Symantec™ Storage Foundation 6.1 Release Notes - Linux

July 2014

Symantec

Symantec™ Storage Foundation Release Notes

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

Product version: 6.1

Document version: 6.1 Rev 6

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Storage Foundation Release Notes

This document includes the following topics:

About this document

This document provides important information about Symantec Storage Foundation (SF) version 6.1 for Linux. Review this entire document before you install or upgrade SF.

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

This is "Document version: 6.1 Rev 6" of the *Symantec Storage Foundation Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

https://sort.symantec.com/documents

Component product release notes

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

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

`/docs/product_name`

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

About Symantec Storage Foundation

Symantec Storage Foundation by Symantec (SF) is a storage management solution to enable robust, manageable, and scalable storage deployment. SF maximizes your storage efficiency, availability, agility, and performance across heterogeneous server and storage platforms.

Symantec Storage Foundation consists of product components and features that can be used individually and together to improve performance, resilience and ease of management for your storage and applications.

Table 1-1 describes the components of Symantec Storage Foundation.

Table 1-1 Symantec Storage Foundation components

Component	Description
Dynamic Multi-Pathing (DMP)	Manages the I/O performance and path availability of the physical devices that are configured on the system. DMP creates DMP metadevices for all of the paths to each LUN. DMP uses the DMP metadevices to manage path failover and load balancing across the paths to the physical devices. DMP metadevices provide the foundation for Veritas Volume Manager (VxVM) and Veritas File System (VxFS). DMP also supports native operating system volumes and file systems on DMP devices.
Veritas Volume Manager (VxVM)	Provides a logical layer between your operating system devices and your applications. VxVM enables you to create logical devices called volumes on the physical disks and LUNs.The applications such as file systems or databases access the volumes as if the volumes were physical devices but without the physical limitations. VxVM features enable you to configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Additional VxVM features enhance fault tolerance and fast recovery from disk failure or storage array failure.
Veritas File System (VxFS)	Provides a high-performance journaling file system. VxFS is designed for use in operating environments that deal with large amounts of data and that require high performance and continuous availability. VxFS features provide quick-recovery for applications, scalable performance, continuous availability, increased I/O throughput, and increased structural integrity.
Symantec Replicator (VR)	Enables you to maintain a consistent copy of application data at one or more remote locations. Symantec Replicator provides the flexibility of block-based continuous replication with Volume Replicator (VVR) and file-based periodic replication with File Replicator (VFR). Symantec Replicator option is a separately-licensable feature of Symantec Storage Foundation.
I/O fencing	Protects the data on shared disks when nodes in a cluster detect a change in the network cluster membership with a potential split brain condition.

A related product, Veritas Operations Manager, provides a centralized management console that you can use with Symantec Storage Foundation and High Availability products.

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/TECH211540
- 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/TECH213121

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 6.1

This section lists the changes in Symantec Storage Foundation 6.1.

Changes related to installation and upgrades

The product installer includes the following changes in 6.1.

Support for SFHA 6.1 installations from any supported operating system to any other supported operating system

You can use the Deployment Server or the web-based installer to install your 6.1 Symantec products on a target system that runs any supported UNIX or Linux platform, even if the source system and target system are running on different UNIX or Linux platforms. Prior to 6.1, releases still require the same platform, architecture, distribution, and version of the operating system.

See the *Installation Guide* for more information.

Improved patching and updating process

You can now download product maintenance releases and public hot fix releases directly from the Symantec Operations Readiness Tools (SORT) website using the installer. When you use the `installer` command with the `-version` option, the installer now lists the available GA releases, maintenance releases, and hot fix releases. If you have Internet access, you can follow the installer prompts to download available patches and hot fixes to your local system.

Downloading patches and hot fixes requires the installer to make outbound networking calls. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network attempts by running the installer using the no Internet patch center (`-noipc`) option. When using the `-noipc` option, the installer does not try to connect to SORT website. For example:

```
# ./installer -version -noipc system1 system2
```

See the *Installation Guide* for more information.

Automatic download of installer hot fixes

If you are running the 6.1 product installer, and your system has Internet access, the installer automatically imports any needed installer hot fix, and begins using it.

If your system does not have Internet access, you can still download installer hot fixes manually using the Symantec Operations Readiness Tools patch finder tool.

Automatic downloading of installer hot fixes requires the installer to make outbound networking calls. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network attempts by running the installer using the no Internet patch center (`-noipc`) option.

See the *Installation Guide* for more information.

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 operating system (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 1-2.

Table 1-2 Deployment Server functionality

Feature	Description
Manage release images	■ View available Storage Foundation releases. ■ Download maintenance and hot fix release images from the Symantec Operations Readiness Tools (SORT) website into a repository. ■ Load the downloaded release image files from FileConnect and SORT into the repository. ■ View and remove release image files stored in the repository.

Table 1-2 Deployment Server functionality *(continued)*

Feature	Description
Check versions	■ Discovers RPMs and patches installed on designated systems and informs you of the product and version installed, including installed hot fixes. ■ Identify base, maintenance, and hot fix level upgrades to your system and download maintenance and hot fix releases. ■ Query SORT for the most recent updates.

Install or upgrade systems	■ Install or upgrade a release stored in the repository on selected systems.
	■ In release 6.1 and later:
	  ■ Install hot fix level releases.
	  ■ Install SFHA from any supported UNIX or Linux operating system to any other supported UNIX or Linux operating system.
	  ■ Automatically load the script-based installer hot fixes that apply to that release.

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 simultaneously installing or upgrading base releases, maintenance patches, and hot fixes

Beginning with version 6.1, Symantec offers you a method to easily install or upgrade your systems directly to a base, maintenance, or hot fix level in one step using Install Bundles. Install Bundles is the ability for installers to merge so customers can install or upgrade directly to maintenance or hot fix levels in one execution. Install Bundles consists of executing the installer from a GA release with a pointer to a higher maintenance or hot fix release. The installer installs them both as if they were combined in the same release image. The various scripts, RPMs, and patch components are merged and multiple releases are installed together as if they are one install entity.

Note: This feature is not supported by the Deployment Server.

There are five possible methods of integration. All upgrades must be executed from the highest level script.

