Symantec™ Storage Foundation Cluster File System High Availability 6.2 Release Notes - Linux

January 2016

VERITAS™

Symantec™ Storage Foundation Cluster File System High Availability Release Notes

Product version: 6.2

Document version: 6.2 Rev 5

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apj@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

Storage Foundation Cluster File System High Availability Release Notes

This document includes the following topics:

About this document

This document provides important information about Symantec Storage Foundation Cluster File System High Availability (SFCFSHA) version 6.2 for Linux. Review this entire document before you install or upgrade SFCFSHA.

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

This is "Document version: 6.2 Rev 5" of the *Symantec Storage Foundation Cluster File System 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 website at:

https://sort.symantec.com/documents

Component product release notes

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

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

`/docs/product_name`

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

This release includes the following component product release notes:

- *Symantec Storage Foundation Release Notes* (6.2)
- *Symantec Cluster Server Release Notes* (6.2)

About Symantec Storage Foundation Cluster File System High Availability

Symantec Storage Foundation Cluster File System High Availability by Symantec extends Symantec Storage Foundation to support shared data in a storage area network (SAN) environment. Using Storage Foundation Cluster File System High Availability, multiple servers can concurrently access shared storage and files transparently to applications.

Symantec Storage Foundation Cluster File System High Availability also provides increased automation and intelligent management of availability and performance.

Symantec Storage Foundation Cluster File System High Availability includes Symantec Cluster Server, which adds high availability functionality to the product.

The Symantec File Replicator feature can also be licensed with this product.

To install the product, follow the instructions in the *Symantec Storage Foundation Cluster File System High Availability Installation Guide*.

For information on high availability environments, read the Symantec Cluster Server documentation.

About Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a website 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 and generate an Installation and Upgrade custom report. ■ List patches by product or platform, and in the order they need to be installed. Display and download the most recent patches or historical patches. ■ Display Array Support Library (ASL) details by vendor, platform, or Storage Foundation and High Availability (SFHA) version. ASLs make it easier to manage arrays that are connected to SFHA-based servers. ■ List VCS and ApplicationHA agents, documentation, and downloads based on the agent type, application, and platform.
Identify risks and get server-specific recommendations	■ Analyze your servers for potential environmental risks. Generate a Risk Assessment custom report with specific recommendations about system availability, storage use, performance, and best practices. ■ Display descriptions and solutions for thousands of Symantec error codes.

Improve efficiency	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), documentation, product releases, Hardware Compatibility Lists (HCLs), and VCS/ApplicationHA agents. ■ Quickly gather installed Symantec product and license key information from across your production environment. Generate a License/Deployment custom report that includes product names, versions, and platforms, server tiers, Symantec Performance Value Units (SPVUs), and End of Service Life dates. ■ List and download Symantec product documentation including product guides, manual pages, compatibility lists, and support articles. ■ Access links to important resources on a single page, including Symantec product support, SymConnect forums, customer care, Symantec training and education, Symantec FileConnect, the licensing portal, and my.symantec.com. The page also includes links to key vendor support sites. ■ Use a subset of SORT features from your iOS device. Download the application at: https://sort.symantec.com/mobile

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/TECH225259
- For the latest patches available for this release, go to: https://sort.symantec.com/
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware, visit the following URL: http://www.symantec.com/docs/TECH211575

- The software compatibility list summarizes each Storage Foundation and High Availability (SFHA) Solutions product stack and the product features, operating system versions, and third-party products it supports. For the latest information on supported software, visit the following URL:
  http://www.symantec.com/docs/TECH225258

Note: Before you install or upgrade SFHA Solutions products, review the current compatibility lists to confirm the compatibility of your hardware and software.

Changes introduced in 6.2

This section lists the changes in Symantec Storage Foundation Cluster File System High Availability 6.2.

Changes related to installation and upgrades

The product installer includes the following changes in 6.2.

VxVM SmartIO support for SFCFSHA installations

VxVM SmartIO is supported for SFCFSHA installations. When SmartIO is enabled on multiple nodes, Group Lock Manager (GLM) library keeps cache on each node coherent.

See the *Symantec™ Storage Foundation and High Availability Solutions SmartIO for Solid State Drives Solutions Guide* for details.

Connecting to the SORT website through a proxy server

The product installer connects to the Symantec Operations Readiness Tools (SORT) website for several purposes, such as downloading the latest installer patches, and uploading installer logs. Deployment Server can connect to SORT to automatically download Maintenance or Patch release images. In this release, before running the product installer or Deployment Server, you can use the following proxy settings to connect to SORT through proxy servers:

```
# https_proxy=http://proxy_server:port
# export https_proxy
# ftp_proxy=http://proxy_server:port
# export ftp_proxy
```

Symantec Storage Foundation Cluster File System High Availability gets installed in secure mode by default

Symantec Storage Foundation Cluster File System High Availability gets installed in secure mode by default. You are advised to install SFCFSHA in secure mode to be able to control guest user access to secure clusters and encrypt communication between SFCFSHA components. You can choose the non-secure mode during installation; however, the product installer warns you during the installation with the following message:

```
Symantec recommends that you install the cluster
in secure mode. This ensures that communication between
cluster components is encrypted and cluster information
is visible to specified users only.
```

The upgrade from non-secure mode continues to happen in non-secure mode. The upgrade from secure mode advises you to control user access to secure clusters.

Package updates

The following lists the RPM changes in this release.

- The `VRTSlvmconv` RPM has been merged with the `VRTSvxvm` RPM.There is no separate RPM for `lvmconvert` now.
- The `VRTSvxvm` RPM adds dependency for bc -1.06.95-13.el7.x86_64, pcre-8.32-12.el7.i686 (pcre(x86-32)), and xz-libs-5.1.2-8alpha.el7.i686 (xz-libs(x86-32)) packages on RHEL 7 distribution. Newly required dependent RPM for `VRTSvxvm` is:

Package name	Version	Architecture
bc	Default version available with RHEL7	x86_64
pcre	Default version available with RHEL7	el7.i686
xz-libs	Default version available with RHEL7	el7.i686

For more information, see the *Installation Guide* for the complete list of RPMs for this release.

Support for centralized installations using the Deployment Server

The Deployment Server is a script that makes it easier to install or upgrade SFHA releases. The Deployment Server lets you store multiple release images in one central location and deploy them to systems of any supported UNIX or Linux platform (6.1 or later). Prior to 6.1, releases still require the same platform, architecture, distribution, and version of the operating system. You can use the Deployment Server if you want to install or upgrade multiple releases and or multiple platforms.

The Deployment Server lets you do the following as described in Table 1-1.

Table 1-1 Deployment Server functionality

Feature	Description
Install or Upgrade systems with Install Bundle and Install Template	■ Install or upgrade systems with an Install Bundle. ■ Install packages on systems based on the information stored in the Install Template.
Define or modify Install Bundles	Define or modify Install Bundles and save them using the Deployment Server.
Create Install Templates	Discover installed components on a running system that you want to replicate on new systems.
Connecting the Deployment Server to SORT using a proxy server	Use a proxy server, a server that acts as an intermediary for requests from clients, for connecting the Deployment Server to the Symantec Operations Readiness Tools (SORT) website.
Platform Filtering	On the Set Preference menu, choose Selected Platforms to filter the platforms that are currently being used in the deployment environment.

Note: The Deployment Server is available only for the script-based installer, not the web-based installer.

See the *Installation Guide* for more information.

Support for installation using the Red Hat Satellite server

You can install SFCFSHA using the Red Hat Satellite server. Red Hat Satellite is supported for Red Hat Enterprise Linux 6 (RHEL6) and Red Hat Enterprise Linux 7 (RHEL7). You can install RPMs and rolling patches on the systems which the Red Hat Satellite server manages.

In a Red Hat Satellite server, you can manage the system by creating a channel. A Red Hat Satellite channel is a collection of software packages. Using channels, you can segregate the packages by defining some rules.

Behavioral changes in RHEL 7 as compared with previous releases

Note the following behavioral changes in RHEL 7:

- The XFS file system is not supported for the Root Disk Encapsulation (RDE) feature.
  RDE is not supported if the root partition is mounted with the XFS file system.
- Enclosure-based naming (EBN) is not supported for RDE.
  RDE, mirroring, splitting, and joining operations on root disks are not supported if the naming scheme is set to EBN.

Support for setting up ssh and rsh connection using the pwdutil.pl utility

The password utility, `pwdutil.pl`, is bundled in the 6.2 release under the scripts directory. The users can run the `pwdutil.pl` utility to set up the `ssh` and `rsh` connection automatically.

Changes related to Symantec Storage Foundation Cluster File System High Availability (SFCFSHA)

Symantec Storage Foundation Cluster File System High Availability (SFCFSHA) includes the following changes in 6.2:

Auto deport of disabled FSS disk groups

In this release, when all the storage of a shared disk group is lost and the disk group has become disabled, the CVMVoldg agent in VCS will try to deport the disk group during clean and offline entry points. However, note that if there are any pending I/Os or open counts on volumes, then the disk group will not be deported and will remain in the disabled state.

Prior to this feature, the disk group used to remain in the disabled state and you had to manually deport it.

See the SFRAC or SFCFS admin guides for more information.

Support for Red Hat Enterprise Virtualization Environment (RHEV)

In RHEV environments, Symantec Storage Foundation can be configured as the backend storage for guest virtual machines. Symantec provides the `rhevadm` utility on RHEV Manager to configure storage for virtual machines. With SF as backend storage, you can leverage Flexible Storage Sharing (FSS) feature to commission commodity hardware in place of costlier storage arrays. Veritas Volume Replicator (VVR) and Veritas File Replicator (VFR) provide volume and file level replication which enables you to perform disaster recovery of virtual machines.

Support for Red Hat Enterprise Linux (RHEL) 7 platform

Support for RHEL 7 is added in this release. As part of this support, there is also an addition of new package VRTSveki. This package will be responsible for inter module communication across all kernel modules in the SFCFSHA stack.

Note: SmartIO with Oracle on RHEL 7 (and the Oracle plugin for SmartIO) is not supported.

Device name format changes in RHEL 7 environments after encapsulation

With RHEL 7, the format of volumes in the `/etc/fstab` file after root disk encapsulation has changed.

Table 1-2 lists the changes in RHEL 7 environments.

Table 1-2 Volume formats changes in RHEL 7 environments

Before RHEL 7	With RHEL 7
Volume format: `/dev/vx/dsk/bootdg/<volume>`	Volume format: `/dev/vx_dsk_bootdg_<volume>`

Table 1-2 Volume formats changes in RHEL 7 environments *(continued)*

Before RHEL 7	With RHEL 7

Contents of `/etc/fstab` file where the rootdisk has two partitions, namely, `/` and `swap`:
 	
`# cat /etc/fstab`
 	
`/dev/vx/dsk/bootdg/rootvol \`
`/ ext4 defaults 1 1`
`/dev/vx/dsk/bootdg/swapvol \`
`swap swap defaults 0 0`
 	
`#NOTE: volume rootvol (/) \`
`encapsulated partition sda1`
`#NOTE: volume swapvol (swap) \`
`encapsulated partition sda2`
	Contents of `/etc/fstab` file where the rootdisk has two partitions, namely, `/` and `swap`:
	
	`# cat /etc/fstab`
	
	`/dev/vx_dsk_bootdg_rootvol \`
	`/ ext4 defaults 1 1`
	`/dev/vx_dsk_bootdg_swapvol \`
	`swap swap defaults 0 0`
	
	`#NOTE: volume rootvol (/) \`
	`encapsulated partition sda1`
	`#NOTE: volume swapvol (swap) \`
	`encapsulated partition sda2`

Note: Though the format of the device names in the `/etc/fstab` has changed, there is no change in the output of the `mount` utility. The `mount` utility still displays the mounted volumes in the old format.

Collecting application and daemon core data for debugging

If a Storage Foundation application or daemon encounters a problem, it may produce a core file. This release introduces the `vxgetcore` script which lets you efficiently collect the core file, binary file, library files, and any related debugging information and generate a tar file. You can then send the tar file to Symantec Technical Support for analysis.

For more information, see the *Symantec Storage Foundation and High Availability Solutions Troubleshooting Guide*.

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 6.2:

Changes in array names for Fusion-io devices

Prior to this release, the generic array name, fusionio, was used for all Fusion-io devices. Starting in this release, the array name indicates the type of Fusion-io card. For example, the ioDrive cards display names such as fiodrive0_0.

Use the `vxdisk list` command to display the array name.

For example:

```
# vxdisk list

fiodrive0_0  auto:cdsdisk    -            -            online ssdtrim
fiodrive0_1  auto:cdsdisk    -            -            online ssdtrim
```

Layered volume enhancements for recovery and snapshots

In this release, a new enhancement is done for layered volumes so that when storage disconnection and subsequent reconnection happen, only inconsistent regions in the affected sub-volume are synchronized using the FastResync feature. In case of a storage failure, the mirror of the sub-volume on that storage will be detached and the future IOs on the sub-volume will be tracked by the DCO associated with the parent volume. When such a detached mirror is reattached after restoring storage connectivity, only regions that are inconsistent in the mirror would be synchronized using the FastResync feature.

Prior to this release, for a layered volume, if the storage within a mirror of a sub-volume became inaccessible, it led to full synchronization of that mirror when the storage was reconnected.

For more information about FastResync, see the *Symantec Storage Foundation Cluster File System High Availability Administrator's Guide*.

Read policy enhancement

In this release, to optimize the read performance, changes have been made in the plex read policies on VxVM volumes. When there are more than one mirror available to serve the read IO, VxVM will select the set of mirrors that will provide the optimal performance and round robin between those. In selecting the set of mirrors, the internal logic will take into account various factors such as site locality, disk connectivity, media type, layout(striping), etc. You can override the logic and set any plex as the preferred mirror or set a round-robin read policy to round robin between all the mirrors of a volume.

For more information about read policies, see the *Administrator's Guide*.

The `vxattachd` daemon added as a VCS resource

In this release, the automatic site reattachment daemon, `vxattachd`, has been added in the list of resources monitored by VCS. The ProcessOnOnly agent in VCS will now monitor the `vxattachd` daemon. If the `vxattachd` process is not running, then in the next monitor cycle this agent will detect and restart it.

For more information about the `vxattachd` daemon, see the *Administrator's Guide*.

VOM integration with FSS

The Flexible Storage Sharing (FSS) feature in VxVM has been integrated with the Veritas Operations Manager (VOM) version 6.1. All the FSS operations can be done through the VOM console.

See the *Veritas™ Operations Manager Management Server 6.1 User Guide* for details.

FSS Performance improvements

In this release, the remote I/O performance on the volumes of Flexible Storage Sharing (FSS) disk groups has been improved. As there are remote devices in the FSS environment, the I/Os destined on remote disks used to give less throughput. In this release, changes have been done in the Low Latency Transport (LLT) component and Cluster Volume Manager (CVM) to remove bottlenecks at different stages. With these changes, the I/Os on remote disks give improved performance, and can reach a maximum of 90% of the local disks' performance.

For more information about FSS, see the *Administrator's Guide*.

Disk Support for FSS

In this release, in addition to the disks supported in the HCL, those disks that have the capability of generating unique IDs, have been extended support in the FSS environment. The VxVM operations such as exporting a disk, adding a disk to an FSS disk group, and creating an FSS disk group have been enhanced to internally check the compatibility of disks involved in these operations. You can use the standalone utility `vxddladm checkfss` *`disksname`* to check whether the disk is supported for FSS.

For more information about FSS disk support, see the *Administrator's Guide*.

Changes in default layout of cachearea volume used by SmartIO

When cachearea is created on multiple devices, stripe layout is used by default instead of concat for creating the cachearea volume. A new option is added to sfcache CLI to override this behavior.

Synchronize existing volumes that may have been created without synchronization

The `vxvol` command `sync` attribute lets you synchronize existing volumes that may have been created without synchronization. You should run `vxvol sync` when the volume is idle.

