hash function. The issue with Storage Checkpoints occurs because the Storage Checkpoints are marked as read-only during the upgrade.

Workaround: Before upgrading a VxFS file system with disk layout Version 8 to Version 9, use the following procedure to avoid this error message.

To avoid the system error message

1   Disable the partitioned directories feature if the feature is enabled by setting the pdir_enable tunable to 0.

See the vxtunefs(1M) manual page.

2   Remove all Storage Checkpoints before the upgrade.

See the fsckptadm(1M) manual page.

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 (2655788)

If you use cross-platform data sharing to convert a file system that has more than 32k nlinks, the conversion process does not update the `vx_maxlink` and `maxlink_enable` tunables on the target file system.

Workaround: After the cross-platform data sharing conversion completes, validate the values of the `vx_maxlink` and `maxlink_enable` tunables. If the file system had more than 32k nlinks before the conversion, ensure that the these tunables are updated on the target file system before mounting the file system.

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.

System hang when using ls, du and find (2598356)

The system sometimes hangs when using the `ls`, `du`, or `find` commands. The hang occurs in the following stack:

```
schedule_timeout
vx_iget
vx_dirlook
vx_lookup
do_lookup
do_path_lookup
```

Workaround: There is no workaround for this issue.

Expanding a 100% full file system can cause a panic (2599590)

Expanding a 100% full file system can cause a panic with the following stack trace:

```
bad_kern_reference()
$cold_vfault()
vm_hndlr()
bubbledown()
vx_logflush()
vx_log_sync1()
vx_log_sync()
vx_worklist_thread()
kthread_daemon_startup()
```

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.

vfradmin abortjob does not kill the file replication process (2527916)

The `vfradmin abortjob` command does not kill the file replication process and may leave the checkpoints mounted.

Workaround: There is no workaround for this issue.

vradmin syncvol command compatibility with IPv6 addresses (2075307)

The `vradmin syncvol` command does not work with the compressed form of IPv6 addresses. In IPv6 environments, if you run the `vradmin syncvol` command and identify the target host using compressed form of the IPv6 address, the command fails with following error message:

```
# vradmin -s -full syncvol vol1 fe80::221:5eff:fe49:ad10:dg1:vol1
VxVM VVR vradmin ERROR V-5-52-420 Incorrect format for syncvol.
```

Also, if you run the `vradmin addsec` command and you specify the Secondary host using the compressed IPv6 address, the `vradmin syncvol` command also fails – even if you specify the target as `hostname`.

Workaround: When you use the `vradmin addsec` and `vradmin syncvol` commands, do not specify compressed IPv6 addresses; instead, use hostnames.

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

When VVR is used for replicating shared disk groups in an SFCFS or SFRAC environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg` commands:

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

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

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

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

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

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

While vradmin commands are running, vradmind may temporarily lose heart beats (2162625, 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 (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

1   Pause or stop the applications.

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

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

3   Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

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

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

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

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.

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

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

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

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

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

Workaround:

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

1   Pause or stop the applications.

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

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

3   Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

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.

Incorrect error message if wrong host name is provided (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.

Workaround

Provide the name of a host that has the repository database, with the -r option of vxsfadm.

FlashSnap validate reports snapshot unsplittable (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.
```

Workaround

Ensure that snapshot plexes for data volumes and snapshot plexes for archive log volumes reside on separate set of disks.

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.

`dbed_vmclonedb` **ignores new clone SID value after cloning once (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.

Workaround

You can use one of the following workarounds:

■ After the snapshot is resynchronized, delete the snapplan using the `dbed_vmchecksnap -o remove` command. You can then use a new clone SID by creating a new snapplan, which may have the same name, and using the snapplan for taking more snapshots.

- Use the `vxsfadm` command to take the snapshot again and specify the clone SID with the snapshot operation so that the clone operation can be done with the new clone SID.

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.

User authentication fails (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>
```

Reattempting the operation fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0372 AT broker failed to start:
```

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.

Workaround

If you have not done authentication setup, set umask to a less strict value before running the `sfae_auth_op -o setup` or `sfae_auth_op -o import_broker_config` commands.

To set umask to a less strict value

◆   Use the command:

    # **umask 022**

If you have already done authentication setup, perform the following steps.

To resolve the problem if you have already done authentication setup

1   Shut down the authentication broker, if it is running.

    # **/opt/VRTSdbed/at-broker/bin/sfaeatd.sh stop**

2   Change the permissions for files and directories that are required to be
    readable by non-root users.

    # **chmod o+r /etc/vx/vxdbed/admin.properties**
    # **chmod o+rx /var/vx/vxdba/auth/users**
    # **find /opt/VRTSdbed/at-broker -type d -exec chmod o+rx {} \;**

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 (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.

When upgrading from SFHA version 5.0 to SFHA 6.0 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 (1764885)

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

Workaround

There is no workaround for this issue.

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.

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 115.

Veritas Volume Manager software limitations

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

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 1-17 describes the DMP tunable parameters and the new values.

Table 1-17   DMP settings for NetApp storage attached environment

Parameter name	Definition	New value	Default value
dmp_restore_internal	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_internal=60**

    # **vxdmpadm settune dmp_path_age=120**

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

    # **vxdmpadm gettune dmp_restore_internal**

    # **vxdmpadm gettune dmp_path_age**

Veritas File System software limitations

The following are software limitations in the 6.0 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.

Replication in a shared environment

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

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

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

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

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

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

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

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 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.

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 if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 6.0: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 6.0.

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

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

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

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

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

Documentation errata

The following sections cover additions or corrections for Document version: 6.0.5 of the product documentation. These additions or corrections may be included in later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).

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

See "Documentation" on page 115.

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

Veritas Storage Foundation Administrator's Guide

The following errata applies to the *Veritas Storage Foundation and High Availability Administrator's Guide*.

"VxFS Version 9 disk layout" section in the "Disk layout" appendix

The following text should be deleted:

The Version 8 disk layout supports group quotas.

See "About quota files on Veritas File System" on page x.

Documentation

Product guides are available in the PDF format on the software media in the `/product_name/docs` 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 1-18 lists the documentation for Veritas Storage Foundation and High Availability.

Table 1-18   Veritas Storage Foundation and High Availability documentation

Document title	File name
Veritas Storage Foundation and High Availability Release Notes	sfha_notes_60_lin.pdf
Veritas Storage Foundation and High Availability Installation and Configuration Guide	sfha_install_60_lin.pdf

Table 1-19 lists the documents for Veritas Cluster Server.

Table 1-19   Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_60_lin.pdf
Veritas Cluster Server Release Notes	vcs_notes_60_lin.pdf
Veritas Cluster Server Administrator's Guide	vcs_admin_60_lin.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_60_lin.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev_60_unix.pdf
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_60_lin.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_60_lin.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_60_lin.pdf

Table 1-20 lists the documentation for Veritas Storage Foundation.

Table 1-20   Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_60_lin.pdf
Veritas Storage Foundation Installation Guide	sf_install_60_lin.pdf
Veritas Storage Foundation Administrator's Guide	sf_admin_60_lin.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sf_adv_ora_60_lin.pdf
Veritas File System Programmer's Reference Guide	vxfs_ref_60_lin.pdf

Table 1-21 lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-21   Veritas Storage Foundation and High Availability Solutions products documentation

Document title	File name
Veritas Storage Foundation and High Availability Solutions Solutions Guide	sfha_solutions_60_lin.pdf
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	sfha_virtualization_60_lin.pdf
Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide	sf_replication_admin_60_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
```