- Base + maintenance
- Base + hot fix
- Maintenance + hot fix
- Base + maintenance + hot fix
- Base or maintenance + multiple hot fixes

See the *Installation Guide* for more information.

Changes related to Symantec Storage Foundation (SF)

Symantec Storage Foundation includes the following changes in 6.1:

SmartIO: Support for caching on Solid-State Drives on Linux

Solid-State Drives (SSDs) are devices that do not have spinning disks. Today's solid-state technologies, such as DRAM and NAND flash, provide faster data access, are more efficient, and have a smaller footprint than traditional spinning disks. The data center uses solid-state technologies in many form factors: in-server, all flash arrays, all flash appliances, and mixed with traditional HDD arrays. Each form factor offers a different value proposition. SSDs also have many connectivity types: PCIe, FC, SATA, and SAS.

Due to the current cost per gigabyte of SSD devices, the best value of SSDs is not as high capacity storage devices. The benefit of adopting SSDs is to improve performance and reduce the cost per I/O per second (IOPS). Data efficiency and placement is critical to maximizing the returns on any data center's investment in solid state.

The SmartIO feature of Storage Foundation and High Availability Solutions (SFHA Solutions) enables data efficiency on your SSDs through I/O caching. Using SmartIO to improve efficiency, you can optimize the cost per IOPS. SmartIO does not require in-depth knowledge of the hardware technologies underneath. SmartIO uses advanced, customizable heuristics to determine what data to cache and how that data gets removed from the cache. The heuristics take advantage of SFHA Solutions' knowledge of the characteristics of the workload.

SmartIO supports read and write caching for VxFS file systems mounted on VxVM volumes, in several caching modes and configurations.

- Read caching for applications running on VxVM volumes
- Read caching for applications running on VxFS file systems
- Writeback caching on applications running on VxFS file systems
- Database caching on VxFS file systems

- Database caching on VxVM volumes

To use SmartIO, you set up a cache area on the target device. You can do this task simply with one command, while the application is online. When the application issues an I/O request, SmartIO checks to see if the I/O can be serviced from the cache. As applications access data from the underlying volumes or file systems, certain data is moved to the cache based on the internal heuristics. Subsequent I/Os are processed from the cache.

You can also customize which data is cached, by adding advisory information to assist the SmartIO feature in making those determinations.

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

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 6.1:

DMP support for thin reclamation commands

In this release, Dynamic Multi-Pathing (DMP) adds support for the `UNMAP` command for thin reclamation. The Array Support Library (ASL) for each array uses the most suitable reclamation method supported for the array. In previous releases, DMP performed reclamation with the WRITE_SAME method for SCSI and the TRIM method for SSD devices. You can use the `vxdisk -p list` command to show the reclaim interface that is supported for a particular device.

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

Enhancements to the disk cloning operations

In this release, the following enhancements are made to the VxVM support for hardware clone disks:

- When you import a disk group, the disks with the `udid_mismatch` flag display the `clone_disk` flag regardless of whether the system sees the original source disks. In previous releases, the `clone_disk` flag was hidden if the source disks were not visible to the system.
- By default, VxVM now prevents the import of a partial set of disks in a clone disk group when the `-o updateid` option is specified. This behavior prevents the missing disks from being permanently detached from the new disk group. You can specify the `-f` option to partially import the clone disk group with `-o updateid`.
- When you import a set of clone disks with the `-o updateid` option and specify a new disk group name, the disk group becomes a standard disk group with

updated disk and disk group identifiers. This operation clears the `udid_mismatch` flag or the `clone_disk` flag from the disks.

- When you import a set of clone disks with the `-o updateid` option, you can use the `vxdg import` with the `-c` option to convert the existing disk group to a standard disk group with updated disk and disk group identifiers. This operation clears the `udid_mismatch` flag or the `clone_disk` flag from the disks. You cannot perform this operation if the source disk group is already imported on the same host.
- You can now update the UDID for a disk and remove the `udid_mismatch` flag and the `clone_disk` flag with a single operation. Updating the UDID aligns it with the UDID detected by the DDL.
  `vxdisk -c updateudid diskname`
- You cannot create disk groups on `udid_mismatch` or `clone_disk` disks.
- If disks are falsely marked as `udid_mismatch`, you can use `vxdg -c init` option to create disk groups on them.
- If the disk group has multiple clone copies, and you import the disk group with a tagname, the disks that have tags set will be selected. The tag-based import operation gives higher priority to disks with the tags set rather than the set of disks that were last imported. In previous releases, if multiple clone copies had the same disk group id, the import operation gave preference to the last import time.

Enhancements to the Dynamic Reconfiguration tool

This section describes enhancements to the Dynamic Reconfiguration tool in this release. The Dynamic Reconfiguration tool now:

- Does not display internal devices or LVM devices for removal. The Dynamic Reconfiguration tool now removes the stale entries properly.
- Prompts you to rename devices during a Dynamic Reconfiguration operation, if appropriate, and if avid=no in the naming scheme. If you agree, the tool renames the devices and refreshes the device list.
  For example, if you have removed the LUN named xyz_8, which leaves the entries xyz_7 and xyz_9. The DR tool prompts you whether you want to rename the LUNs. If you agree, xyz_9 is renamed to xyz_8.
- Logs messages for each use of the tool, in the format `dmpdr_yyyymmdd_HHMM.log`.
- Accepts a file containing a list of devices as input to the removal operation.
- Displays all LUNs that are not operating as candidates for removal.

- Supports pattern matching to select disks for removal. For example, you can use an asterisk (*) to match multiple characters and a question mark (?) to match a single character. This functionality replaces the option to specify a range of devices.
- If you quit a disk removal operation without physically removing the disks, the Dynamic Reconfiguration tool prompts you to run `vxdisksetup` over the selected disks to avoid data corruption.

Changes related to Veritas File System

Veritas File System (VxFS) includes the following changes in 6.1:

Support for 64-bit quotas

Starting in release 6.1, 64-bit quotas are supported on disk layout Version 10. Users were earlier limited to set quota usage limits only up to 1 terabyte, restricting functionality in high data usage environments. With the support for 64-bit quotas, the quota usage limit can be set up to 4 exabytes.

As for 32-bit quotas, this continues to be supported on disk layout Version 9 or earlier. The same quota commands can be used for both 32-bit and 64-bit quotas.