For more information, see the `vxvol`(1M) man page.

Changes related to Veritas File System

There are no changes related to VxFS in this release.

Changes related to SFDB tools

The following sections describe the changes related to Storage Foundation for Databases (SFDB) tools in 6.2.

Support for multitenant databases

SFDB tools support operations on Oracle 12c multitenant databases. The SFDB tools do not support operations on individual Pluggable Databases (PDB).

Managing OEM using the Symantec Storage plug-in

Symantec Storage plug-in provides a graphical interface to efficiently manage and view your Storage Foundation and VCS objects through Oracle Enterprise Manager 12c (OEM).

The plug-in has the following three tabs:

- SmartIO - provides a gateway to manage the objects that use Storage Foundation's SmartIO feature, which is an advanced caching solution.
- Snapshot - enables you to apply the SFDB's point-in-time copy technologies to the selected database objects, such as datafiles, tablespaces.
- Cluster - extracts various configuration-specific information from the Symantec Cluster Server and manifests them in a tabular format.

For details on downloading and using the plug-in, visit

https://www-secure.symantec.com/connect/downloads/sfha-solutions-62-symantec-storage-plug-oem-12c

Changes to LLT, GAB, and I/O fencing

This section covers new features or enhancements made to LLT, GAB, and I/O fencing.

Changes to LLT

Symantec Cluster Server includes the following changes to LLT in 6.2:

Changes to GAB

Symantec Cluster Server (VCS) includes the following changes to GAB in 6.2:

Changes to I/O fencing

Symantec Cluster Server (VCS) includes the following changes to I/O fencing in 6.2:

I/O fencing supports majority-based fencing mechanism, a new fencing mechanism that does not need coordination points

I/O fencing supports a new fencing mode called majority-based I/O fencing. Majority-based I/O fencing mode does not need coordination points to provide protection against data corruption and data consistency in a clustered environment. Use majority-based I/O fencing when there are no additional servers and or shared SCSI-3 disks to be used as coordination points. It provides a reliable arbitration method and does not require any additional hardware setup, such as CP Servers or shared SCSI3 disks.

In the event of a network failure, the majority sub-cluster wins the fencing race and survives the race. Note that even if the majority sub-cluster is hung or unresponsive, the minority sub-cluster loses the fencing race and the cluster panics. The cluster remains unavailable till the issue is resolved.

For more information, refer to the *Installation Guide* and *Administrator's Guide*.

Clear coordination point server registrations using the vxfenclearpre utility

The vxfenclearpre utility is enhanced to clear registrations from coordination point servers for the current cluster in addition to the existing functionality to remove SCSI3 registrations and reservation keys from the set of coordinator disks and shared data disks. The local node from where you run the utility must have the UUID of the current cluster at `/etc/vx/.uuids` directory in the `clusuuid` file.

Note that you may experience delays while clearing registrations on the coordination point servers because the utility tries to establish a network connection with IP addresses used by the coordination point servers. The delay may occur because of a network issue or if the IP address is not reachable or is incorrect.

For more information, refer to the *Administrator's Guide*.

Raw disk I/O fencing policy is not supported

Symantec does not support raw disk policy for I/O fencing. Use DMP as the I/O fencing policy for coordinator disks that have either a single hardware path or multiple hardware paths to nodes.

For more information, refer to the *Installation Guide* and *Administrator's Guide*.

SCSI3 fencing support inside Virtual Machines

SCSI3 fencing is now supported inside KVM and VMware Virtual Machines.

For more information, refer to *Virtualization Guide*.

Release level terminology changes

With the 6.2 release, terms that are used to describe patch-based releases have changed as follows:

Table 1-3 Release level terminology changes

Pre 6.0.1	6.0.x, 6.1, 6.1.x	6.2 and forward	Status	Available from
P-Patch	Public hot fix	Patch	Official	SORT
Hot fix	Private hot fix	Hot fix	Unofficial	Customer support

Official patch releases are available from SORT. This release was previously referred to as a P-Patch or a Public hot fix and is now referred to as a Patch. Unofficial patch releases are available from customer support. Hot fix is the only unofficial patch release.

No longer supported

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

- Raw disk I/O fencing policy is no longer supported.

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems

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

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

Table 1-4 Supported operating systems

Operating systems	Supported level and kernel version
Red Hat Enterprise Linux 7	3.10.0-123.el7

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

Operating systems	Supported level and kernel version
Red Hat Enterprise Linux 6	Update 3 (2.6.32-279.el6) Update 4 (2.6.32-358.el6) Update 5 (2.6.32-431.el6) Update 6 (2.6.32-504.el6)
SUSE Linux Enterprise 11	SP2 (3.0.13-0.27.1) SP3 (3.0.76-0.11.1)
Oracle Linux 6 (RHEL compatible mode)	Update 3 (2.6.32-279.el6) Update 4 (2.6.32-358.el6) Update 5 (2.6.32-431.el6) Update 6 (2.6.32-504.el6)
Oracle Linux 7 (RHEL compatible mode)	3.10.0-123.el7 **Note:** SF Oracle RAC has not yet announced support for Oracle Linux 7. You may find information pertaining to OL 7 in the installation and administrator guides. Note that this information will become relevant only after SF Oracle RAC announces support when due certification efforts are complete. Refer to the following TechNote for the latest information on the supported operating systems and Oracle RAC database versions. http://www.symantec.com/docs/DOC4848

Note: All subsequent kernel updates are supported, but you should check the Symantec Operations Readiness Tools (SORT) website for additional information that applies to the exact kernel version for which you plan to deploy.

Note: Only 64-bit operating systems are supported on the AMD Opteron or the Intel Xeon EM64T (x86_64) Processor line.

Note: Configuring LLT over RDMA is not supported with Oracle Linux Unbreakable Enterprise Kernel 2 that is 2.6.39-400.17.1.el6uek.x86_64.

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

http://www.symantec.com/docs/TECH211575

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 Symantec software. Consult the Red Hat, SUSE, or Oracle documentation for more information on upgrading or reinstalling your operating system.

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

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

Required Linux RPMs for SFCFSHA

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

Note: Some required RHEL RPMs have different version numbers between RHEL update versions.

Table 1-5 lists the RPMs that SFCFSHA requires for a given Linux operating system.

Table 1-5 Required RPMs

Operating system	Required RPMs
RHEL 7 **Note:** Symantec recommends that you install RHEL 7 as the operating system of Server GUI.	bc -1.06.95-13.el7.x86_64 coreutils-8.22-11.el7.x86_64 ed-1.9-4.el7.x86_64 findutils-4.5.11-3.el7.x86_64 gcc-c++-4.8.2-16.el7.x86_64 gcc-4.8.2-16.el7.x86_64 glibc-2.17-55.el7.i686 glibc-2.17-55.el7.x86_64 glibc-headers-2.17-55.el7.x86_64 glib-networking-2.36.2-3.el7.x86_64 glibmm24-2.36.2-4.el7.x86_64 glibc-common-2.17-55.el7.x86_64 glibc-devel-2.17-55.el7.x86_64 glibc-devel-2.17-55.el7.i686 glib2-2.36.3-5.el7.x86_64 glibc-utils-2.17-55.el7.x86_64 kmod-14-9.el7.x86_64 ksh-20120801-19.el7.x86_64

Table 1-5 Required RPMs *(continued)*

Operating system	Required RPMs
RHEL 7 (continued)	libacl-2.2.51-12.el7.i686
	libacl-2.2.51-12.el7.x86_64
	libaio-devel-0.3.109-12.el7.x86_64
	libaio-devel-0.3.109-12.el7.i686
	libaio-0.3.109-12.el7.i686
	libaio-0.3.109-12.el7.x86_64
	libgcc-4.8.2-16.el7.i686
	libgcc-4.8.2-16.el7.x86_64
	libstdc++-4.8.2-16.el7.i686
	libstdc++-4.8.2-16.el7.x86_64
	lsof-4.87-4.el7.x86_64
	ncompress-4.2.4.4-3.el7.x86_64
	ncurses-libs-5.9-13.20130511.el7.x86_64
	nss-softokn-freebl-3.15.4-2.el7.i686
	pam-1.1.8-9.el7.i686
	parted-3.1-17.el7.x86_64
	pcre-8.32-12.el7.i686 (pcre(x86-32))
	policycoreutils-2.2.5-11.el7.x86_64
	prelink-0.5.0-6.el7.x86_64
	screen-4.1.0-0.19.20120314git3c2946.el7.x86_64
	systemd-libs-208-11.el7.i686
	systemd-libs-208-11.el7.x86_64
	xz-libs-5.1.2-8alpha.el7.i686 (xz-libs(x86-32))

Table 1-5 Required RPMs *(continued)*

Operating system	Required RPMs
OL 6	coreutils-8.4-19.el6.x86_64.rpm
	ed-1.1-3.3.el6.x86_64.rpm
	findutils-4.4.2-6.el6.x86_64.rpm
	glibc-2.12-1.80.el6.i686.rpm
	glibc-2.12-1.80.el6.x86_64.rpm
	ksh-20100621-16.el6.x86_64.rpm
	libacl-2.2.49-6.el6.x86_64.rpm
	libgcc-4.4.6-4.el6.i686.rpm
	libgcc-4.4.6-4.el6.x86_64.rpm
	libstdc++-4.4.6-4.el6.i686.rpm
	libstdc++-4.4.6-4.el6.x86_64.rpm
	module-init-tools-3.9-20.0.1.el6.x86_64.rpm
	ncurses-libs-5.7-3.20090208.el6.x86_64.rpm
	nss-softokn-freebl-3.12.9-11.el6.i686.rpm
	openssl-1.0.0-20.el6_2.5.x86_64.rpm
	pam-1.1.1-10.el6_2.1.i686.rpm
	parted-2.1-18.el6.x86_64.rpm
	perl-5.10.1-127.el6.x86_64.rpm
	policycoreutils-2.0.83-19.24.0.1.el6.x86_64.rpm
	readline-6.0-4.el6.x86_64.rpm

Table 1-5 Required RPMs *(continued)*

Operating system	Required RPMs
RHEL 6	coreutils-8.4-19.el6.x86_64.rpm
	ed-1.1-3.3.el6.x86_64.rpm
	findutils-4.4.2-6.el6.x86_64.rpm
	glibc-2.12-1.80.el6.i686.rpm
	glibc-2.12-1.80.el6.x86_64.rpm
	ksh-20100621-16.el6.x86_64.rpm
	libacl-2.2.49-6.el6.x86_64.rpm
	libgcc-4.4.6-4.el6.i686.rpm
	libgcc-4.4.6-4.el6.x86_64.rpm
	libstdc++-4.4.6-4.el6.i686.rpm
	libstdc++-4.4.6-4.el6.x86_64.rpm
	mksh-39-7.el6.x86_64.rpm
	module-init-tools-3.9-20.el6.x86_64.rpm
	ncurses-libs-5.7-3.20090208.el6.x86_64.rpm
	nss-softokn-freebl-3.12.9-11.el6.i686.rpm
	openssl-1.0.0-20.el6_2.5.x86_64.rpm
	pam-1.1.1-10.el6_2.1.i686.rpm
	parted-2.1-18.el6.x86_64.rpm
	policycoreutils-2.0.83-19.24.el6.x86_64.rpm
	readline-6.0-4.el6.x86_64.rpm
	zlib-1.2.3-27.el6.x86_64.rpm

Table 1-5 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 11 SP2	coreutils-8.12-6.19.1.x86_64.rpm
	ed-0.2-1001.30.1.x86_64.rpm
	findutils-4.4.0-38.26.1.x86_64.rpm
	glibc-2.11.3-17.31.1.x86_64.rpm
	glibc-32bit-2.11.3-17.31.1.x86_64.rpm
	ksh-93u-0.6.1.x86_64.rpm
	libacl-2.2.47-30.34.29.x86_64.rpm
	libacl-32bit-2.2.47-30.34.29.x86_64.rpm
	libgcc46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm
	libgcc46-4.6.1_20110701-0.13.9.x86_64.rpm
	libncurses5-5.6-90.55.x86_64.rpm
	libstdc++46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm
	libstdc++46-4.6.1_20110701-0.13.9.x86_64.rpm
	module-init-tools-3.11.1-1.21.1.x86_64.rpm
	pam-32bit-1.1.5-0.10.17.x86_64.rpm
	parted-2.3-10.21.18.x86_64.rpm
	zlib-32bit-1.2.3-106.34.x86_64.rpm

Table 1-5 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 11 SP3	coreutils-8.12-6.25.27.1.x86_64.rpm
	ed-0.2-1001.30.1.x86_64.rpm
	findutils-4.4.0-38.26.1.x86_64.rpm
	glibc-2.11.3-17.54.1.x86_64.rpm
	glibc-32bit-2.11.3-17.54.1.x86_64.rpm
	ksh-93u-0.18.1.x86_64.rpm
	libacl-2.2.47-30.34.29.x86_64.rpm
	libacl-32bit-2.2.47-30.34.29.x86_64.rpm
	libgcc_s1-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libgcc_s1-4.7.2_20130108-0.15.45.x86_64.rpm
	libncurses5-5.6-90.55.x86_64.rpm
	libstdc++6-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libstdc++6-4.7.2_20130108-0.15.45.x86_64.rpm
	module-init-tools-3.11.1-1.28.5.x86_64.rpm
	pam-32bit-1.1.5-0.10.17.x86_64.rpm
	parted-2.3-10.38.16.x86_64.rpm
	zlib-1.2.7-0.10.128.x86_64.rpm
	zlib-32bit-1.2.7-0.10.128.x86_64.rpm

Additional RPMs required for Veritas Volume Manager

You must install the 32-bit `libudev` RPM before you install Veritas Volume Manager.

Table 1-6 lists the required RPMs.

Table 1-6 Additional RPMs required for Veritas Volume Manager

Operating system	Required RPMs
RHEL 7	systemd-libs-208-11.el7.i686.rpm
RHEL 6 Update 5	libudev-147-2.51.el6.i686.rpm
RHEL 6 Update 4	libudev-147-2.46.el6.i686.rpm
RHEL 6 Update 3	libudev-147-2.41.el6.i686.rpm

Table 1-6 Additional RPMs required for Veritas Volume Manager *(continued)*

Operating system	Required RPMs
SLES 11 SP3	libudev0-32bit-147-0.84.1.x86_64.rpm
SLES 11 SP2	libudev0-32bit-147-0.47.2.x86_64.rpm

Symantec Storage Foundation Cluster File System High Availability hardware requirements

The following hardware requirements apply to Symantec Storage Foundation Cluster File System High Availability.

Table 1-7 Hardware requirements for Symantec Storage Foundation Cluster File System High Availability

Requirement	Description
Memory	2 GB of memory.
CPU	A minimum of 2 CPUs.
Node	All nodes in a Cluster File System must have the same operating system version.
Shared storage	Shared storage can be one or more shared disks or a disk array connected either directly to the nodes of the cluster or through a Fibre Channel Switch. Nodes can also have non-shared or local devices on a local I/O channel. It is advisable to have `/`, `/usr`, `/var` and other system partitions on local devices. In a Flexible Storage Sharing (FSS) environment, shared storage may not be required.
Fibre Channel or iSCSI storage	Each node in the cluster must have a Fibre Channel I/O channel or iSCSI storage to access shared storage devices. The primary component of the Fibre Channel fabric is the Fibre Channel switch.
Cluster platforms	There are several hardware platforms that can function as nodes in a Symantec Storage Foundation Cluster File System High Availability (SFCFSHA) cluster. For a cluster to work correctly, all nodes must have the same time. If you are not running the Network Time Protocol (NTP) daemon, make sure the time on all the systems comprising your cluster is synchronized.

Table 1-7 Hardware requirements for Symantec Storage Foundation Cluster File System High Availability *(continued)*

Requirement	Description
SAS or FCoE	Each node in the cluster must have an SAS or FCoE I/O channel to access shared storage devices. The primary components of the SAS or Fibre Channel over Ethernet (FCoE) fabric are the switches and HBAs.