As for 64-bit quotas, there are two new quotas files. For group quotas the file name is `quotas.grp.64` and for user quotas the file name is `quotas.64`. These files will be created on each file system after the disk layout version upgrade is completed.

See the *Administrator's Guide* for more information on quota files on Veritas File System.

See the *Installation Guide* for more information on upgrading disk layout versions.

maxlink support

Added support for more than 64K sub-directories. If maxlink is disabled on a file system, the sub-directory limit will be 32K by default. If maxlink is enabled on a file system, this allows you to create up to 4294967295(2^32 – 1) sub-directories.

By default maxlink is enabled.

See the *Administrator's Guide*.

Disk layout Version 10

In this release, disk layout Version 10 is now the default version.

Version 10 disk layout enables support for SmartIO and maxlink.

See the *Administrator's Guide*.

vxfsstat command can display per file system memory and VOP statistics

The `vxfsstat` command can now display per file system memory and VOP statistics. The following options display the statistics:

`-B` Displays per file system metadata buffer cache statistics.

`-I` Displays per file system inode cache and DNLC statistics.

`-x` An already existing option that displays per file system statistics, and now additionally displays the newly added memory and VOP counters. VOP counters include VOP time and VOP count.

Changes related to SFDB tools

The following sections describe the changes related to Storage Foundation for Databases (SFDB) tools in 6.1.

Reverse Resync for Oracle database recovery

In this release, the SFDB tools reintroduce the Reverse Resync feature for Oracle database recovery.

Reverse Resynchronization or Reverse Resync process helps in recovering a database from its volume snapshots using FlashSnap service.

Storage Foundation Database FlashSnap service is used to reverse resynchronize an online point-in-time copies image of a database in an Oracle environment.

Reverse Resync feature was supported in 5.X release. This feature was discontinued for 6.0 and 6.0.1 releases. In the current release, Reverse Resync feature is reintroduced with the following changes:

- You can perform ReverseResyncBegin operation after ReverseResyncAbort operation
- You can control the database recovery in ReverseResyncBegin operation using the new (optional) parameters:

  ```
  Reverse_Resync_Recovery

  Reverse_Resync_Archive_Log
  ```

Use the following commands for reverse resynchronization of the snapshot volume:

- `vxsfadm -o rrbegin` to start the Reverse Resync operation
- `vxsfadm -o rrcommit` to commit the Reverse Resync changes

- `vxsfadm -o rrabort` to abort or cancel the Reverse Resync operation and to go back to the original data volumes

Note: Reverse resync is not supported for RAC databases.

Supported DB2 configurations

In this release, SFDB tools are supported with DB2 10.1 release.

Supported Oracle configurations

In 6.1 release, SFDB tools support Oracle 12c release for Oracle databases.

Note: For Oracle 12c, the SFDB tools do not support the Multitenant database features, including the CDB and PDB databases.

Support for instant mode snapshots for Oracle RAC databases

In 6.1, the SFDB tools support instant mode snapshots for Oracle RAC databases.

Changes related to replication

Symantec Storage Foundation and High Availability Solutions includes the following changes related to replication in 6.1:

New vfradmin job promote and vfradmin job recover commands simplify changing replication direction

The `vfradmin` command now has the `job promote` keyword and `job recover` keyword that enable you to change the replication direction with a single command instantiation. You use `job recover` after a disaster occurred, and `job promote` under normal circumstances.

See the `vfradmin`(1M) manual page.

VVR replication performance improvements using bulk transfer

To effectively use network bandwidth for replication, data is replicated to a disaster recovery (DR) site in bulk at 256 KB. This bulk data transfer reduces Volume Replicator (VVR) CPU overhead and increases the overall replication throughput. With compression enabled, bulk data transfer improves the compression ratio and reduces the primary side CPU usage. Bulk data transfer is not supported with bunker replication, and in cross-platform replication.

VVR I/O throughput improvements using batched writes

Batched writing of multiple application writes to the SRL increases application I/O throughput and lowers VVR CPU utilization. This is achieved by allocating a log location for a set of application writes, and then batching the writes together to form a single write to the SRL, and therefore replacing the multiple writes to the SRL at the primary RVG.

Changes related to product name branding

Beginning with the 6.1 release, Storage Foundation and High Availability Solutions product names are rebranded.

Table 1-3 lists the rebranded Storage Foundation and High Availability Solutions products.

Table 1-3 Rebranded Storage Foundation and High Availability Solutions products

Old product name	New product names with Symantec branding
Veritas Storage Foundation	Symantec Storage Foundation (SF)
Veritas Dynamic Multi-Pathing	Symantec Dynamic Multi-Pathing (DMP)
Veritas Replicator Option	Symantec Replicator Option
Veritas File Replicator Option	Symantec File Replicator Option (VFR)
Veritas Volume Replicator	Symantec Volume Replicator (VVR)
Veritas Storage Foundation Cluster File System HA	Symantec Storage Foundation Cluster File System HA (SFCFSHA)
Veritas Storage Foundation for Oracle RAC	Symantec Storage Foundation for Oracle RAC (SFRAC)
Veritas Storage Foundation for Sybase ASE CE	Symantec Storage Foundation for Sybase ASE CE
Veritas Storage Foundation HA	Symantec Storage Foundation HA (SFHA)
Veritas Cluster Server	Symantec Cluster Server (VCS)
Veritas Disaster Recovery Advisor	Symantec Disaster Recovery Advisor (DRA)
Veritas Storage Foundation and High Availability Solutions	Symantec Storage Foundation and High Availability Solutions (SFHAS)

Table 1-3 Rebranded Storage Foundation and High Availability Solutions products *(continued)*

Old product name	New product names with Symantec branding
Veritas High Availability Agent Pack	Symantec High Availability Agent Pack
Veritas File System Software Development Kit	Symantec File System Software Development Kit

Symantec rebranding does not apply to the following:

- Product acronyms
- Command names
- Error messages
- Alert messages
- Modules and components
- Feature names
- License key description
- Veritas Operations Manager product branding

No longer supported

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

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

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

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

- Storage Checkpoint policy and Storage Checkpoint quotas
- Interactive modes in clone and rollback