Storage Foundation for Databases features supported in database environments

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

Table 1-8 SFDB features supported in database environments

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

Notes:

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

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

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

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

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

Disk space requirements

Before installing any of the Symantec Storage Foundation products, confirm that your system has enough free disk space.

Use the "Perform a Preinstallation Check" (P) menu or the `-precheck` option of the product installer to determine whether there is sufficient space.

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

Number of nodes supported

SFCFSHA supports cluster configurations with up to 64 nodes.

Flexible Storage Sharing (FSS) only supports cluster configurations with up to 8 nodes.

SmartIO writeback caching only supports cluster configurations with up to 2 nodes.

Fixed issues

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

Installation and upgrades fixed issues

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

Table 1-9 Fixed issues related to installation and upgrades

Incident	Description
3326196	Rolling upgrade may encounter a problem if open volumes from different disk groups have the same name.

Table 1-9 Fixed issues related to installation and upgrades *(continued)*

Incident	Description
3326639	CP server service group fails to come online with the default database path after the CP server is upgraded from 6.0 to 6.2 on a multi-node cluster.
3442070	If you select rolling upgrade task from the Install Bundles menu, the Installer exits with an error.

Symantec Storage Foundation Cluster File System High Availability fixed issues

This section describes the incidents that are fixed in Symantec Storage Foundation Cluster File System High Availability in this release.

See "Veritas File System fixed issues" on page 36.

See "Veritas Volume Manager fixed issues" on page 39.

Table 1-10 Symantec Storage Foundation Cluster File System High Availability fixed issues

Incident	Description
3642894	VxFS agents for VCS should honor the customization of `VCS_LOG`.
3592783	On the FSS partially shared storage configuration, after you run `hastop -all` and then `hastart`, remote disks, which are not part of the disk group, are not visible in `vxdisk -o alldgs list`.
3582470	Data change object (DCO) volume gets created using the same node's disks for both plexes.
3573908	Multiple `cbrbk.tmp$$` files in the /var/tmp folder on each node do not clean up properly.
3552008	The `vxconfigrestore` (vxvol resync) operation hangs on the master node while recovering a stripe-mirror volume.
3551050	The `vx*` commands are not able to connect to vxconfigd.
3538683	After a master panic, the new master does not start plex resync when the older master comes online and joins the cluster.

Table 1-10 Symantec Storage Foundation Cluster File System High Availability fixed issues *(continued)*

Incident	Description
3537519	The `vxdisk unexport` command hangs and then `vxconfigd` gets fault in an asymmetric array connection setup.
3534779	Internal stress testing on Cluster File System (CFS) hits a debug assert.
3523731	VxVM command check for device compliance prior to doing FSS operations on disks.
3511444	Panics occur while checking memory pressure.
3508390	DCO object is unnecessarily marked BADLOG mode in cascade failure scenarios, and it results in requiring a full-recovery and can result in lost snapshots as well.
3505017	When the `da name` is the same as the existing `dm name`, the `vxdg addisk` operation from slave fails.
3496673	On a Flexible Storage Sharing (FSS) disk group, the read I/O performance is impacted.
3495811	When you create a disk group, the disk shows LMISSING state when the SCSI PGR operation fails.
3489167	Plex(es) on remote disks goes to DISABLED state because of a plex I/O error encountered after slave node reboot in cluster volume manger.
3484570	Accessing CVM message after decrementing reference count causes a panic.
3482026	The `vxattachd`(1M) daemon reattaches plexes of manually detached site.
3449152	The `vxtunefs`(1M) command fails to set the `thin_friendly_alloc` tunable in cluster file systems.
3444771	Internal noise test on cluster file system hits an debug assert when you are creating a file.
3430256	Space allocation for Volume: Single DAS disk in a disk group takes more preference than shared storage in that disk group.

Table 1-10 Symantec Storage Foundation Cluster File System High Availability fixed issues *(continued)*

Incident	Description
3422704	Unique prefix generation algorithm redesign for forming cluster wide consistent name (CCN).
3411438	The `preferred` and `round-robin` read policy settings should be honoured irrespective of local connectivity to the plexes.
3394940	Anomaly numbers are displayed in the `vxstat` output.
3383625	When a cluster node that contributes the storage to the Flexible Storage Sharing (FSS) disk group rejoins the cluster, the local disks brought back by that node do not get reattached.
3373747	Adding new nodes to the 22-node cluster causes Cluster File System (CFS) failures after CVM deletes 2 nodes.
3370753	Internal tests with SmartIO writeback SSD cache hit debug asserts.
3370722	Operating system (OS) installation on virtual machines fails in two-node clusters with writeback caching enabled.
3368361	When siteconsistency is configured within a private disk group and Cluster Volumne Manager (CVM) is up, then the reattach operation of a detached site fails.
3329603	The `vxconfigd` related error messages are observed in system log files on every node for large cluster setups.
3301085	The `vxdg adddisk` operation fails when adding nodes containing disks with the same name.
3300418	VxVM volume operations on shared volumes cause unnecessary read I/Os.
3286030	Vxattachd debug messages get displayed on the console during a reboot.
3283518	Disk group deport operation reports messages in the syslog for remote disks.
3283418	Writes from the source node hang due to heavy workload on the target node.

Table 1-10 Symantec Storage Foundation Cluster File System High Availability fixed issues *(continued)*

Incident	Description
3283418	Remote writes hang due to heavy sync workload on the target node in FSS environments.
3281160	When autoreminor is set off, no error is thrown when you import a disk group having the same minor number as that of the existing imported disk group.
3214542	Disk group creation or addition of a new disk to an existing disk group fails with "VxVM vxdg ERROR V-5-1-16087 Disk for disk group not found" error when the command is executed from the slave node.
3213411	A node join fails if a "disabled" Flexible Storage Sharing disk group exists in the cluster.
3198590	SmartIO VxVM caching is not supported for CFS.
3191807	CVM requires the T10 vendor provided ID to be unique.
3152304	When connectivity to some of the plexes of a volume is lost from all nodes, an I/O hang occurs.
3117153	Disk group remains in deported state for remote disks after destroying the disk group from a node that is not exporting the disks.
3079819	`vxconfigbackup` fails on Flexible Storage Sharing disk groups.
2705055	The `preferred` and `round-robin` read policy settings should be honoured irrespective of local connectivity to the plexes.
1203819	In some cases, inode and allocation maps inconsistencies occur in the event of a node crash in clusters.
640213	The node panics in case of overlapping reconfigurations due to race conditions.

Veritas File System fixed issues

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

Table 1-11 Veritas File System fixed issues

Incident	Description
3641719	The `fallocate` may allocate a highly fragmented file when the size of the file is extremely large.
3613048	VxFS does not correctly support IOCB_CMD_PREADV and IOCB_CMD_PREADV, which causes an error in the kernel code. Now the support for the vectored Asynchronous I/O commands is added to fix this issue.
3597482	The `pwrite(2)` function fails with the `EOPNOTSUPP` error.
3589264	The `fsadm` command shows an incorrect option for the file system type in usage.
3564076	The MongoDB noSQL database creation fails with an `ENOTSUP` error.
3563796	The file system `fullfsck` flag is set when the inode table overflows.
3560968	The `delicache_enable` tunable is not persistent in the Cluster File System (CFS) environment.
3560187	The kernel may panic when the buffer is freed in the `vx_dexh_preadd_space()` function with the message `Data Key Miss Fault in kernel mode`.
3557009	Run the `fallocate` command with `-l` option to specify the length of the reserve allocation. The file size is not expected, but multiple of file system block size.
3550103	After you upgrade or restart the system, mismatch in SSD cache usage may occur.
3523316	The writeback cache feature does not work for write size of 2MB.
3520349	When there is a huge number of dirty pages in the memory, and a sparse write is performed at a large offset of 4 TB or above on an existing file that is not null, the file system hangs
3506485	The code is modified to not allow writeback caching on a volume with Veritas Volume Replicator (VVR) enabled.
3486726	Veritas File Replicator (VFR) logs too much data on the target node.

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

Incident	Description
3484336	The `fidtovp()` system call can panic in the `vx_itryhold_locked()` function.
3473390	The multiple stack overflows with Veritas File System (VxFS) on Red Hat Enterprise Linux (RHEL) 6 lead to panics or system crashes.
3471245	The `mongodb` fails to insert any record.
3469644	The system panics in the `vx_logbuf_clean()` function.
3466020	The file system is corrupted with the error message `vx_direrr: vx_dexh_keycheck_1`.
3457803	The file system gets disabled intermittently with metadata I/O errors.
3444775	Internal noise testing on cluster file system results in a kernel panic in `vx_fsadm_query()` function with an error message.
3444154	Reading from a de-duped file-system over NFS can result in data corruption seen on the NFS client.
3436326	The attribute validation (pass 1d) of the `full fsck` operation takes too much time to complete.
3434811	The `vxfsconvert`(1M) command in VxFS 6.1 hangs.
3430461	The nested unmounts fail if the parent file system is disabled.
3424564	`fsppadm` fails with `ENODEV` and `file is encrypted or is not a database` errors.
3417076	The `vxtunefs`(1M) command fails to set tunables when the file contains blank lines or white spaces.
3415639	The type of the `fsdedupadm`(1M) command always shows as MANUAL even when it is launched by the `fsdedupschd` daemon.
3412667	The RHEL 6 system panics with a stack overflow.
3394803	A panic is observed in the VxFS routine `vx_upgrade7`() function while running the `vxupgrade` command (1M).
3370727	Internal tests with SmartIO writeback SSD cache hit debug asserts.

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

Incident	Description
3356947	When there are multi-threaded writes with `fsync` calls between them, VxFS becomes slow.
3352883	During the rename operation, many nfsd threads hang.
3340286	After a file system is resized, the tunable setting `dalloc_enable` is reset to a default value.
3337806	The `find`(1) command may panic the systems with Linux kernels with versions greater than 3.0.
3335272	The `mkfs` (make file system) command dumps core when the log size provided is not aligned.
3332902	While shutting down, the system running the `fsclustadm`(1M) command panics.
3317368	File system operations needing a file system freeze may take longer in the presence of file level snapshots and when there is a heavy I/O load.
3297840	VxFS corruption is detected during a dynamic LUN resize operation.
3294074	The `fsetxattr()` system call is slower on the Veritas File System (VxFS) than on the ext3 file system.
3285927	The `vfradmin` command does not validate job and consistency group names. It displays an improper error `Stale NFS file handle` for invalid job/consistency group.
3121933	There is a DB2 crash or corruption when `EOPNOTSUPP` is returned from VxFS.

Veritas Volume Manager fixed issues

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

Table 1-12 Veritas Volume Manager fixed issues

Incident	Description
3584341	The `vxdmpadm listapm` command prints error message VxVM ERROR V-5-2-14196.

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

Incident	Description
3584311	The vxconfigd daemon hangs with "vol_rv_transaction_prepare+0005C8" on secondary site.
3580962	A panic occurs in VxDMP under high I/O load, and may cause complete storage disconnection.
3577957	Database instance is terminated while rebooting a node.
3577586	The server panics in dmp_send_scsipkt_req().
3567823	System reboot causes all PP_EMC devices to be lost if DMP devices are excluded.
3566493	Orphan cpmap objects cannot be removed after you disassociate unfinished snap plexes.
3565212	I/O failure occurs during controller giveback operations with Netapp FAS31700 array.
3564260	The `vxrlink pause` command hangs on the primary master node.
3555230	The `vxconfigd` daemon hangs in Veritas Volume Replicator (VVR) when writing to SRL volume during replication.
3554608	Mirroring a volume creates a larger plex than the original on a CDS disk.
3544972	620:dmp:coredump while rebooting the OS after dmp installation.
3543284	Storage devices are not visible in the vxdisk list or the vxdmpadm getdmpnode outputs.
3542713	The `vxdmpadm listenclosure` all displays a different enclosure from array console.
3542272	The `vxconfigbackupd` daemon never exits after reboot. The daemon remains active for a disk group because configuration has been changed after the backup is initiated.
3539548	Duplicate disks and I/O error occurs after dynamic LUN allocation.
3531385	Asynchronous access to per dmpnode request queues may cause system panic.
3526500	DMP I/O getting timeout lot earlier than io timeout value if I/O statistics daemon is not running.

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

Incident	Description
3521726	When using Symantec Replication Option, system panics happens due to double freeing IOHINT memory.
3520991	The `vxconfigd`(1M) daemon dumps core due to memory corruption.
3513392	Secondary panics when rebooted while heavy IOs are going on primary.
3503852	With multiple Replicated Volume Groups (RVGs), if you detach storage on a secondary master then reattach it back, rlinks are not able to connect. The rlink state is different on one of the three rlinks. .
3502923	ESX panic while running add/remove devices from smartpool with no license installed on server.
3498228	The `vxconfigd` core dump occurs after port disable or enable operation with migration from PP to DMP.
3495553	DV:6.1 The `vxconfigd` daemon hangs on secondary in vol_ru_transaction_prepare.
3490458	After managing class under PP, some of the devices are seen in error state.
3489572	Slave nodes panic when volume with DCO hits storage failure while volume is online.
3482026	The `vxattachd`(1M) daemon reattaches plexes of manually detached site.
3478019	When VxVM fails to assign a unique name to a new DCL volume, the `vxsnap prepare` command fails silently without giving an error.
3466160	DMP does not coexist with `scsi_dh_emc` module on SLES11.
3456153	When Veritas Volume Replicator (VVR) replication is in progress, a Cluster Volume Manager (CVM) slave node reboot causes an I/O hang.
3455460	The `vxfmrshowmap` and the `verify_dco_header` utilities fail.
3450758	The slave node was not able to join CVM cluster and resulted in panic.
3446415	A pool may get added to the file system when the file system shrink operation is performed on FileStore.
3440790	The `vxassist` command with parameter mirror and the vxplex command(1M) with parameter att hang.
3433503	Due to an incorrect memory access, the `vxconfigd` daemon cores with a stack trace.

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

Incident	Description
3428025	When heavy parallel I/O load is issued, the system that runs Symantec Replication Option (VVR) and is configured as VVR primary crashes.
3417044	System becomes unresponsive while creating a VVR TCP connection.
3415188	I/O hangs during replication in Veritas Volume Replicator (VVR).
3411668	Network and host endian difference is not handled in the nmcom_print_sock_storage() function.
3403390	After a crash, the linked-to volume goes into NEEDSYNC state.
3399323	The reconfiguration of DMP database fails.
3399131	For PowerPath (PP) enclosure, both DA_TPD and DA_COEXIST_TPD flags are set.
3390162	nmap scanning UDP port 4145 will cause vxnetd to consume 100% CPU and rlink disconnection resulting in system hangs.
3385905	Data corruption occurs after VxVM makes cache area offline and online again without a reboot.
3385753	Replication to the Disaster Recovery (DR) site hangs even though Replication links (Rlinks) are in the connected state.
3380481	When you select a removed disk during the "5 Replace a failed or removed disk" operation, the vxdiskadm(1M) command displays an error message.
3374200	A system panic or exceptional IO delays are observed while executing snapshot operations, such as, refresh.
3373208	DMP wrongly sends the SCSI PR OUT command with APTPL bit value as ‘0’ to arrays.
3368361	When siteconsistency is configured within a private disk group (with LUNs mapped only to local server) and CVM is up, then the reattach operation of a detached site fails.
3336714	The slab of I/O request in Linux may get corrupted.
3326964	VxVM hangs in Clustered Volume Manager (CVM) environments in the presence of FMR operations.
3317430	The `vxdiskunsetup` utility throws error after upgrade from 5.1SP1RP4.