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems

This section lists the supported operating systems for this release of 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	Levels	Kernel version
Red Hat Enterprise Linux 6	Update 3 Update 4 Update 5 **Note:** Update 5 is supported if you install the required VxFS, LLT, and ODM patches. See the section called "Support for RHEL 6.5 and OL 6.5" on page 25.	2.6.32-279.el6 2.6.32-358.el6 2.6.32-431.el6
Red Hat Enterprise Linux 5	Update 5 Update 6 Update 7 Update 8 Update 9 Update 10	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5 2.6.18-308.el5 2.6.18-348.el5 2.6.18-371.el5
SUSE Linux Enterprise 11	SP2 SP3	3.0.13-0.27.1 3.0.76-0.11.1

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

Operating systems	Levels	Kernel version
Oracle Linux 6	Update 3	2.6.32-279.el6
	Update 4	2.6.32-358.el6
	Update 5 **Note:** Update 5 is supported if you install the required VxFS, LLT, and ODM patches. See the section called "Support for RHEL 6.5 and OL 6.5" on page 25.	2.6.32-431.el6
Oracle Linux 5	Update 5	2.6.18-194.el5
	Update 6	2.6.18-238.el5
	Update 7	2.6.18-274.el5
	Update 8	2.6.18-308.el5
	Update 9	2.6.18-348.el5
	Update 10	2.6.18-371.el5

Note: Oracle Linux is supported with Red Hat Enterprise Linux compatible kernel only. Oracle Linux Unbreakable Enterprise Kernel is not supported.

Note: All subsequent kernel versions and patch releases on the supported operating system levels 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: SmartIO is not supported with SLES11 SP3 for Fusion-io SSD cards as the driver support for these SSD cards is not available.

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.

Support for RHEL 6.5 and OL 6.5

Symantec Storage Foundation and High Availability Solutions (SFHA) 6.1 by default does not support RHEL 6.5 and Oracle Linux (OL) 6.5 due to incompatibility in the kernel interface in the Veritas File System (VxFS), Low Latency Transport (LLT), and Oracle Disk Manager (ODM) components of SFHA.

To support RHEL 6.5 and OL 6.5, you must install SFHA 6.1, and then install the required VxFS, LLT, and ODM patches as described in Table 1-5 to resolve the kernel incompatibility.

Table 1-5 Required VxFS, LLT, and ODM patches

RPM name	Minimum patch level	Platform
VRTSvxfs	6.1.0.200	RHEL 6
VRTSllt	6.1.0.100	RHEL 6
VRTSodm	6.1.0.100	RHEL 6

You can obtain the required VxFS, LLT, and ODM patches from the Symantec Operations Readiness Tools (SORT) Patch Finder page at:

https://sort.symantec.com/patch/finder

Required Linux RPMs for SF

Make sure you install the following operating system-specific RPMs on the systems where you want to install or upgrade SF. SF 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 1-6 lists the RPMs that SF requires for a given Linux operating system.

Table 1-6 Required RPMs

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 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 1-6 Required RPMs *(continued)*

Operating system	Required RPMs
RHEL 5	coreutils-5.97-23.el5_4.2.x86_64.rpm
	ed-0.2-39.el5_2.x86_64.rpm
	findutils-4.2.27-6.el5.x86_64.rpm
	glibc-2.5-58.i686.rpm
	glibc-2.5-58.x86_64.rpm
	libacl-2.2.39-6.el5.i386.rpm
	libacl-2.2.39-6.el5.x86_64.rpm
	libgcc-4.1.2-50.el5.i386.rpm
	libgcc-4.1.2-50.el5.x86_64.rpm
	libstdc++-4.1.2-50.el5.i386.rpm
	libstdc++-4.1.2-50.el5.x86_64.rpm
	module-init-tools-3.3-0.pre3.1.60.el5_5.1.x86_64.rpm
	ncurses-5.5-24.20060715.x86_64.rpm
	openssl-0.9.8e-12.el5_5.7.x86_64.rpm
	pam-0.99.6.2-6.el5_5.2.i386.rpm
	parted-1.8.1-27.el5.i386.rpm
	parted-1.8.1-27.el5.x86_64.rpm
	policycoreutils-1.33.12-14.8.el5.x86_64.rpm
	readline-5.1-3.el5.x86_64.rpm
	zlib-1.2.3-3.x86_64.rpm

Table 1-6 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
	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.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 1-6 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
	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-1.2.3-106.34.x86_64.rpm

Table 1-6 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
	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

Storage Foundation for Databases features supported in database environments

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

Table 1-7 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

Table 1-7 SFDB features supported in database environments *(continued)*

Symantec Storage Foundation feature	DB2	Oracle	Oracle RAC	Sybase
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.

Symantec Storage Foundation memory requirements

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

Fixed issues

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

Installation and upgrades fixed issues

There are no install and upgrade fixed issues in this release.

Symantec Storage Foundation fixed issues

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

See "Veritas File System fixed issues" on page 32.

See "Veritas Volume Manager fixed issues" on page 37.

Veritas File System fixed issues

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

Table 1-8 Veritas File System fixed issues

Incident	Description
3331134	File system hangs due to a race condition when inodes are re-used from the delicache list.
3331125	Enhancement to handle partial compressed extents during dedupe operation.
3331120	Aborting a replication job leaves behind a checkpoint.
3331109	Additional checks in `fsck` to prevent file system metadata corruption with filesnap.
3331105	The `fsck` command does not validate if multiple reorg inodes point to the same source inode.
3331095	The `fsppadm` utility dumps core when an incorrect policy is specified during enforcement.
3331071	The `fsppadm` query and enforcement should honor the `-P` option to exclude private files.
3331066	Newly created consistency group for replication is not preserved after a checkpoint promotion on target.
3331047	Memory leak in `vx_followlink`.
3331045	The system panics in `vx_unlockmap` due to null pointer dereference.