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

Incident	Description
3287940	LUNs from any EMC CLARiiON arrays that have Not Ready state are shown in the "online invalid" state by Veritas Volume Manager (VxVM).
3279932	The `vxdisksetup` and `vxdiskunsetup` utilities fail for disks that are part of a deported disk group, even if "-f" option is specified.
3236772	Heavy I/O loads on primary sites result in transaction/session timeouts between the primary and secondary sites.
3133854	SmartIO cache area does not come online after I/O error to the disk.
3124698	VxDMP memory allocation mechanism affects system performance, due to excessive swapping.
2882312	If an SRL fault occurs in the middle of an I/O load, and you immediately issue a read operation on data written during the SRL fault, the system returns old data.
2573229	On RHEL6, the server panics when DMP executes PERSISTENT RESERVE IN command with REPORT CAPABILITIES service action on powerpath controlled device.
2049371	DMP behavior on Linux SLES11 when connectivity to a path is lost.
1390029	The `vxconfigrestore` command fails when there is a dot in the disk group name, i.e., test.2

LLT, GAB, and I/O fencing fixed issues

Table 1-13 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-13 LLT, GAB, and I/O fencing fixed issues

Incident	Description
3335137	Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups.
3473104	When virtual NICs are configured under LLT without specifying the MTU size 1500 in lltab, cluster does not function properly. For example, VCS engine commands may hang and print below message in the engine logs: `VCS CRITICAL V-16-1-51135 GlobalCounter not updated`

Table 1-13 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
3031216	The dash (-) in a disk group name causes vxfentsthdw(1M) and Vxfenswap(1M) utilities to fail.
3471571	Cluster nodes may panic if you stop the HAD process by force on a node and reboot that node.
3410309	LLT driver fails to load and logs the following message in the syslog when a mismatch is observed in the RDMA-specific symbols. llt: disagrees about version of symbol rdma_connect llt: Unknown symbol rdma_connect llt: disagrees about version of symbol rdma_destroy_id llt: Unknown symbol rdma_destroy_id
3532859	The Coordpoint agent monitor fails if the cluster has a large number of coordination points.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues

Table 1-14 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-14 SFDB tools fixed issues

Incident	Description
2869266	Checkpoint clone fails if the archive log destination is same as the datafiles destination.
3313775	SmartIO options are not restored after Reverse Resync Commit operation is performed.
3615735	During a Reverse Resync Begin operation, a mismatch in database control file version is observed.
3615745	For thin storage setups, the snapshot operation reports that the diskgroup cannot be split.
3615764	The flashSnap operation fails to create a symlink on a Symantec Volume Replicator (VVR) secondary site.

Known issues

This section covers the known issues in this release.

Installation known issues

This section describes the known issues during installation and upgrade.

`installer -requirements` does not list RHEL 6 Update 6 and Oracle Linux 7 as supported platforms (3657260)

The `installer -requirements` command does not list RHEL 6 Update 6 and Oracle Linux 7 as supported platforms though they are qualified with version 6.2.

Workaround: The correct supported list is mentioned in the latest version of the product Release Notes. See the latest Release Notes on the Symantec website for the updated list.

https://sort.symantec.com/documents

Installer reports incorrect minimal version for several required Oracle Linux 7 RPMs (3653382)

During installation, the product installer reports incorrect minimum version for the following required Oracle Linux 7 RPMs:

```
systemd-libs-208-11.el7.i686
coreutils-8.22-11.el7.x86_64
policycoreutils-2.2.5-11.el7.x86_64
```

The correct minimum version required for the RPMs is as follows:

```
systemd-libs-208-11.0.1.el7.i686.rpm
coreutils-8.22-11.0.1.el7.x86_64
policycoreutils-2.2.5-11.0.1.el7.x86_64
```

Workaround: Install the required operating system RPMs using native methods, such as yum, or install them manually.

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups (2574731)

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

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

To unfreeze the service groups manually

1 List all the frozen service groups

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

2 Unfreeze all the frozen service groups:

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

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

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to Symantec Storage Foundation (SF) 6.2, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/openv`), then while upgrading to SF 6.2, 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.

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
```

XFS file system is not supported for RDE

The Root Disk Encapsulation (RDE) feature is not supported if the root partition is mounted with XFS file system.

Workaround: There is no workaround available.

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, 2116264)

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

After performing the first phase of a rolling upgrade, make sure the CVM is online on all nodes without errors (2595441)

Make sure that the CVM is online on all nodes without errors after you perform the first phase of a rolling upgrade. The CVM protocol version will not upgrade successfully on the nodes during rolling upgrade phase two where CVM is offline or has errors.

If the CVM protocol version does not upgrade successfully, upgrade the CVM protocol on the CVM master node.

To upgrade the CVM protocol on the CVM master node

1. Find out which node is the CVM master:

   ```
   # vxdctl -c mode
   ```

2. On the CVM master node, upgrade the CVM protocol:

   ```
   # vxdctl upgrade
   ```

Upgrading from Storage Foundation 5.1 Service Pack 1 Rolling Patch 2 to 6.2 with rootability enabled fails (2581313)

While upgrading from Storage Foundation (SF) 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 to 6.2 with intermediate patches (if any), if you use an encapsulated root disk, the upgrade fails because the `initrd` image creation fails during the VxVM upgrade.

Workaround: To upgrade from 5.1 SP1 RP2 to6.2 with intermediate patches (if any) using an encapsulated root disk, you must reinstall the nash utility on the system prior to the upgrade from 5.1 SP1 RP2.

To upgrade from 5.1 SP1 RP2 to 6.2 with intermediate patches (if any) using an encapsulated root disk:

1. Reinstall the nash utility.
2. Upgrade to the SF 6.2 release.

During upgrade from 5.1SP1 to 6.2 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 SFCFSHA 5.1 SP1 to SFCFSHA 6.2 with an encapsulated root disk, splitting the root mirror fails if the target disk group name for the split operation is used by an existing deported disk group.

Workaround:

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

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

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

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

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

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

Workaround: Do one of the following:

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

After upgrade from VxVM version 5.1 SP1RP3 or version 6.0 with an encapsulated boot disk, the system fails to boot (2750782)

On Red Hat Enterprise Linux 6 (RHEL6), during the Veritas Volume Manager (VxVM) upgrade from version 5.1 SP1RP3 or version 6.0 to a higher version, the RPM runs the installation scripts of the VxVM higher version first. Then the RPM runs the uninstallation scripts of the existing VxVM version. Due to a defect in the 5.1 SP1RP3 or 6.0 uninstallation script, it corrupts the file installed by the higher version. This leads to boot failure.

Workaround:

1. Unroot the encapsulated root disk.
2. Uninstall the existing `VRTSvxvm` (5.1 SP1RP3 or 6.0) RPM.
3. Install `VRTSvxvm` of the higher version.

Erroneous resstatechange trigger warning [2277819]

You may encounter the following warning when you restart resources:

```
CPI WARNING V-9-40-4317 The installer has detected that resstatechange
trigger is configured by setting TriggerResStateChange attributes.
```

Workaround: In future releases, the resstatechange trigger will not be invoked when a resource is restarted. Instead, the resrestart trigger will be invoked if you set the TriggerResRestart attribute. The resrestart trigger is available in the current release. Refer to the VCS documentation for details.

The uninstaller does not remove all scripts (2696033)

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

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

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

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

Rolling upgrade of VCS from pre-6.0 versions fails with CP server in secure mode [3262900]

If the CP server is configured in secure mode, rolling upgrade of VCS from versions lower than 6.0 to 6.1 is not supported. Since the `vxcpserv` process is not compatible with shared authentication, CP server service group fails to come online after performing phase 1 of the rolling upgrade.

Workaround: Use full upgrade or phased upgrade instead of rolling upgrade.

Incorrect VVR tunable settings after upgrade to version 6.2 from 6.0 [3581543]

The `vol_min_lowmem_sz` and `vol_max_nmpool_sz` tunables may be set to a value less than their default values after you upgrade to version 6.2. Additionally, the `vxtune` command may allow the tunable value to be thus modified without displaying an error.

Workaround:

The problem has no critical functionality impact. However, for performance considerations, it is recommended that you verify that the value of the `vol_min_lowmem_sz` and `vol_max_nmpool_sz` tunables are set to at least the default value. Use the `vxtune` command to modify the tunable value.

Symantec Storage Foundation Cluster File System High Availability known issues

This section describes the known issues in this release of Symantec Storage Foundation Cluster File System High Availability (SFCFSHA).

See "Veritas File System known issues" on page 78.

See "Veritas Volume Manager known issues" on page 62.

NFS lock failover is not supported on Linux [3331646]

If a file is locked from an NFS client, the other client may also get the lock on the same file after failover of the NFS share service group. This is because of the changes in the format of the lock files and the lock failover mechanism.

Workaround: No workaround.

On CFS, SmartIO is caching writes although the cache appears as nocache on one node (3760253)

On CFS, SmartIO is caching writes although the `sfcache list` output shows the cache in `nocache` mode on one node. The OS `mount` command also shows the file systems as unmounted. This issue is due to a known bug that is documented in the Linux `mount` manual page. The `/etc/mtab` file and the `/proc/mounts` file, which are expected to have entries for all the mounted file systems, do not match. When the `sfcache list` command displays the list of file systems that are mounted in writeback mode, `sfcache list` refers to the `/etc/mtab` entries for the mount status of the file systems. As a result, `sfcache list` may sometimes show a writeback enabled file system as umounted while in reality the file system is still mounted. The `/proc/mounts` file correctly shows the file systems as mounted.

Workaround:

Verify that the file system is mounted through the contents of the `/proc/mounts` file.

Cache area is lost after a disk failure (3158482)

SmartIO supports one VxFS cache area and one VxVM cache area. If you create one cache area, and the disk fails, the cache area becomes disabled. If you attempt to create a second cache area of the other type before the cache disk group is enabled, then the first cache area is lost. It cannot be brought online.

For example, first you created a VxFS cache area. The disk failed and the cache area is disabled. Now create the VxVM cache area. While creating VxVM cache area, SmartIO looks for an existing default cache area. Due to the failed disk, the existing cache area cannot be found. So SmartIO creates a VxVM cache area with the same name. Now even if disk containing VxFS cache area comes up, SmartIO cannot access the original cache area. In this scenario, the VxFS cache area is lost. Losing the cache area in this case does not result into any data loss or data inconsistency issues.

Workaround:

Create a new VxFS cache area.

Incorrect usage message displays for sfcache app oracle command (3617893)

In some cases, the usage message that displays for the `sfcache app oracle` command may be incorrect.

`Workaround`:

Refer to the `sfcache`(1m) manual page for correct command usage.

CFS commands might hang when run by non-root (3038283)

The CFS commands might hang when run by non-root.

Workaround

To resolve this issue

- Use `halogin` command to save the authentication information before running any CFS commands on a non-root session.

  When you run the `halogin` command, VCS stores encrypted authentication information in the user's home directory.

NFS resource might not come online while configuring CNFS share (2488685)

If SELinux is configured as `enforcing` or `permissive`, NFS resource might not come online and go into `FAULTED` state while configuring CNFS share `cfsnfssg` service group.

Sample output:

```
# hastatus -sum

-- SYSTEM STATE
-- System               State                Frozen

A  swlx14               RUNNING              0

-- GROUP STATE
-- Group           System   Probed     AutoDisabled    State

B  cfsnfssg        swlx14   Y          N               OFFLINE|FAULTED
B  cfsnfssg_dummy  swlx14   Y          N               OFFLINE
B  cvm             swlx14   Y          N               ONLINE
B  vip1            swlx14   Y          N               OFFLINE
```

```
-- RESOURCES FAILED
-- Group           Type                 Resource             System

D  cfsnfssg        NFS                  nfs                  swlx14
```

Workaround

To resolve this issue you need to add the Ethernet port into the trusted list for SELinux.

- In the System Setup->Firewall configuration, select customize.
- In the Trusted device, select the Ethernet port.

Miscalculated file set usage (2123429)

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

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

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

Workaround:

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

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

Application group attempts to come online on primary site before fire drill service group goes offline on the secondary site (2107386)

The application service group comes online on the primary site while the fire drill service group attempts to go offline at the same time, causing the application group to fault.

Workaround: Ensure that the fire drill service group is completely offline on the secondary site before the application service group comes online on the primary site.

Panic due to null pointer de-reference in vx_bmap_lookup() (3038285)

If you use the `fsadm -b` command on a CFS secondary node to resize the file system, it might fail with the following error message printed in the syslog:

```
Reorg of inode with shared extent larger than 32768 blocks
can be done only on the CFS Primary node
```

Workaround: Resize the file system with the `fsadm` command from the primary node of the cluster.

tail -f run on a cluster file system file only works correctly on the local node (2613030)

When using the `tail -f` command to monitor a file on a cluster file system, changes to the file made on remote nodes are not detected. This is due to the `tail` command now utilizing `inotify`. Symantec is currently unable to support `inotify` with a cluster file system due to GPL restrictions.

Workaround: To revert to the old behavior, you can specify the `---disable-inotify` option with the `tail` command.

"Configuration must be ReadWrite : Use haconf -makerw" error message appears in VCS engine log when hastop -local is invoked (2609137)

A message similar to the following example appears in the `/var/VRTSvcs/log/engine_A.log` log file when you run the `hastop -local` command on any system in a SFCFSHA cluster that has `CFSMount`resources:

```
2011/11/15 19:09:57 VCS ERROR V-16-1-11335 Configuration must be
ReadWrite : Use haconf -makerw
```

The `hastop -local` command successfully runs and you can ignore the error message.

Workaround: There is no workaround for this issue.

Issues observed with force unmounting a parent cluster file system mount before unmounting a nested child VxFS or cluster file system mount (2621803)

When you have nested mounts in which a secondary VxFS file system is mounted in the name space of the primary file system in the cluster, if the primary file system gets force umounted before unmounting the secondary, then unmounting the secondary at a later time can cause unpredictable issues.

Workaround: There is no workaround for this issue.

Performance degradation seen on a CFS filesystem while reading from a large directory (2644485)

Performance degradation is seen on a CFS filesystem while reading from a large directory.

Workaround: There is no workaround.

vxdisk resize from slave nodes fails with "Command is not supported for command shipping" error (3140314)

When running the `vxdisk resize` command from a slave node for a local disk, the command may fail with the following error message:

```
VxVM vxdisk ERROR V-5-1-15861 Command is not supported for command
shipping.
Operation must be executed on master
```

Workaround: Switch the master to the node to which the disk is locally connected and run the `vxdisk resize` on that node.

Default volume layout with DAS disks spans across different disks for data plexes and DCO plexes (3087867)

The default volume layout for Flexible Storage Sharing disk groups is a two-way mirrored volume with a DCO log. The DCO log plexes may be allocated using host class instances that may be different from the ones used for the data plexes.

Workaround: Use the command line `alloc` attributes to explicitly set allocation requirements. For example, you can specify `alloc=host:host1,host:host2` during volume creation on an FSS disk group, and allocate DCO plexes on the same host (failure domain) as the data plexes.

SG_IO ioctl hang causes disk group creation, CVM node joins, and storage connects/disconnects, and vxconfigd to hang in the kernel (3193119)

In RHEL 5.x, the `SG_IO ioctl` process hangs in the kernel. This causes disk group creation and CVM node joins to hang. The `vxconfigd` thread hangs in the kernel during storage connects/disconnects and is unresponsive.

Workaround: This issue is fixed in RHEL 6.3. Upgrade to RHEL 6.3.

Preserving Flexible Storage Sharing attributes with vxassist grow and vxresize commands is not supported (3225318)

Preservation of FSS attributes using `vxasssist grow` and `vxresize` is not supported. FSS attributes include the attributes that are specified on the command line as well as the attributes implicitly assumed for FSS disk groups. These attributes are not reused with further `vxassist` orpations on the volume such as the `growby` and the `vxresize` commands.

Workaround:

There is no workaround for this issue.

Flexible Storage Sharing export operation fails when nodes in the cluster are joined in parallel (3327028)

When two or more nodes join the cluster in parallel in an FSS environment, the remote disk creation on some nodes may fail with the following message in the syslog:

```
vxvm:vxconfigd: V-5-1-12143 CVM_VOLD_JOINOVER command received for node(s) 1
vxvm:vxconfigd: V-5-1-3866 node 1: vxconfigd not ready
vxvm:vxconfigd: V-5-1-3866 node 1: vxconfigd not ready
vxvm:vxconfigd: V-5-1-18321 Export operation failed : Slave not joined
...
vxvm:vxconfigd: V-5-1-4123 cluster established successfully
```

The automatic reattach of subdisks and plexes may not occur, causing some resources to remain in the offline or faulted state. User intervention is required to remove the fault and bring the resources online.

Workaround:

Manually reattach the disks from the node that has connectivity to the disks:

```
# vxreattach diskname
```

If the resources are faulted, clear the fault and online the service group:

```
# hagrp -clear service_group

# hagrp -online service_group -any
```

In a Flexible Storage Sharing disk group, the default volume layout is not mirror when creating a volume with the mediatype:ssd attribute (3209064)

As per current VxVM behavior, specifying a subset of disks, such as mediatype:ssd, as a command line argument during volume creation, takes precedence over internal FSS attributes. VxVM does not implicitly apply by default the mirrored volume layout for a FSS volume.

Workaround: Explicitly specify the `layout=mirror` attribute during volume creation.