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

Incident	Description
3331010	File system full `fsck` fails as it erroneously accesses freed memory during RCT processing.
3330985	Support for changing the direction of replication in VFR is added.
3310755	`fsck` fix to handle ZFOD extents while processsing the `VX_RCQ_OP_DEC_ALL` operation.
3308673	A fragmented FS may get disabled when delayed allocations are enabled.
3298041	With the delayed allocation feature enabled on a locally mounted file system, observable performance degradation might be experienced when writing to a file and extending the file size.
3291635	The file system hangs when RCQ is full.
3271892	VFR jobs are not scalable, and they may fail if the same PID is associated with the multiple jobs working on different target file systems.
3261462	Mapbad corruption due to buffer overrun of VX_TYPED_4 to VX_TYPED_DEV8 conversion.
3253210	The file system hangs when it has reached the space limit.
3252983	During the test after having ported from 2486597, you see a dead loop situation where CPU is taken 100%, and the system barely responds.
3249958	When /usr is mounted as a separate file system, the VxFS fails to load.
3233284	The `fsck` (1M) command hangs while checking Reference Count Table (RCT).
3228955	Some `fsck` enhancements to check that invalid extops are not present in older file system layouts.
3224101	After the optimization is enabled for updating the `i_size` across the cluster nodes lazily, the system panics.

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

Incident	Description
3214816	With the DELICACHE feature enabled, frequent creation and deletion of the inodes of a user may result in corruption of the user quota file.
3194635	File system metadata corruption involving ZFOD extents and filesnap or compression.
3189562	Oracle daemons get hang with the `vx_growfile()` kernel function.
3164418	Data corruption happens due to the ZFOD split during ENOSPC conditions.
3153919	The `fsadm` shrink may hang, waiting for the hlock ownership while structural file set reorg is in progress.
3152313	With the Partitioned Directories feature enabled, removing a file may panic the system.
3150368	The `vx_writesuper()` function causes the system to panic in `evfsevol_strategy()`.
3142045	With Oracle 12c version, Veritas ODM library gives a version mismatch error.
3140990	Requirement for the ability to turn off VxFS's invalidation of pages for some Network File System (NFS) workloads.
3137886	Thin Provisioning Logging does not work for reclaim operations triggered via `fsadm`.
3101418	The current time returned by the operating system (Oracle error code ORA-01513) during Oracle startup is invalid.
3096834	Intermittent `vx_disable` messages display in the system log.
3089211	When you add or remove CPUs, Veritas File System (VxFS) may crash with the Data Storage Interrupt (DSI) stack trace.
3089210	The message `V-2-17: vx_iread_1 <filesystem> file system inode <inode number> marked bad incore` is displayed in the system log.
3069695	Default Access Control Lists (ACLs) are handled on named attributes.

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

Incident	Description
3068902	In case of stale NFS mounts, the `statfs()` function calls on non-VxFS file systems may cause df commands to hang.
3066116	The system panics due to the NULL pointer dereference at the `vx_worklist_process()` function.
3042485	The fix to address file system metadata corruption involves named attribute directories.
3040944	The file system hangs due to a deadlock between the `dalloc` flusher thread and `dalloc` freeze under ENOSPC conditions.
3029093	The `fsck` command fails to repair a file system with inconsistencies in RCT/RCQ records.
3011959	The system may panic because of the file system locking or unlocking using the `fsadm`(1M) or the `vxumount`(1M) command.
3003679	When you run the `fsppadm`(1M) command and remove a file with the named stream attributes `(nattr)` at the same time, the file system does not respond.
2999493	The file system check validation fails after a successful full `fsck` during the internal testing with the following message: `run_fsck : First full fsck pass failed, exiting`.
2983248	The `vxrepquota`(1M) command dumps core.
2977697	A core dump is generated while the clone is being removed.
2966277	The high file system activities like read, write, open and lookup may panic the system.
2963763	When the `thin_friendly_alloc()` and `deliache_enable()` functionality is enabled, VxFS may enter a deadlock.
2926684	In rare cases, the system may panic while performing a logged write.
2924447	Full `fsck` performance needs to be improved to reduce the amount of disk I/O.

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

Incident	Description
2923105	Removal of the VxFS module from the kernel takes a longer time.
2916691	The `fsdedup` command hangs with an infinite loop in `vx_dedup_extents`.
2908391	It takes a long time to remove checkpoints from the VxFS file system, when there are a large number of files present.
2906018	The `vx_iread` errors are displayed after successful log replay and mount of the file system.
2905820	If the file is being read via the NFSv4 client, then removing the same file on the NFSv4 server may hang if the file system is VxFS.
2893551	The file attribute value is replaced with question mark symbols when the Network File System (NFS) connections experience a high load.
2885592	The `vxdump` operation is aborted on file systems which are compressed using the `vxcompress` command.
2881211	File ACLs are not preserved in checkpoints properly if the file has a hardlink.
2878164	VxFS consumes too much pinned heap.
2864471	The file system hangs during clone removal with Partition directory turned on.
2858683	For files greater than 8192 bytes, the reserve-extent attribute is changed after you run the command `vxrestore`.
2841059	The file system is marked for a full `fsck` operation and the attribute inode is marked as `bad ondisk`.
2839871	On a system with DELICACHE enabled, several file system operations may hang.
2825125	VxFS does not support for sub-directories larger than 64K.
2781552	Mount detects the file system not being clean and hence sets the `fullfsck` flag. `fsck` is not able to clean the system.
2773383	The read and write operation on a memory mapped files hangs.

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

Incident	Description
2750860	Performance of the write operation with small request size may degrade on a large file system.
2720034	The `vxfsckd` daemon does not restart after being manually killed.
2706900	`llseek()` implementation of VxFS for SEEK_DATA and SEEK_HOLE on Linux needs to be enhanced.
2667658	The `fscdsconv endian` conversion operation fails because of a macro overflow.
2641438	After a system is restarted, the modifications that are performed on the username space-extended attributes are lost.
2624262	System panics while executing dedup operation.
2444146	The Oracle Disk Manager read returns EINTR while running unspecified Oracle jobs.
2417858	VxFS quotas do not support 64 bit limits.
2414266	The `fallocate(2)` system call fails on VxFS file systems in the Linux environment.
2244932	Possible assertion failure in `the vx_freeze_block_threads_all()` call when the pdir_threshold tunable is set to 1.

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 1-9 Veritas Volume Manager fixed issues

Incident	Description
3325371	When using snapshots, there is a panic in `vol_multistepsio_read_source`.
3312162	VVR:DV: Verification of the remote volumes found differences with `vradmin verifydata`.