```
# vxassist -g diskgroup make volume size mediatype:ssd layout=mirror
```

FSS Disk group creation with 510 exported disks from master fails with Transaction locks timed out error (3311250)

Flexible Storage Sharing (FSS) Disk group creation for local disks that are exported may fail if the number of disks used for disk group creation is greater than 150, with the following error message:

```
VxVM vxdg ERROR V-5-1-585 Disk group test_dg: cannot create: Transaction
    locks timed out
```

A similar error can be seen while adding more that 150 locally exported disks (with `vxdg adddisk`) to the FSS disk group, with the following error message:

```
VxVM vxdg ERROR V-5-1-10127 associating disk-media emc0_0839 with emc0_0839:
        Transaction locks timed out
```

Workaround:

Create an FSS disk group using 150 or less locally exported disks and then do an incremental disk addition to the disk group with 150 or less locally exported disks at a time.

vxconfigrestore is unable to restore FSS cache objects in the pre-commit stage (3461928)

While restoring a Flexible Storage Sharing (FSS) disk group configuration that has cache objects configured, the following error messages may display during the pre-commit phase of the restoration:

```
VxVM vxcache ERROR V-5-1-10128  Cache object meta-data update error
VxVM vxcache ERROR V-5-1-10128  Cache object meta-data update error
VxVM vxvol WARNING V-5-1-10364 Could not start cache object
VxVM vxvol ERROR V-5-1-11802 Volume volume_name cannot be started
VxVM vxvol ERROR V-5-1-13386 Cache object on which Volume volume_name
 is constructed is not enabled
VxVM vxvol ERROR V-5-1-13386 Cache object on which Volume volume_name
 is constructed is not enabled
```

The error messages are harmless and do not have any impact on restoration. After committing the disk group configuration, the cache object and the volume that is constructed on the cache object are enabled.

Change in naming scheme is not reflected on nodes in an FSS environment (3589272)

In a Flexible Storage Sharing (FSS) environment, if you change the naming scheme on a node that has local disks, the remote disk names are not reflected with the corresponding name change. If you change the naming scheme on a node where exported disks are present, to reflect the updated remote disk names, you must either export the disks again or restart the node where the remote disks are present

Workaround:

There is no workaround for this issue.

vxassist does not create data change logs on all mirrored disks, if an FSS volume is created using DM lists (3559362)

When a Flexible Storage Sharing (FSS) volume is created using DM lists, the `vxassist` command does not create data change logs on all the mirrored disks; the number of DCO mirrors is not equal to the number of data mirrors. The `vxassist` command creates a two-way DCO volume.

Workaround:

Manually add a DCO mirror using the `vxassist -g diskgroup mirror dco_volume` command.

Intel SSD cannot be initialized and exported (3584762)

Initializing an Intel SSD with the Flexible Storage Sharing (FSS) export option may fail with the following error message:

```
VxVM vxedpart ERROR V-5-1-10089 partition modification failed: Device
or resource busy
```

Workaround:

Initialize the private region of the SSD disk to 0 and retry the disk initialization operation.

For example:

```
# dd if=/dev/zero of=/dev/vx/dmp/intel_ssd0_0 bs=4096 count=1

# vxdisksetup -i intel_ssd0_0 export
```

VxVM may report false serial split brain under certain FSS scenarios (3565845)

In a Flexible Storage Sharing (FSS) cluster, as part of a restart of the master node, internal storage may become disabled before network service. Any VxVM objects on the master node's internal storage may receive I/O errors and trigger an internal transaction. As part of this internal transaction, VxVM increments serial split brain (SSB) ids for remaining attached disks, to detect any SSB. If you then disable the network service, the master leaves the cluster and this results in a master takeover. In such a scenario, the master takeover (disk group re-import) may fail with a false split brain error and the `vxsplitlines` output displays 0 or 1 pools.

For example:

```
Syslog: "vxvm:vxconfigd: V-5-1-9576 Split Brain. da id is 0.2,
while dm id is 0.3 for dm disk5mirr
```

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
   ```

The fsappadm subfilemove command moves all extents of a file (3258678)

This issue occurs under following conditions:

- You run the `fsppadm subfilemove` command from a cluster file system (CFS) secondary node.
- You specify a range of extents for relocation to a target tier.

If the extent size is greater than or equal to 32768, the `fsppadm subfilemove` command moves all extents of the specified table to the target tier. The expectation is to move a specified range of extents.

Workaround:

- On the CFS primary node, determine the primary node using one of the following commands:

```
# fsclustadm showprimary mountpoint

# fsclustadm idtoname nodeid
```

A volume remains in DETACHED state even after storage nodes join back to the cluster (3628933)

This issue occurs in an FSS configuration, in the following scenario:

1. The volume is part of an FSS disk group with storage from only a subset of the nodes.
2. The storage nodes fail or are rebooted while I/O is in progress on all the nodes.
3. The node in the cluster that is not contributing to storage becomes the master.
4. The storage nodes come up and join the cluster.

The issue is that the volume remains in a detached state even after the storage nodes rejoin the cluster. Trying to start the volume manually with the following command generates an error:

```
# vxvol start -g dg_name vol_name
VxVM vxvol ERROR V-5-1-10128  DCO experienced
IO errors during the operation.
Re-run the operation after ensuring that DCO is accessible
```

Workaround:

Deport the disk group and then import the disk group.

The cluster may hang due to a known lock hierarchy violation defect (2919310)

If VxFS File Change Log (FCL) is turned ON in Cluster File System (CFS) environments, a known lock hierarchy violation defect may lead to the cluster hang.

Workaround:

There is no workaround for this issue.

Certain I/O errors during clone deletion may lead to system panic. (3331273)

Certain I/O errors during clone deletion may lead to system panic.

Workaround:

There is no workaround for this issue.

Veritas Volume Manager known issues

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

The vxconfigd daemon fails to start after machine reboot (3566713)

The `shutdown -r` command makes sure that the file contents on the OS file system are written properly to the disk before a reboot. The `volboot` file is created in the OS file system, and is used to bring up the `vxconfigd` daemon after the system reboot. If the machine reboots for any reason without proper shutdown, and the `volboot` file contents are not flushed to the disk, `vxconfigd` will not start after the system reboots.

Workaround:

You must rerun the `vxinstall` script to re-create the `volboot` file and to start the `vxconfigd` daemon and other daemons.

Root disk encapsulation and its dependent operations are not supported for enclosure-based naming scheme (3623681)

When the naming scheme for Veritas Volume Manager (VxVM) is set to enclosure-based naming (EBN), root disk encapsulation is not supported. Operations such as mirroring, splitting, or joining on encapsulated root disks are also not supported.

Root disk encapsulation is only supported if the device naming scheme is set to operating system-based naming (OSN) and the persistence attribute is set to yes.

Workaround:

Before encapsulating a root disk, use the following command to set the VxVM device naming scheme to OSN and the persistence attribute to yes.

```
# vxddladm set namingscheme=osn persistence=yes
```

Restarting the vxconfigd daemon on the slave node after a disk is removed from all nodes may cause the disk groups to be disabled on the slave node (3591019)

The issue occurs if the storage connectivity of a disk is removed from all the nodes of the cluster and the `vxconfigd` daemon is restarted on the slave node before the disk is detached from the slave. All the disk groups are in the `dgdisabled` state on the slave node, but show as `enabled` on the other nodes.

If the disk was detached before the `vxconfigd` daemon is restarted, the issue does not occur.

In a Flexible Storage Sharing (FSS) environment, removing the storage connectivity on a node that contributes DAS storage to a shared disk group results in global connectivity loss because the storage is not connected elsewhere.

Workaround:

To prevent this issue:

Before restarting the `vxconfigd` daemon, if a disk in a shared disk group has lost connectivity to all nodes in the cluster, make sure that the disk is in the `detached` state. If a disk needs to be detached, use the following command:

```
# vxdisk check diskname
```

To resolve the issue after it has occurred:

If `vxconfigd` is restarted before the disks got detached, remove the node from the cluster and rejoin the node to the cluster.

SmartIO VxVM cache invalidated after relayout operation (3492350)

If a relayout operation is done on a volume that has SmartIO VxVM caching enabled, the contents of the cache for the volume may be invalidated.

Workaround:

This behavior is expected. There is no workaround.

Creating a disk group with a large number of objects or splitting, joining, or moving such a disk group reports an out of kernel memory error (3069711)

When you create a disk group with an extremely large number of objects (volumes, snapshots, plexes, disks), you may see the following error:

```
ERROR-V-5-1-10128 Out of kernel memory
```

You may also see the error when you perform operations like split/join/move on such a disk group.

Each object has a record which is used for its description and state. These records are stored in the private region of every disk group. The default private region size is 32 MB which can accommodate a sufficient number of objects. If the private region of disk group does not have space to create a new record, the operation fails with the above error message. Typical use cases would not hit this condition.

Workaround:

The best practice is not to have an extremely large number of objects in the disk group. Instead, split the disk group into multiple disk groups.

Refer to the section “Reorganizing the contents of disk groups” in the *Administrator's Guide* for information about splitting disk groups.

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

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

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

Workaround:

To set the `disktimeout` to 200:

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

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 *Symantec 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.

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

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

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

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

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

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

- The shared disk groups on the disconnected storage are marked as dgdisabled on the slave node only.
- The shared disk groups are available to rest of the cluster nodes but no transactions, such as VxVM configuration changes, are possible on any shared disk group.
- Attempts to deport such shared disk groups will fail.

Workaround:

Do one of the following:

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

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

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

Re-enabling connectivity if the disks are in local failed (lfailed) state (2425977)

In a Cluster Volume Manager (CVM) cluster, you can disable connectivity to the disks at the controller or enclosure level with the `vxdmpadm disable` command. In this case, CVM may place the disks into the `lfailed` state. When you restore connectivity with the `vxdmpadm enable` command, CVM may not automatically clear the `lfailed` state. After enabling the controller or enclosure, you must run disk discovery to clear the locally failed state.

To run disk discovery

- Run the following command:

```
# vxdisk scandisks
```

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.

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

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

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

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

Workaround:

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

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

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

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

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

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

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

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

1. Unroot the encapsulated root disk:

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

2. Upgrade the kernel:

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

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

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

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

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

Workaround: The following procedure resolves this issue.

To resolve this issue

1 Set the `iotimeout` tunable to 600:

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

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

```
# vxdctl enable
```

Upgrading from Symantec Storage Foundation Cluster File System High Availability 5.x to 6.2 may fail for IBM XIV Series arrays (2715119)

Starting in the Symantec Storage Foundation Cluster File System High Availability 5.1 SP1 release, the Array Support Library (ASL) for the IBM XIV enclosures converts the LUN Serial Number from hexadecimal to decimal. Because of this change, the enclosure names differ from releases prior to the 5.1 SP1 releases. When you upgrade Symantec Storage Foundation Cluster File System High Availability from a release prior to that release to the current 6.2 release, XIV LUNs may go into an error state. Note that the latest RPs on 5.1/5.1SP1 are already modified to use the same logic for enclosure naming.

Workaround:

After the upgrade, run `vxddladm assign names`.

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

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

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

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

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

Workaround:

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

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

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

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

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

Workaround

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

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

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

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

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

Workaround:

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

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

Disk group import of BCV LUNs using -o updateid and -ouseclonedev options is not supported if the disk group has mirrored volumes with DCO or has snapshots (2831658)

VxVM uses guid stored in configuration to uniquely identify all objects. The data change object (DCO) volume stores the guid of mirrors and snapshots. If the disk group is imported with `-o updateid` and `-o useclonedev`, it changes the guid of objects in VxVM configuration database and the guids stored in the DCO volume are not updated. The operations involving DCO cannot find objects with the stored guid. This could lead to failure of certain operations involving DCO or could lead to unexpected behavior.

Workaround:

No workaround available.

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

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

Workaround:

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

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

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

Workaround:

There is no workaround available.

Dynamic LUN expansion is not supported for EFI disks in simple or sliced format and non-EFI disks greater than 1TB in simple or sliced format.(2836798)

Dynamic LUN expansion is not supported for EFI (Extensible Firmware Interface) disks in simple or sliced format and non-EFI disks greater than 1TB in simple or sliced format. The recommended format is the Cross-platform Data Sharing (CDS) disk format.

Workaround:

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

CVMVolDg agent may fail to deport CVM disk group when CVMDeportOnOffline is set to 1

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

Workaround: If multiple CVMVolDg resources are configured for a shared disk group and the disk group is required to be deported during offline, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

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

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

Workaround: There is no workaround for this issue.

vxresize does not work with layered volumes that have multiple plexes at the top level (3301991)

If a layered volume has multiple plexes at the top level, `vxresize` does not work. For example, if you add a mirror to a concat-mirror volume for a third mirror snapshot. The `vxresize` operation fails with the following message:

```
VxVM vxassist ERROR V-5-1-2528 Volume volname built on layered volumes
have multiple plexes
```

```
VxVM vxresize ERROR V-5-1-4703 Problem running vxassist command for
volume volname, in diskgroup dgname
```

Workaround:

To resize the volume

1. After adding the mirror to the volume, take a snapshot using the plex.
2. Grow the volume and snapshot volume with `vxresize`
3. Reattach the snapshot volume to the source volume.

Importing a clone disk group fails after splitting pairs (3134882)

When you import a clone disk group with the `-o updateid` option, the GUIDs of all the objects are assigned new values. However, these values are not updated on the maps in the data change object (DCO). When you initiate a volume recovery, it fails on the volumes having instant DCO (version >= 20) because it does not find the objects corresponding to the GUIDs. In this situation, the DCO is considered corrupt and the volume remains inaccessible.

Workaround: You mainly need the `-o updateid` option when you import the clone disk group on the same host as the primary disk group. You can avoid using the option by doing one of the following:

- Import the clone disk group on a different host.
- Deport the primary disk group before you import the clone disk group.

If the import of the clone disk group with `-o updateid` option or the recovery of volume thereafter fails with a message about the DCO being corrupted, this error occurs because the GUIDs are not being updated on the DCO implicitly. If the workaround is not acceptable and you need to access the volume, you can remove the DCO. You can dissociate or remove the snapshots and then remove the DCO manually to let the recovery proceed.

The DMP EMC CLARiiON ASL does not recognize mirror view not ready LUNs (3272940)

On hosts that have EMC CLARiiON mirror view not ready LUNs, if you enable or disable the switch port and then issue the `vxdisk scandisks` or `vxdctl enable` command, I/O error messages are written continuously in the syslog.

The dynamic multi-pathing (DMP) request for providing information to identify mirror view not ready LUNs through in-band SCSI command is pending with EMC engineering. Not ready LUNs are special kind of LUNs which reject all kinds of I/O requests.