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

Incident	Description
3283525	Data Change Object (DCO) corruption after volume resize leads to `vxconfigd` hang.
3271595	VxVM should not allow turning off disk reclaim flag when there are pending reclaims on the disk.
3261601	`dmp_destroy_dmpnode` trying to free an already freed address.
3254311	The system panics when reattaching the site to a site-consistent disk group having a volume larger than 1 TB.
3249264	Disks get into the `ERROR` state after being destroyed with the command `vxdg destroy` *`dg-name`*.
3240858	The `/etc/vx/vxesd/.udev_lock` file may have different permissions at different instances.
3237503	System hang may happen after creating space-optimized snapshot with large size cache volume.
3236773	Multiple error messages of format `vxdmp V-5-3-0 dmp_indirect_ioctl: Ioctl Failed` can be seen during set/get failover-mode for EMC ALUA disk array.
3235350	System panic by `vxiod` process.
3220929	The `vxvmconvert`(1M) tool fails to convert the Logical Volume Manager (LVM) volume to Veritas Volume Manager (VxVM) volume.
3218013	Dynamic Reconfiguration (DR) Tool does not delete stale OS (operating system) device handles.
3199398	Output of the command `vxdmpadm pgrrereg` depends on the order of DMP node list where the terminal output depends on the last LUN (DMP node).
3199056	Veritas Volume Replicator (VVR) primary system panics in the `vol_cmn_err` function due to the VVR corrupted queue.
3194358	Continuous I/O error messages on OS device and DMP node can be seen in the syslog associated with the EMC Symmetrix not-ready (NR) LUNs.
3188154	`vxconfigd` is down after enabling native support on and reboot.
3186971	The logical volume manager (LVM) configuration file is not correctly set after turning on DMP native support. As a result, the system is unbootable.

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

Incident	Description
3186149	On Linux system with LVM version 2.02.85, on enabling `dmp_native_support` LVM volume groups disappear.
3182350	If there are more than 8192 paths in the system, the `vxassist`(1M) command hangs when you create a new VxVM volume or increase the existing volume's size.
3182175	The `vxdisk -o thin,fssize list` command can report incorrect file system usage data.
3178029	The value of "different blocks" is more than 100% while syncing the rvg.
3177758	Performance degradation is seen after upgrade from SF 5.1SP1RP3 to SF 6.0.1 on Linux.
3162418	The `vxconfigd`(1M) command dumps core due to an incorrect check in the `ddl_find_cdevno()` function.
3146715	Rlinks do not connect with Network Address Translation (NAT) configurations on Little Endian Architecture.
3130379	The `vxplex` command core dumped under random memory allocation failures.
3126204	[VVR] : machine panics when SRL is full.
3125631	With latest train snapshot fails for dbdst setup with error `vxsnap ERROR V-5-1-6433 Component volume has changed`.
3114134	Smart(sync) Autosync fails to work and instead replicates the entire volume size for larger sized volumes.
3111062	Make the `vxrsync` socket connection mechanism more robust.
3107741	`vxrvg snapdestroy` fails with a `Transaction aborted waiting for io drain` error and `vxconfigd` hangs for about 45 minutes.
3091978	The `lvm.conf` variable preferred_names is set to use DMP even if the dmp_native_support tunable is ‘off’.
3090667	System panics/hangs while executing `vxdisk -o thin,fssize list` as part of VOM SF discovery.
3088059	On Red Hat Enterprise Linux 6.x (RHEL6.x), the type of host bus adapter (HBA) is reported as SCSI instead of FC.

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

Incident	Description
3086627	The `VxVM vxdisk ERROR V-5-1-16282 Cannot retrieve stats: Bad address` error message displays while using vxdisk -o thin,fssize list for hitachi_usp-vm0 enclosure on the array configured for truecopy P-VOLs.
3076093	The patch upgrade script `installrp` can panic the system while doing a patch upgrade.
3063378	Some VxVM commands run slowly when EMC PowerPath presents and manages "read only" devices such as EMC SRDF-WD or BCV-NR.
3041014	Beautify error messages which are seen on the execution of relayout command.
3012929	`vxconfigbackup` keeps old disk names in its files and gives errors, when disk names are changed.
3011405	The `vxtune -o export` command failed with `V-5-1-8826 (EXDEV).`
3010191	Previously excluded paths are not excluded after upgrade to VxVM 5.1SP1RP3.
3002770	While issuing a SCSI inquiry command, NULL pointer dereference in DMP causes system panic.
2994976	BAD TRAP panic in `vxio:vol_mv_pldet_callback`.
2992667	When new disks are added to the SAN framework of the Virtual Intelligent System (VIS) appliance and the Fibre Channel (FC) switcher is changed to the direct connection, the `vxdisk list` command does not show the newly added disks even after the `vxdisk scandisks` command is executed.
2988593	Mirror resync is slower than expected.
2979824	The `vxdiskadm`(1M) utility bug results in the exclusion of the unintended paths.
2970368	Enhance handling of SRDF-R2 Write-Disabled devices in DMP.
2969844	The device discovery failure should not cause the DMP database to be destroyed completely.
2959733	Handling the device path reconfiguration in case the device paths are moved across LUNs or enclosures to prevent the `vxconfigd`(1M) daemon coredump.
2959325	The `vxconfigd`(1M) daemon dumps core while performing the disk group move operation.

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

Incident	Description
2948172	Executing the `vxdisk -o thin,fssize list` command can result in panic.
2946440	Add back the support for "INF" for LSI and ENGENIO VIDs to the LSI ASL.
2945658	If the Disk label is modified for an Active/Passive LUN, then the current passive paths do not reflect this modification after a failover.
2943637	DMP IO statistic thread may cause out of memory issue so that OOM (Out Of Memory) killer is invoked and causes system panic.
2940446	A full file system check (fsck) hangs on I/O in Veritas Volume Manager (VxVM) when the cache object size is very large.
2928764	SCSI3 PGR registrations fail when dmp_fast_recovery is disabled.
2925893	Make changes to Huawei APM to skip re-registering the keys on Secondary during failover.
2919714	On a thin Logical Unit Number (LUN), the `vxevac`(1M) command returns 0 without migrating the unmounted-VxFS volumes.
2915836	`vxnotify` does not report volume enabled message.
2911040	The restore operation from a cascaded snapshot leaves the volume in unusable state if any cascaded snapshot is in the detached state.
2899173	The `vxconfigd`(1M) daemon hangs after executing the `vradmin stoprep` command.
2898547	The `vradmind` process dumps core on the Veritas Volume Replicator (VVR) secondary site in a Clustered Volume Replicator (CVR) environment, when Logowner Service Group on VVR Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes
2898324	UMR errors reported by Purify tool in `vradmind migrate` command.
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.
2880981	Thin Reclamation of EMC Symmetrix array with Microcode 5876 could fail with error EIO.
2878876	The `vxconfigd` daemon dumps core in `vol_cbr_dolog()` due to race between two threads processing requests from the same client.

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

Incident	Description
2876706	VxVM commands hang when a LUN is changed to `not_ready` state from the array.
2869514	Issue with a configuration with large number of disks when the joining node is missing disks.
2866299	The `vxrecover` command does not automatically recover layered volumes in an RVG.
2859470	The Symmetrix Remote Data Facility R2 (SRDF-R2) with the Extensible Firmware Interface (EFI) label is not recognized by Veritas Volume Manager (VxVM) and goes in an error state
2857044	System crashes while resizing a volume with Data Change Object (DCO) version 30.
2851403	System panics while unloading `vxio` module when SmartMove feature is used and the `vxportal` module is reloaded (for example, during VxFS RPM upgrade).
2839059	`vxconfigd` logged warning `cannot open /dev/vx/rdmp/cciss/c0d device to check for ASM disk format`.
2815517	The `vxdg adddisk` command allows mixing of clone and non-clone disks in a disk group.
2779580	The `vradmin repstatus` operation may display a configuration error after cluster reconfiguration in a CVR environment.
2762147	I/O hangs on the primary node when running the `vxrvg` snaprestore operation.
2753954	At cable disconnect on port1 of dual-port FC HBA, paths via port2 marked SUSPECT.
2751423	`vxconfigd` core dumps in `ddl_migration_devlist_removed` during execution of internal testing.
2737686	The `vxddladm list [devices|hbas|ports|targets]` command shows invalid output in some platforms and in some platforms the output fields are empty.
2715124	The `vxrecover` command does not handle RAID 5 volumes correctly.
2643506	`vxconfigd` core dumps when different LUNs of same enclosure are configured with different array modes.

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

Incident	Description
2567618	The VRTSexplorer dumps core in `vxcheckhbaapi/print_target_map_entry`.
2510928	The extended attributes reported by `vxdisk -e list` for the EMC SRDF luns are reported as `tdev mirror`, instead of `tdev srdf-r1`.
2398954	The system panics while performing I/O on a VxFS mounted instant snapshot with the Oracle Disk Manager (ODM) SmartSync enabled.
2366066	The VxVM (Veritas Volume Manager) `vxstat` command displays absurd statistics for READ & WRITE operations on VxVM objects.
2152830	In a multilevel clone disks environment, a regular disk group import should be handled properly. In the case of a disk group import failure, it should report the correct error message.
2149922	Importing a disk group using clone disks fails with a "wrong usage" or "invalid attribute" error.
2000585	`vxrecover` does not start remaining volumes if one of the volumes is removed during `vxrecover` command run.
1973983	The `vxunreloc`(1M) command fails when the Data Change Object (DCO) plex is in DISABLED state.
1953257	Panic in voldiodone, because a disk with hung IO is moved out of the disk group.
1952197	Running vxtrace against a volume shows response times as negative.
1942051	I/O hangs on master node after disabling secondary paths from slave node and rebooting slave node.
1903700	Removing mirror using `vxassist` does not work.
1902483	Unique PGR key per group is not needed.
1765916	VxVM socket files do not have proper write protection.
1557628	`vxvm-startup` script does not update apm key links if a different kernel is booted in a dual boot kernel system.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues

Table 1-10 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-10 SFDB tools fixed issues

Incident	Description
2591463	Database Storage Checkpoint unmount may fail with device busy.
2534422	FlashSnap validate reports snapshot unsplittable.
2580318	dbed_vmclonedb ignores new clone SID value after cloning once.
2579929	User authentication fails.
2479901	FlashSnap resync fails if there is an existing space-optimized snapshot.
2869268	Checkpoint clone fails in a CFS environment if cloned using same checkpoint and same clone name on both nodes.
2849540	Very long off-host cloning times for large number of data files.
2715323	SFDB commands do not work with the ZHS16GBK character set.

Known issues

This section covers the known issues in this release.

Installation known issues

This section describes the known issues during installation and upgrade.

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) 6.1, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/openv`), then while upgrading to SF 6.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure RPMs `VRTSpbx`, `VRTSat`, and `VRTSicsco`. This causes NetBackup to stop working.

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the `/usr/openv/netbackup/bin/version` file and

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

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

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

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

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

Error messages in syslog (1630188)

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

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

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

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

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

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

Or

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

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

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

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

Upgrading from Storage Foundation 5.1 Service Pack 1 Rolling Patch 2 to 6.1 with rootability enabled fails (2581313)

Upgrading from Storage Foundation (SF) 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 to 6.1 while using an encapsulated root disk fails because the post installation scripts of Veritas Volume Manager (VxVM) are unable to start the `initrd` daemon.

Workaround: To upgrade from 5.1 SP1 RP2 to 6.1 while using an encapsulated root disk, you must reinstall the nash utility on the system prior to the upgrade.

To upgrade from 5.1 SP1 RP2 to 6.1 while using an encapsulated root disk

1. Encapsulate the root disk.
2. Reinstall the nash utility.
3. Upgrade to the SF 6.1 release.

During upgrade from 5.1SP1 to 61 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 SF 5.1 SP1 to SF 6.1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name for the split operation is used by an existing deported disk group.

Workaround:

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

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

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

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

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.

The uninstaller does not remove all scripts (2696033)

After removing SF, 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

The installer -license sys1 command does not show the check result if the system already has the SF Basic license key registered (3343592)

For SF Basic, you can use the **`installer -license sys1`** command to register the SF Basic key. After the command completes, however, and you try to run the same command again, the product installer does not tell you that the system has a registered license key.