Because DMP does not recognize not ready LUNs, Veritas Volume Manager (VxVM) tries to bring them online. As part of the online process, VxVM issues I/Os to read the disk private region. These I/Os fail and generate error messages in syslog.

Because of events that are generated as part of the online process, the `vxattachd` script triggers the vxdisk scandisks command again. This cycle causes continuous I/O error messages. This problem can also cause other commands to run slowly because the VxVM configuration daemon (`vxconfigd`) is busy servicing `vxdisk scandisks`.

Workaround: Stop the `vxattachd` script and set EMC CLARiiON values, as follows:

1. Disable the `vxattachd` process.

   For more information on how to disable `vxattachd` and what features you lose if `vxattachd` is disabled, see the `vxattachd` man page

2. Set the following EMC CLARiiON values:

   - `recoveryoption=fixedretry`
   - `retrycount=5`

   Enter:

   `vxdmpadm setattr enclosure` *`enclosure_name`* `recoveryoption=fixedretry \ retrycount=5`

When multiple backups are taken with the vxconfigbackup command within a short period of time, the vxconfigrestore operation may restore an older configuration (3331769)

If you take multiple backups with the `vxconfigbackup` command within a short period of time, the `vxconfigrestore` command may restore an older configuration of the disk group.

Workaround:

There is no workaround.

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

The Veritas Volume Manager (VxVM) 6.2 includes several array names that differ from the array names in releases 5.1SP1 or prior. Therefore, if you upgrade to VxVM 6.2 from a release 5.1SP1 or earlier, 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 6.2.

Workaround:

Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-15 shows the Hitachi arrays that have new array names.

Table 1-15 Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC
TagmaStoreUSPV	Hitachi_USP-V
TagmaStoreUSPVM	Hitachi_USP-VM
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 6.2. 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

Running the vxdisk *disk* set clone=off command on imported clone disk group luns results in a mix of clone and non-clone disks (3338075)

If you do not specify a disk group name, the `vxdisk set` operation works on the `dmname` rather than the `daname`. If a `dmname` is the same as an existing `daname`, the vxdisk set operation reflects on the dm name.

Workaround: Use the following command syntax to set the attributes:

```
vxdisk -g diskgroup_name set dmname clone=off
```

For example:

```
vxdisk -g dg1 set eva4k6k0_12 clone=off
```

vxunroot cannot encapsulate a root disk when the root partition has XFS mounted on it (3614362)

If the root partition has the XFS file system mounted on it, you cannot change the root partition's Universally Unique IDentifier (UUID). However, changing the UUID of the partitions of the root disk is necessary in root disk encapsulation. Given the limitation above, Symantec does not support root disk encapsulation where the root partition has an XFS file system.

Workaround:

None.

VxVM fails to recognize iSCSI LUNs after a restart (3631990)

During a restart, you may see a large number of failover messages. After the restart, the output of the `vxdisk list` command may not show iSCSI LUNs. In Red Hat Enterprise Linux (RHEL) 7, the `systemd` system and service manager may stop the `iscsid` service before Veritas Volume Manager (VxVM) services stop. As a result, the LUNs are not visible to VxVM immediately before shutdown, causing a large number of error messages. Similarly, `systemd` may start the `iscsid` service after VxVM services. As a result, VxVM does not discover the LUNs.

Workaround:

The workaround is applicable only if root disk encapsulation is not in use and `dmp_native_support` is disabled.

1. Create the iSCSI directory as follows:

   ```
   # mkdir /etc/systemd/system/iscsi.service.d/
   ```

2. Write the unit file for the iSCSI daemon:

   ```
   # cat /etc/systemd/system/iscsi.service.d/vxvm.conf
   [Unit]
   Before=vxvm-boot.service
   ```

3. Reload `systemd`:

   ```
   # systemctl daemon-reload
   ```

Veritas File System known issues

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

The command tab auto-complete fails for the `/dev/vx/` file tree; specifically for RHEL 7 (3602082)

The command tab auto-complete operation fails because the following RPM is installed on the machine:

"bash-completion-2.1-6.el7.noarch"

This somehow overwrites the default auto-complete rules. As a result, some issues are observed with the VxFS commands. However, the issue is not observed with all the VxFS commands. The issue is observed with the `mkfs`(1M) command, but is not observed with the `mount`(1M) command.

Workaround: Please remove the "bash-completion-2.1-6.el7.noarch" RPM, so that the command tab auto-complete does not fail for the `/dev/vx/` file tree.

When hard links are present in the file system, the sfcache list command shows incorrect cache usage statistics (3059125)

If a hard link is present for a file that is loaded in the cache, the `sfcache list` command shows the cache usage for both files: the original file and the hard link. The resulting statistics are incorrect, because the cache usage is shown to be twice the actual usage.

For example:

```
# sfcache list -r /mnt1
/mnt1:
CACHE-USED(MB) MODE PINNED NAME
0 read no /mnt1/test_10
0 read no /mnt1/test_20
0 read no /mnt1/test_50
0 read no /mnt1/test_100
0 read no /mnt1/test_200
0 read no /mnt1/test_300
0 read no /mnt1/test_400
500 read yes /mnt1/test_500
0 read no /mnt1/test_1024
500 read yes /mnt1/dir/hardlink
```

```
500 read no /mnt1/dir
1000 read no /mnt1

# sfcache list fs1

Cachearea: fs1
Assoc Type: AUTO
Type: VxFS
Size: 1.00g
State: ONLINE
/dev/vx/dsk/sfcache_defaultdg/fs1:
FSUUID CACHE-USED(MB) MODE MOUNTPOINT
23642651-81a5-0d00-1a26-0000911ec26c 1000 read /mnt1
```

Workaround: There is no workaround for this issue.

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 RHEL5 and RHEL6 (2560357)

For RHEL5 and 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 the sleep locks. However, the task is not hung and the messages can be safely ignored.

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

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

Deduplication can fail with error 110 (3741016)

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.

The system panics with the panic string "kernel BUG at fs/dcache.c:670!" (3323152)

The umount of the file system under high-memory-pressure condition may lead to a system panic. The panic string is displayed as following: "kernel BUG at fs/dcache.c:670!"

Workaround: There is no workaround for this issue.

A restored volume snapshot may be inconsistent with the data in the SmartIO VxFS cache (3193525)

The data in a volume snapshot may have data that is inconsistent with the VxFS level SmartIO cache. When the volume snapshot is restored and mounted, then before using that file system you should purge the corresponding cache data. Or, disable the caching for that file system.

Workaround:

Purge the file system data from the SmartIO cache after restoring the volume snapshot.

```
# sfcache purge {mount_point|fsuuid}
```

Full file system check takes over a week (2628207)

On a large file system with many Storage Checkpoints, a full file system check using the`fsck_vxfs`(1M) command might appear to be hung. The `fsck` command is not actually hung; the process can take an extremely long time to complete.

Workaround: There is no workaround for this issue.

The file system may hang when it has compression enabled (3331276)

In a VxFS file system that has compression enabled, a deadlock in page fault handler can lead to the file system hang.

Workaround:

There is no workaround for this issue.

The file system may hang due to file system full conditions when file level snapshots are present (2746259)

In the presence of file level snapshots, file system full conditions may lead to the file system hang. Following a reboot, a mount may hang as well.

Workaround:

There is no workaround for this issue.

The file system may be marked for full fsck during a clone removal (2977828)

Under low memory conditions, a clone removal may lead to file system being marked for full fsck.

Workaround:

A full fsck of the file system will be required to recover the file system.

I/O errors on the file system may lead to data inconsistency (3331282)

If there are writable clones on the file system, I/O errors may lead to data inconsistency.

Workaround:

Run a full `fsck` to recover the file system.

Forcing the system to unmount during heavy I/O load may result in system panic in vx_is_fs_disabled_impl (3331284)

Forcing the system to unmount during heavy I/O load may result in system panic in `vx_is_fs_disabled_impl` .

Workaround:

There is no workaround for this issue.

During a deduplication operation, the spoold script fails to start (3196423)

This issue occurs because a port is not available during the operation; therefore the spoold script fails to start with the the following error:

```
DEDUP_ERROR INIT: exec spoold failed (1024)
```

Workaround:

Check the `spoold.log` file for specific error messages, and if the log indicates a port is not available, you can check if the port is in use with the `netstat/lsof` command. If the port is not open, you can retry the deduplication operation; if the port is open, you can wait for the port to close, and then try the deduplication operation again.

For example, the following error message in the `spoold.log` file indicates that port `51003` is not available:

```
ERR [140399091685152]: -1: NetSetup: NetBindAndListen returned error,
unable to bind to localhost:51003
```

When in-place and relocate compression rules are in the same policy file, file relocation is unpredictable (3278193)

You cannot have in-place compress/uncompress rules and relocate compress/uncompress rules in the same policy file. If they are in the same file, file relocation is unpredictable.

Workaround: Create a different policy file for each policy, and enforce the policy as per the required sequence.

“rpc.statd” in the “nfs-utils” RPM in the various Linux distributions does not properly cleanse the untrusted format strings (3335691)

"rpc.statd” in the “nfs-utils” RPM in various Linux distributions does not properly cleanse untrusted format strings. This vulnerability may allow remote attackers to gain root privileges.

Workaround: Update to version 0.1.9.1 of the “nfs-utils” RPM to correct the problem.

In a CFS cluster, that has multi-volume file system of a small size, the fsadm operation may hang (3348520)

In a CFS cluster, that has multi-volume file system of a small size, the fsadm operation may hang, when the free space in the file system is low.

Workaround: There is no workaround for this issue.

The file system deduplication operation fails with the error message "DEDUP_ERROR Error renaming X checkpoint to Y checkpoint on filesystem Z error 16" (3348534)

The file system deduplication operation fails with the error message "DEDUP_ERROR Error renaming X checkpoint to Y checkpoint on filesystem Z error 16", due to the failure in unmounting the checkpoint.

Workaround: Retry the deduplication operation to resolve the problem.

Replication known issues

This section describes the replication known issues in this release of Symantec Storage Foundation Cluster File System High Availability.

Transactions on VVR secondary nodes may timeout waiting for I/O drain [3236772]

If the VVR secondary node receives updates out of order from the Primary, and a transaction starts on the secondary site, then the transaction may timeout waiting for I/O drain. This issue may occur in situations where the gaps created by out of order updates are not filled within the transaction timeout period.

Workaround:

Pause replication and make configuration changes.

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

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

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

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

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

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

Workaround:

To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

vradmin syncvol command compatibility with IPv6 addresses (2075307)

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

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

RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2036605)

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" (2036644)

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
   ```

vxassist relayout removes the DCM (145413)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

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

3. Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7. Start the RVG. Enter the following:

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

8. Resume or start the applications.

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

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

Workaround:

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

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

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

3. Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7. Start the RVG. Enter the following:

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

8. Resume or start the applications.

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

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

vradmin functionality may not work after a master switch operation (2163712)

In certain situations, if you switch the master role, `vradmin` functionality may not work. The following message displays:

```
VxVM VVR vxrlink ERROR V-5-1-15861  Command is not supported for
command shipping. Operation must be executed on master
```

Workaround:

To restore vradmin functionality after a master switch operation

1 Restart `vradmind` on all cluster nodes. Enter the following:

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

2 Re-enter the command that failed.

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

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

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

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

Workaround: Manually edit the following files to update the script:

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

In each file, modify the following line:

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

to

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

RLINK name cannot exceed 31 characters

The `vradmin` utility truncates the RLINK name to 31 characters, as the `vxmake` utility does not support the creation of RLINK names that are longer than 31 characters.

Workarounds:

- Specify the `prlink` and `srlink` attributes using the `vradmin addsec` command, so you can choose the RLINK name in the `addsec` command line.
- If using IPv6 addresses, create host name aliases for the IPv6 addresses and specify the aliases in the `addsec` command line.

Plex reattach operation fails with unexpected kernel error in configuration update (2791241)

In a VVR environment with layered volumes, if a DCM plex becomes detached because of a storage failure, reattaching the plex after fixing the storage issue fails with the following error:

```
VxVM vxplex ERROR V-5-1-10128  Unexpected kernel error in configuration
update
```

Workaround:

There is no workaround for this issue.

While vradmin commands are running, vradmind may temporarily lose heartbeats (3347656)

This issue may occasionally occur when you use `vradmin` commands to administer Volume Replicator (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 the 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 replicated data set (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 stop
   # /etc/init.d/vras-vradmind.sh start
   ```

The vradmin repstatus command does not show that the SmartSync feature is running (3345984)

In a Volume Replicator (VVR) environment, after you start the initial synchronization with the `vradmin -a startrep` command with file system mounted on the primary data volumes, the `vradmin repstatus` command does not show that the SmartSync feature is running. This is an only issue with the output of the `vradmin repstatus` command.

Workaround:

To confirm that SmartSync is running, enter:

```
vxrlink status rlink
```

Write I/Os on the primary logowner may take a long time to complete (2622536)

Under a heavy I/O load, write I/Os on the Volume Replicator (VVR) primary logowner take a long time to complete.

Workaround:

There is no workaround for this issue.

Bunker replay does not occur with volume sets (3329970)

There are issues with bunker replication using Volume Replicator (VVR) with volume sets. Do not upgrade to SFCFSHA 6.2 if you have configured or plan to configure bunker replication using VVR with volume sets.

Workaround:

Contact Symantec Technical Support for a patch that enables you to use this configuration.

SmartIO does not support write-back caching mode for volumes configured for replication by Volume Replicator (3313920)

SmartIO does not support write-back caching mode for volumes that are configured for replication by Volume Replicator (VVR).

Workaround:

If you have configured volumes for replication by VVR, do not enable write-back caching

During moderate to heavy I/O, the vradmin verifydata command may falsely report differences in data (3270067)

While an application is online at the Volume Replicator primary site, the `vradmin verifydata` command may fail. The command output shows the differences between the source data volume and the target data volume.

Workaround:

The reason for this error is that the cache object that is used for the verification might be under allocated. You might need to allocate more space for the shared cache object. For guidelines on shared cache object allocation, see the section "Creating a shared cache object" in the *Symantec Storage Foundation Administrator's Guide*.

After performing a CVM master switch on the secondary node, both rlinks detach (3642855)

If the VVR logowner (master) node on the secondary site goes down during initial synchronization, then during the RVG recovery (initiated on any secondary site node as a result of node crash), the replication links detach with the following error:

```
WARNING: VxVM VVR vxio V-5-0-187 Incorrect magic number or unexpected
upid (1) rvg rvg1
WARNING: VxVM VVR vxio V-5-0-287 rvg rvg1, SRL srl1: Inconsistent log
 - detaching all rlinks.
```

Workaround:

Restart replication using the autosync operation.

`vradmin -g dg repstatus` rvg displays the following configuration error: vradmind not reachable on cluster peer (3648854)

`vradmin -g dg rep` status rvg displays the following configuration error:

vradmind is not reachable on the cluster peer

However, replication is an ongoing process. The reason is that an unclean disconnect left the `vradmind` port open and in the **TIME_WAIT** state. An instance is as following:

```
# netstat -n | grep 8199
tcp        0      0 1:44781    1:8199
TIME_WAIT
```

```
tcp        0      0 1:44780     1:8199
TIME_WAIT
```

The following error message appear in **/var/vx/vras/log/vradmind_log_A**:

```
VxVM VVR Notice V-5-20-0 TAG_D IpmHandle:recv peer closed errno=0
VxVM VVR Debug V-5-20-8690 VRASCache TAG_E Cache_RLink
repstatus UPDATE message created for rlink rlk_192.168.111.127_rvg1
VxVM VVR Warning V-5-20-0 TAG_C IpmHandle::handleTo
vvr_sock_host_serv failed for l111031
VxVM VVR Warning V-5-20-0 TAG_C IpmHandle::open: getaddrinfo
error(could not resolve srchost l111032, error: Connection refused)
```

Workaround: Restart the `vradmind` daemon.

```
/etc/init.d/vras-vradmind.sh stop
/etc/init.d/vras-vradmind.sh start
```

LLT known issues

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

Dependency may get overruled when uninstalling multiple RPMs in a single command [3563254]

When performing uninstallation of multiple RPMs through a single comment, the system identifies and follows the specified dependency among the RPMs as the uninstallation progresses. However, if the pre-uninstallation script fails for any of the RPMs, the system does not abort the task but instead uninstalls the remaining RPMs.

For example, if you run `rpm -e VRTSllt VRTSgab VRTSvxfen` where the RPMs have a dependency between each other, the system bypasses the dependency if the pre-uninstallation script fails for any RPM.

Workaround: Uninstall the RPMs independently.

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 `lltab` file, and update the file before you restart LLT.

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

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

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

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

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

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

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

GAB known issues

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

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

After you unregister the gtx port with `-nodeinit` option, the `gabconfig -C` command shows refcount as 1. But when forceful `deinit` option (`gabdebug -R`

`GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

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

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

Workaround: There is no workaround for this issue.

Cluster panics during reconfiguration (2590413)

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

Workaround: There is no workaround for this issue.

I/O fencing known issues

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

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 *Symantec Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Workaround: Start VxFEN again after some time.

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

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

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

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

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

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

Unable to customize the 30-second duration (2551621)

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

Workaround: There is no workaround for this issue.

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

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSat RPM is not removed from the system, the `cpsadm` command loads the old security libraries present on the system. As the installer runs the `cpsadm` command on the CP server to add or upgrade the SFCFSHA cluster (application cluster), the installer also fails.

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

To resolve this issue

1 Rename `cpsadm` to `cpsadmbin`:

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

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

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

3 Change the permissions of the new file to 775:

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

Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later [3226290]

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

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

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

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

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

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

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

Workaround: To resolve this issue:

1. Activate the storage domain in RHEV-M.
2. Check that the data center is in the up state.

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

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

Workaround:

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

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

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

Workaround: There is no workaround for this issue.

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

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

Workaround:

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

The upper bound value of FaultTolerance attribute of CoordPoint agent should be less than the majority of the coordination points. (2846389)

The upper bound value of `FaultTolerance` attribute of `CoordPoint` agent should be less than the majority of the coordination points. Currently this value is less than the number of coordination points.

The vxfenswap utility deletes comment lines from the `/etc/vxfemode` file, if you run the utility with hacli option (3318449)

The vxfenswap utility uses RSH, SSH, or hacli protocol to communicate with peer nodes in the cluster. When you use vxfenswap to replace coordination disk(s) in disk-based fencing, vxfenswap copies `/etc/vxfenmode` (local node) to `/etc/vxfenmode` (remote node).

With the hacli option, the utility removes the comment lines from the remote `/etc/vxfenmode` file, but, it retains comments in the local `/etc/vxfenmode` file.

Workaround: Copy the comments manually from local `/etc/vxfenmode` to remote nodes.

When you configure CP server only for HTTPS-based communication, the `engine_A.log` displays a misleading message (3321101)

The `engine_A.log` file displays the following message when you configure CP server only for HTTPS-based communication but not for IPM-based communication.

```
No VIP for IPM specified in /etc/vxcps.conf
```

Workaround: Ignore the message.

The CoordPoint agent faults after you detach or reattach one or more coordination disks from a storage array (3317123)

After you detach or reattach a coordination disk from a storage array, the CoordPoint agent may fault because it reads an older value stored in the I/O fencing kernel module.

Workaround: Run the `vxfenswap` utility to refresh the registration keys on the coordination points for both server-based I/O fencing and disk-based I/O fencing.

But, even if the registrations keys are not lost, you must run the `vxfenswap` utility to refresh the coordination point information stored in the I/O fencing kernel module.

For more information on refreshing registration keys on the coordination points for server-based and disk-based I/O fencing, refer to the *Symantec Cluster Server Administrator's Guide*.

CP server does not allow adding and removing HTTPS virtual IP or ports when it is running [3322154]

CP server does not support adding and removing HTTPS virtual IPs or ports while the CP server is running. However, You can add or remove the IPM virtual IPs or ports.

Workaround: No workaround. If you want to add a new virtual IP for HTTPS, you must follow the entire manual procedure for generating HTTPS certificate for the CP server (server.crt), as documented in the *Symantec Cluster Server Installation Guide*.

The `vxfentsthdw` utility may not run on systems installed with partial SFHA stack [3333914]

The `vxfentsthdw` utility runs if the SFHA stack and VCS are fully installed with properly configured SF and VxVM. It also runs if the entire SFHA stack and VCS are not installed. However, partial installs where SF is installed and configured but VCS is not installed is not supported. The utility will display an error with the `-g` or `-c` options.

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

When a client node goes down, for reasons such as node panic, I/O fencing does not come up on that client node after node restart (3341322)

This issue happens when one of the following conditions is true:

- Any of the CP servers configured for HTTPS communication goes down.
- The CP server service group in any of the CP servers configured for HTTPS communication goes down.
- Any of the VIPs in any of the CP servers configured for HTTPS communication goes down.

When you restart the client node, fencing configuration starts on the node. The fencing daemon, vxfend, invokes some of the fencing scripts on the node. Each of

these scripts has a timeout value of 120 seconds. If any of these scripts fails, fencing configuration fails on that node.

Some of these scripts use `cpsadm` commands to communicate with CP servers. When the node comes up, `cpsadm` commands try to connect to the CP server using VIPs for a timeout value of 60 seconds. So, if the multiple `cpsadm` commands that are run within a single script exceed the timeout value, then the total timeout value exceeds 120 seconds, which causes one of the scripts to time out. Hence, I/O fencing does not come up on the client node.

Note that this issue does not occur with IPM-based communication between CP server and client clusters.

Workaround: Fix the CP server.

Symantec Storage Foundation for Databases (SFDB) tools known issues

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

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SFCFSHA.

Workaround:

There is no workaround at this point of time.

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.

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

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

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

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

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

- BALANCE
- CHECKPOINT
- METADATA

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

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

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

Workaround:

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

- For FlashSnap, resync the snapshot and try the clone operation again.
- For FileSnap and Database Storage Checkpoint, 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.

Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.2 (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 6.2.

When upgrading from SFCFSHA version 5.0 to SFCFSHA 6.2 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

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

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

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

If you have a parameter, such as `log_archive_dest_1`, in single line in the `init.ora` file, then `dbed_vmclonedb` works but `dbed_vmcloneb` fails if you put in multiple lines for parameter.

Workaround: Edit the PFILE to arrange the text so that the parameter values are on a single line. If the database uses a spfile and some parameter values are spread across multiple lines, then use the Oracle commands to edit the parameter values such as they fit in a single line.

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.

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

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

You might encounter an Oracle error such as the following:

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

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

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

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

Workaround: Retry the cloning operation until it succeeds.

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

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

There is no workround at this point of time.

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

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

Workaround: There is no workaround for this issue.

Flashsnap clone fails under some unusual archivelog configuration on RAC (2846399)

In a RAC environment, when using FlashSnap, the archive log destination to snapshot must be a shared path, and must be the same across all the nodes. Additionally, all nodes must use the same archive log configuration parameter to specify the archive log destination. Configurations similar to the following are not supported:

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_2='location=/tpcc_arch'
tpcc3.log_archive_dest_3='location=/tpcc_arch'
```

Where tpcc1, tpcc2, and tpcc3 are the names of the RAC instances and /tpcc_arch is the shared archive log destination.

Workaround: To use FlashSnap, modify the above configuration to *.log_archive_dest_1='location=/tpcc_arch'. For example,

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_1='location=/tpcc_arch'
tpcc3.log_archive_dest_1='location=/tpcc_arch'
```

The dbdst_obj_move(1M) command moves all the extents of a database table (3277003)

The `dbdst_obj_move(1M)` command moves all the extents of a database table when:

- The `dbdst_obj_move(1M)` command is run from the CFS secondary node.
- The object is an Oracle database table (-t option)
- A range of extents is specified for movement to a target tier (-s and -e options). The `dbdst_obj_move(1M)` command moves all extents of the specified table to a target tier when the extent size is greater than or equal to 32768. However, the expectation is to move only a specified range of extents.

Workaround: Run the `dbdst_obj_move(1M)` command from the CFS primary node.

Use the `fsclustadm showprimary` *<mountpoint>* and `fsclustadm idtoname` *<nodeid>*commands to determine the mode of a CFS node.

When you attempt to move all the extents of a table, the `dbdst_obj_move`(1M) command fails with an error (3260289)

When you attempt to move all the extents of a database table, which is spread across multiple mount-points in a single operation, the `dbdst_obj_move(1M)` command fails. The following error is reported:

```
bash-2.05b$ dbdst_obj_move -S sdb -H $ORACLE_HOME -t test3 -c MEDIUM
FSPPADM err : UX:vxfs fsppadm: WARNING: V-3-26543: File handling failure
on /snap_datadb/test03.dbf with message -
SFORA dst_obj_adm ERROR V-81-6414 Internal Error at fsppadm_err
```

Note: To determine if the table is spread across multiple mount-points, run the dbdst_obj_view(1M) command

---

Workaround: In the `dbdst_obj_move(1M)` command, specify the range of extents that belong to a common mount-point. Additionally, if your table is spread across "n" mount-points, then you need to run the `dbdst_obj_move(1M)` command "n" times with a different range of extents.

The ReverseResyncBegin (RRBegin) operation with recovery option as AUTO fails (3076583)

The RRBegin operation with the recovery option as AUTO fails when you perform the following sequence of operations:

1. Validate the FlashSnap setup using the validate operation.
2. In the database, take the tablespace offline.
3. Perform a snapshot operation.

4 Bring the tablespace online which was taken offline in 2.

5 Perform the Reverse Resync Begin operation.

Note: This issue is encountered only with Oracle version 10gR2.

Workaround: Perform one of the following:

- Make sure to bring the tablespace online only after performing the RRBegin and RRCommit operations. Otherwise, perform the Reverse Resync Begin operation while the tablespace is in the offline mode.
- To recover a database, specify the recovery option as **AUTO_UNTIL_SCN** in the RRBegin operation.

The ReverseResyncBegin (RRBegin) operation fails when performed on multiple snapshot configurations (3066532)

When you perform a Reverse Resync operation on multiple snapshot configurations, SFDB reports the following error message:

```
$ vxsfadm -a oracle -s flashsnap --name \
man -o rrbegin

SFDB vxsfadm ERROR V-81-0943 Repository already relocated to alternate
location.
```

As per the Reverse Resync design, the first RRBegin operation relocates the SFDB repository to a backup location, and the ReverseResyncAbort and ReverseResyncCommit operations restore it to the original location. When the second RRBegin operation attempts to relocate the same repository which is already relocated, SFDB reports the error message.

Workaround: Make sure to perform the RRAbort or RRCommit operation using the snapshot configuration that is in the RRBegin state.

Note: You must complete Reverse Resync operations for a particular configuration before you start with another configuration.

The ReverseResyncBegin (RRBegin) operation fails and reports an error message due to a missing binary control file (3157314)

When the RRBegin operation cannot find the binary control file that is used to recover a database instance, it reports the following error message:

```
[oracle@testbox ~]$ vxsfadm -a oracle -s flashsnap -name man -o rrbegin

SFDB vxsfadm ERROR V-81-0949 Binary Control file is not available for
recovery purposes
```

This issue is observed in the third-mirror break-off type (FlashSnap) snapshots that are created using the older SFDB version, which did not include the binary control file in the snapshot images.

Workaround: There is no workaround for this issue.

Sometimes SFDB may report the following error message: SFDB remote or privileged command error (2869262)

While using SFDB tools, if you attempt to run commands, such as `dbed_update` then you may observe the following error:

```
$ /opt/VRTSdbed/bin/dbed_update
No repository found for database faildb, creating new one.
SFDB vxsfadm ERROR V-81-0450 A remote or privileged command could not
be executed on swpa04

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

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

Workaround: There is no workaround for this issue.

In the cloned database, the seed PDB remains in the mounted state (3599920)

In Oracle database version 12.1.0.2, when a container database (CDB) is cloned, the **PDB$SEED** pluggable database (PDB) remains in the mounted state. This behavior is observed because of the missing datafiles in the cloned database for all point-in-time copies.

When you attempt to open the cloned seed database, the following error is reported:

```
"ORA-01173" oracle error.
...
SFDB vxsfadm ERROR V-81-0564 Oracle returned error.
```

```
Reason: ORA-01122: database file 15 failed verification check
ORA-01110: data file 15: '/tmp/test1/data/sfaedb/newtbs1.dbf'
ORA-01202: wrong incarnation of this file - wrong creation time
...
```

Workaround: There is no workaround for this issue.

Cloning of a container database may fail after a reverse resync commit operation is performed (3509778)

After a reverse resync operation is performed, the cloning of a container database may fail with the following error message:

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

Reason: ORA-01503: CREATE CONTROLFILE failed
ORA-01189: file is from a different RESETLOGS than previous files
ORA-01110: data file 6: '/tmp/testRecoverdb/data/sfaedb/users01.dbf'
```

Workaround: There is no workaround for this issue.

If one of the PDBs is in the read-write restricted state, then cloning of a CDB fails (3516634)

Cloning a container database (CDB) for point-in-time copies fails if some of the pluggable databases (PDBs) are open in the restricted mode. The failure occurs with the following error message:

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

Reason: ORA-65106: Pluggable database #3 (PDB1) is in an invalid state.
```

Workaround: There is no workaround for this issue.

Cloning of a CDB fails for point-in-time copies when one of the PDBs is in the read-only mode (3513432)

For Oracle version 12.1.0.1, cloning a container database (CDB) fails if one of the pluggable databases (PDBs) is in the read-only mode. The failure occurs with the following error message:

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

Reason: ORA-00376: file 9 cannot be read at this time
ORA-01111: name for data file 9 is unknown - rename to correct file
ORA-01110: data file 9: '/ora_base/db_home/dbs/MISSING00009'...
```

Workaround: There is no workaround for this issue.

If a CDB has a tablespace in the read-only mode, then the cloning fails (3512370)

For Oracle version 12.1.0.1, when a container database (CDB) has a tablespace in the read-only mode for all point-in-time copies, cloning of that CDB fails with the following error message:

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

Reason: ORA-01122: database file 15 failed verification check
ORA-01110: data file 15: '/tmp/test1/data/sfaedb/newtbs1.dbf'
ORA-01202: wrong incarnation of this file - wrong creation time
...
```

Workaround: There is no workaround for this issue.

If any SFDB installation with authentication setup is upgraded to 6.2, the commands fail with an error (3644030)

The commands fail with the error message similar to the following:

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

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

Action: Verify that the prodhost is reachable. If it is, verify
that
the vxdbd daemon is enabled and running using the [
/opt/VRTS/bin/sfae_config
status ] command, and enable/start vxdbd using the [
/opt/VRTS/bin/sfae_config
enable ] command if it is not enabled/running. Also make sure you are
authorized to run SFAE commands if running in secure mode.
```

Workaround: Set up the authentication for SFDB again. See *Storage and Availability Management for Oracle Databases* or *Storage and Availability Management for DB2 Databases*.

Virtualization known issues

This section describes the virtualization known issues in this release of Symantec Storage Foundation Cluster File System High Availability (SFCFSHA).

Agent kill on source during migration may lead to resource concurrency violation (3042499)

In the case of a migration initiated outside Symantec Cluster Server (VCS) control, there is a very small window in which the agent restart might not be able to recognize the migration event. As this is initiated outside VCS, there is no way to synchronize the agent restart and the migration., Also, there is no intermediate state in KVM that can indicate that the event was a migration. This problem does not occur in Red Hat Enterprise Virtualization (RHEV), as there are clear states visible that can specify the virtual machine events. This is applicable to KVM environment only.

Workaround: There is no workaround for this issue.

Virtual devices backed by ALUA DMPNODE are not discovered by Veritas Volume Manager running inside KVM guests (3341432)

Inside the KVM guest, during device discovery, Veritas Volume Manager performs the `RTPG IOCTL` process on disks to fetch the disk properties. The `RTPG IOCTL` process fails on virtual devices backed by ALUA DMPNODE, which are exported from the host using the VirtIO-scsi interface. Therefore, the ASL fails to claim the disks inside guest and the disks are not visible to Volume Manager.

For example sdb is the DMPNODE backed disk, where DPMNODE belongs to the ALUA enclosrue and is exported from the host to the guest using the VirtIO-scsi interface. As the corresponding ALUA vendor ASL fails to claim disk, the DDL-STATUS displays "ERROR".

```
[root@guest1 ~]# vxddladm list devices
DEVICE               TARGET-ID    STATE   DDL-STATUS (ASL)
============================================================
sdb                  -            Online  ERROR (libvxxiv.so)
sda                  -            Online  CLAIMED (OTHER_DISKS)
```

Therefore the disk is not visible to volume manager.

```
[root@guest1 ~]# vxdisk list
DEVICE       TYPE            DISK         GROUP        STATUS
sda          auto:none       -            -            online invalid
```

Workaround: Export the underlying subpaths of DMPNODE to the guest using the VirtIO-scsi interface.

Subpaths may be removed from DMP database after I/O error occurs and become invisible inside the KVM guest (3214523)

After an I/O error occurs due to a path failure, devices may become invisible to DMP inside the KVM guest. This issue is caused by the current OS design.

The guest syslog will display the following message for the missing device:

```
detected capacity change from 107374182400 to 0
```

Workaround: When a device is missing from the `vxdmpadm getsubpaths all` output, recover the device.

To recover the missing device

1. Make sure the underlying device is accessible from the KVM host.
2. Inside the guest, re-read the partition table:

   ```
   # blockdev --rereadpt /dev/device_name
   ```

3. Re-scan the devices in the OS device tree:

   ```
   # vxdisk scandisks
   ```

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 128.

Symantec Storage Foundation Cluster File System High Availability software limitations

The following are software limitations in this release of Symantec Storage Foundation Cluster File System High Availability.

cfsmntadm command does not verify the mount options (2078634)

You must confirm the mount options are correct which are then passed to the `cfsmntadm` command. If the mount options are not correct, the mount fails and the CFSMount resource will not come online. You can check the VCS engine log file for any mount failure messages.

Obtaining information about mounted file system states (1764098)

For accurate information about the state of mounted file systems on Linux, refer to the contents of `/proc/mounts`. The `mount` command may or may not reference this source of information depending on whether the regular `/etc/mtab` file has been replaced with a symbolic link to `/proc/mounts`. This change is made at the discretion of the system administrator and the benefits are discussed in the mount online manual page. A benefit of using `/proc/mounts` is that changes to SFCFSHA mount options are accurately displayed for all nodes.

Stale SCSI-3 PR keys remain on disk after stopping the cluster and deporting the disk group

When all nodes present in the SFCFSHA cluster are removed from the cluster, the SCSI-3 Persistent Reservation (PR) keys on the data disks may not get preempted. As a result, the keys may be seen on the disks after stopping the cluster or after the nodes have booted up. The residual keys do not impact data disk fencing as they will be reused or replaced when the nodes rejoin the cluster. Alternatively, the keys can be cleared manually by running the `vxfenclearpre` utility.

For more information on the `vxfenclearpre` utility, see the *Symantec Storage Foundation Cluster File System High Availability Administrator's Guide*.

Unsupported FSS scenarios

The following scenario is not supported with Flexible Storage Sharing (FSS):

Symantec NetBackup backup with FSS disk groups

Veritas File System software limitations

The following are software limitations in this release of Veritas File System.

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, 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 feature of NetBackup 7.5 (or earlier) does not support disk layout Version 8, 9, or 10

The FlashBackup feature of NetBackup 7.5 (or earlier) does not support disk layout Version 8, 9, or 10.

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.

On SUSE, creation of a SmartIO cache of VxFS type hangs on Fusion-io device (3200586)

On SUSE, creating a SmartIO cache of VxFS type hangs on Fusion-io devices. This issue is due to a limitation in the Fusion-io driver.

Workaround:

To workaround the issue

- Limit the maximum I/O size:

```
# vxtune vol_maxio 1024
```

A NetBackup restore operation on VxFS file systems does not work with SmartIO writeback caching

A NetBackup restore operation on VxFS file systems does not work with SmartIO writeback caching.

VxFS file system writeback operation is not supported with volume level replication or array level replication

The VxFS file system writeback operation is not supported with volume level replication or array level replication.

Veritas Volume Manager software limitations

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

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

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

Cloned disks operations not supported for FSS disk groups

In this release, the VxVM cloned disks operations are is not supported with FSS disk groups. If you clone a disk in the FSS disk groups, the cloned device cannot be imported. If you prefer to use hardware mirroring for disaster recovery purposes, you need to make sure that such devices should not be used to create FSS disk groups.

For more information, see the *Administrator's Guide*.

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

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

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

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

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

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

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

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

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 1-16 describes the DMP tunable parameters and the new values.

Table 1-16 DMP settings for NetApp storage attached environment

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

The change is persistent across reboots.

To change the tunable parameters

1 Issue the following commands:

```
# vxdmpadm settune dmp_restore_interval=60

# vxdmpadm settune dmp_path_age=120
```

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

```
# vxdmpadm gettune dmp_restore_interval

# vxdmpadm gettune dmp_path_age
```

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

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

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

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

```
Disk xiv0_612 : Done.
Disk xiv0_613 : Done.
Disk xiv0_614 : Done.
Disk xiv0_615 : Done
```

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

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

VRTSvxvm upgrade using native command fails (3384435)

If you upgrade the `VRTSvxvm` rpm using the native command, `rpm -Uvh VRTSvxvm`, the command fails.

Workaround:

1. Remove the older `VRTSlvmconv` rpm.

   ```
   # rpm -e VRTSlvmconv
   ```

2. Upgrade the `VRTSvxvm` rpm.

   ```
   # rpm -Uvh VRTSvxvm
   ```

Replication software limitations

The following are replication software limitations in this release of Symantec Storage Foundation Cluster File System High Availability.

VVR Replication in a shared environment

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

VVR IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

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

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 6.1 and the prior major releases of Storage Foundation (6.0 and 6.0.1). Replication between versions is supported for disk group versions 170, 180, and 190 only. Both the Primary and Secondary hosts must be using a supported disk group version.

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

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

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

The preferred fencing feature gives preference to more weighted or larger subclusters by delaying the smaller subcluster. This smaller subcluster delay is effective only if the initial RACER node in the larger subcluster is able to complete the race. If due to some reason the initial RACER node is not able to complete the race and the VxFEN driver activates the racer re-election algorithm, then the smaller

subcluster delay is offset by the time taken for the racer re-election and the less weighted or smaller subcluster could win the race. This limitation though not desirable can be tolerated.

Stopping systems in clusters with I/O fencing configured

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

Symantec 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

SFDB tools cannot be used with RAC standby databases. SFDB tools can still be used with the primary database, even in a Data Guard Oracle RAC environment.

Upgrading to Oracle 10.2.0.5 is required if using SFDB tools

If you are running Oracle version 10.2.0.4 and upgrading a Storage Foundation product with SFDB tools to 6.2, you must upgrade the Oracle binaries and database to version 10.2.0.5, before upgrading to 6.2.

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.

Virtualization software limitations

This section describes the virtualization software limitations in this release of Symantec Storage Foundation Cluster File System High Availability (SFCFSHA).

Paths cannot be enabled inside a KVM guest if the devices have been previously removed and re-attached from the host

LUNs are exported to the KVM guest via virtio-scsi interface. When some physical link between the host and the SAN array fails for a certain time (45-60 seconds by default), the HBA driver in the host will remove the timed-out devices. When the link is restored, these devices will be re-attached to the host; however, the access from inside the KVM guest to these devices cannot be automatically restored too without rebooting the system or manually re-attaching the devices. For DMP, these subpaths will remain in DISABLED state.

This is a known limitation of KVM.

Workaround:

From the KVM host, tune the `dev_loss_tmo` parameter of the Fibre Channel ports to a very large value, and set the `fast_io_fail_tmo` parameter to 15.

To restore access to the timed-out devices

1 Add the following lines into `/dev/udev/rules.d/40-kvm-device` file:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add", \
  RUN+="/bin/sh -c 'grep -q off \
  /sys/class/fc_remote_ports/%k/fast_io_fail_tmo;if [ $? -eq 0 ]; \
  then echo 15 > /sys/class/fc_remote_ports/%k/fast_io_fail_tmo 2> \
  /dev/null;fi;'"
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add", \
  RUN+="/bin/sh -c 'echo 8000000 > \
  /sys/class/fc_remote_ports/%k/dev_loss_tmo 2> /dev/null'"
```

2 Create the `/etc/modprobe.d/qla2xxx.conf` file with the following content:

```
options qla2xxx qlport_down_retry=8000000
```

3 Create the `/etc/modprobe.d/scsi_transport_fc.conf` with the following content:

```
options scsi_transport_fc dev_loss_tmo=8000000
```

4 Rebuild the `initrd` file and reboot.

SmartIO software limitations

The following are the SmartIO software limitations in this release.

Cache is not online after a reboot

Generally, the SmartIO cache is automatically brought online after a reboot of the system.

If the SSD driver module is not loaded automatically after the reboot, you need to load the driver and bring the cache disk group online manually.

To bring a cache online after a reboot

1 Load the SSD driver module with the `insmod` command.

See the Linux documentation for details.

2 Perform a scan of the OS devices:

```
# vxdisk scandisks
```

3 Bring the cache online manually:

```
# vxdg import cachedg
```

Writeback caching limitations

In the case of CFS, writeback caching is supported with the cache area created on direct attached storage (DAS) and SAN via a Fibre Channel. The cache area should not be shared between cluster nodes.

Writeback caching is only supported on two-node CFS only.

The sfcache operations may display error messages in the caching log when the operation completed successfully (3611158)

The `sfcache` command calls other commands to perform the caching operations. If a command fails, additional commands may be called to complete the operation. For debugging purposes, the caching log includes all of the success messages and failure messages for the commands that are called.

If the `sfcache` command has completed successfully, you can safely ignore the error messages in the log file.

Documentation

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

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

http://sort.symantec.com/documents

Documentation set

Each product in the Storage Foundation and High Availability Solutions product line includes release notes, an installation guide, and additional documents such as

administration and agent guides. In most cases, you may also need to refer to the documentation for the product's components.

The SFHA Solutions documents describe functionality and solutions that apply across the product line. These documents are relevant whichever SFHA Solutions product you use.

Symantec Storage Foundation Cluster File System High Availability documentation

Table 1-17 lists the documentation for Symantec Storage Foundation Cluster File System High Availability.

The SFHA Solutions documents describe functionality and solutions relevant to the SFCFSHA product.

See Table 1-19 on page 131.

Table 1-17 Symantec Storage Foundation Cluster File System High Availability documentation

Document title	File name	Description
Symantec Storage Foundation Cluster File System High Availability Release Notes	sfcfs_notes_62_lin.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.
Symantec Storage Foundation Cluster File System High Availability Installation Guide	sfcfs_install_62_lin.pdf	Provides information required to install the product.
Symantec Storage Foundation Cluster File System High Availability Administrator's Guide	sfcfs_admin_62_lin.pdf	Provides information required for administering the product.

Symantec Cluster Server documentation

Table 1-18 lists the documents for Symantec Cluster Server.

Table 1-18 Symantec Cluster Server documentation

Title	File name	Description
Symantec Cluster Server Release Notes	vcs_notes_62_lin.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.

Table 1-18 Symantec Cluster Server documentation *(continued)*

Title	File name	Description
Symantec Cluster Server Installation Guide	vcs_install_62_lin.pdf	Provides information required to install the product.
Symantec Cluster Server Administrator's Guide	vcs_admin_62_lin.pdf	Provides information required for administering the product.
Symantec High Availability Solution Guide for VMware	sha_solutions_62_vmware_lin.pdf	Provides information on how to install, configure, and administer Symantec Cluster Server in a VMware virtual environment, by using the VMware vSphere Client GUI.
Symantec Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_62_lin.pdf	Provides information about bundled agents, their resources and attributes, and more related information.
Symantec Cluster Server Generic Application Agent Configuration Guide	vcs_gen_agent_62_lin.pdf	Provides notes for installing and configuring the generic Application agent.
Symantec Cluster Server Agent Developer's Guide (This document is available online only.)	vcs_agent_dev_62_unix.pdf	Provides information about the various Symantec agents and procedures for developing custom agents.
Symantec Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_62_lin.pdf	Provides notes for installing and configuring the DB2 agent.
Symantec Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_62_lin.pdf	Provides notes for installing and configuring the Oracle agent.
Symantec Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_62_lin.pdf	Provides notes for installing and configuring the Sybase agent.

Symantec Storage Foundation and High Availability Solutions products documentation

Table 1-19 lists the documentation for Symantec Storage Foundation and High Availability Solutions products.

Table 1-19 Symantec Storage Foundation and High Availability Solutions products documentation

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions—What's new in this release (This document is available online.)	sfhas_whats_new_62_unix.pdf	Provides information about the new features and enhancements in the release.
Symantec Storage Foundation and High Availability Solutions Getting Started Guide	getting_started.pdf	Provides a high-level overview of installing Symantec products using the script-based installer. The guide is useful for new users and returning users that want a quick refresher.
Symantec Storage Foundation and High Availability Solutions Solutions Guide	sfhas_solutions_62_lin.pdf	Provides information about how SFHA Solutions product components and features can be used individually and in concert to improve performance, resilience and ease of management for storage and applications.
Symantec Storage Foundation and High Availability Solutions Virtualization Guide (This document is available online.)	sfhas_virtualization_62_lin.pdf	Provides information about Symantec Storage Foundation and High Availability support for virtualization technologies. Review this entire document before you install virtualization software on systems running SFHA products.
Symantec Storage Foundation and High Availability Solutions SmartIO for Solid State Drives Solutions Guide	sfhas_smartio_solutions_62_lin.pdf	Provides information on using and administering SmartIO with SFHA solutions. Also includes troubleshooting and command reference sheet for SmartIO.
Symantec Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide (This document is available online.)	sfhas_dr_impl_62_lin.pdf	Provides information on configuring campus clusters, global clusters, and replicated data clusters (RDC) for disaster recovery failover using Storage Foundation and High Availability Solutions products.

Table 1-19 Symantec Storage Foundation and High Availability Solutions products documentation *(continued)*

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions Replication Administrator's Guide	sfhas_replication_admin_62_lin.pdf	Provides information on using Symantec Replicator Option for setting up an effective disaster recovery plan by maintaining a consistent copy of application data at one or more remote locations. Symantec Replicator Option provides the flexibility of block-based continuous replication with Symantec Volume Replicator Option (VVR) and file-based periodic replication with Symantec File Replicator Option (VFR).
Symantec Storage Foundation and High Availability Solutions Troubleshooting Guide	sfhas_tshoot_62_lin.pdf	Provides information on common issues that might be encountered when using Symantec Storage Foundation and High Availability Solutions and possible solutions for those issues.

Veritas Operations Manager (VOM) is a management tool that you can use to manage Symantec Storage Foundation and High Availability Solutions products. If you use VOM, refer to the VOM product documentation at:

https://sort.symantec.com/documents

Manual pages

The manual pages for Symantec 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 Symantec Storage Foundation manual pages:

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

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

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

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

See the `man`(1) manual page.

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

To edit the man(1) configuration file

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

   ```
   export LC_ALL=C
   ```

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

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

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

   where other man paths are specified in the configuration file.

3. Add new section numbers. Change the line:

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

   to

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

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

https://sort.symantec.com/documents