Workaround: You can either run the **`vxlicrep`** command to see the registered license keys on the system, or simply ignore the second run of the **`installer -license sys1`** command.

Symantec Storage Foundation known issues

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

SmartIO cache area does not come online after I/O error to the disk (3133854)

If the paths to the disk used for the SmartIO cache area are disabled or become unavailable, then the cache area goes offline. If you attempt to bring the cache area online before the paths are back online, then the following message displays:

```
Cache area does not exist
```

Workaround:

After the paths are available and enabled, then you can bring the cache area online. The cache area does not come online automatically.

To bring the cache area online

1 Display the cache area:

```
# sfcache list
```

2 Bring the cache area online:

```
# sfcache online cachearea_name
```

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.

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 using the Linux commands 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
```

Not all the objects are visible in the VOM GUI (1791063, 1821803)

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

Workaround:

To resolve this known issue

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

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

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

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

Workaround: To see the placement class tags in the VEA GUI, you must tag the volumes prior to constructing the volume set. If you already constructed the volume set before tagging the volumes, restart `vxsvc` to make the tags visible in the GUI.

Veritas Volume Manager known issues

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

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.

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

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

```
VxVM vxencap INFO V-5-2-5327 Missing file: /boot/grub/device.map
```

Workaround:

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

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

vxrestored daemon fails to restore disabled paths (1663167)

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

Workaround:

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

To enable the mpt_disable_hotplug_remove tunable

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

```
options mptsas mpt_disable_hotplug_remove=0
```

2 Rebuild the initrd image:

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

3 Reboot the system.

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

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

Workaround:

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

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

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

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

Workaround:

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

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

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

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

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

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

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

Workaround:

To recover from this situation:

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

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

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

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

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

Root disk encapsulation issue (1603309)

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

See the vxdmpadm(1M) manual page.

See the "Setting customized names for DMP nodes" section of the *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.

System panic occurs on RHEL6.0 when VxVM is installed on EMC PowerPath managed devices (2573229)

System panic occurs on RHEL6.0 due to page cache issue when VxVM is installed on top of EMC PowerPath managed devices.

Workaround :

Install VxVM before installing and configuring EMC PowerPath

DMP uses OS device physical path to maintain persistence of path attributes from 6.0 (2410716)

From release 6.0, DMP uses OS device physical path instead of logical name to maintain persistence of path attributes. Hence after upgrading to DMP 6.0 or later releases, path attributes are reset to the default values. You must reconfigure the path-level attributes that were defined in the /etc/vx/dmppolicy.info file.

Workaround:

To configure path-level attributes:

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

2 Reset the path attributes.

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

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In SF 6.1, if this command is run while the volumes are online and being actively used, the shown %dirty may lag from actual percentage dirty for instant snap data cache object (DCO) volumes. That is, the command output may show less %dirty than actual.

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 5.x to 6.1 may fail for IBM XIV Series arrays (2715119)

Starting in the Symantec Storage Foundation 5.1 SP1 release, the Array Support Library (ASL) for the IBM XIV enclosures converts the LUN Serial Number from hexadecimal to decimal. Because of this change, the enclosure names differ from releases prior to the 5.1 SP1 releases. When you upgrade Symantec Storage Foundation from a release prior to that release to the current 6.1 release, XIV LUNs may go into an error state. Note that the latest RPs on 5.1/5.1SP1 are already modified to use the same logic for enclosure naming.

Workaround:

After the upgrade, run `vxddladm assign names`.

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

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

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

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

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

Workaround:

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

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

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

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

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

Workaround

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

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

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.

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.

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.

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 other commands to run slowly because the VxVM configuration daemon (`vxconfigd`) is busy servicing `vxdisk scandisks`.

Workaround: Stop the `vxattachd` script and:

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.

LUNs from any EMC CLARiiON arrays that have Not Ready state are shown in "online invalid" state by Veritas Volume Manager (VxVM) (3287940)

LUNs from any EMC CLARiiON arrays that have Not Ready state are shown in "online invalid" state by Veritas Volume Manager (VxVM). They should be shown in "error" state. The EMC CLARiiON arrays do not have a mechanism to communicate the NR (Not Ready) state of the LUNs, so VxVM cannot recognize it. However, the read operation on these LUNs fails and due to defect in disk online operation this read failure is ignored causing disk online to succeed. Thus, these LUNs are shown as "online invalid".

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) 6.1 includes several array names that differ from the array names in releases 5.1SP1 or prior. Therefore, if you upgrade to VxVM 6.1 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 6.1. Manually reconfigure the enclosure attributes to resolve the issue.

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

Table 1-11 Hitachi arrays with new array names

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

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

Veritas File System known issues

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

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 (2591473)

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

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

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

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

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

Workaround: Make more space available on the file system.

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

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

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

Workaround: Rerun the shrink operation after stopping the I/Os.

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

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

Workaround: There is no workaround.

In a VxFS file system that has compression enabled, the file system may get disabled due to the ENOSPC error (3301716)

In a VxFS file system that has compression enabled, the file system may get disabled due to the ENOSPC error. This occurs because of a defect in the delayed allocation feature.

Workaround: Turn off the delayed allocation.

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.

NFSv4 server panics in unlock path (3228646)

In a CFS configuration, if `fcntl(1m)` fails, some NFS specific structures `(I_pid)` are not updated correctly and may point to stale information. This causes the NFSv4 server to panic.

Workaround:

There is no workaround for this issue.

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.

Operations involving the fsetxattr system calls show slow performance (3294074)

On Linux, operations involving the `fsetxattr` system calls such as `cp -rp` show slow performance.

Workaround:

There is no workaround for this issue.

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” package in the various Linux distributions does not properly cleanse the untrusted format strings (3335691)

"rpc.statd” in the “nfs-utils” package 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” package to correct the problem.

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.

In case of scenarios where updates to writable clones are frequent, the clone operation may hang (3348553)

In case of scenarios where updates to writable clones are frequent, the clone operation may hang when a large directory hash is enabled, and inodes are reused aggressively.

Workaround: There is no workaround for this issue.

Replication known issues

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

vradmin syncvol command compatibility with IPv6 addresses (2075307)

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

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

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

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
```

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.

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: None

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 Storage Foundation HA 6.1 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*.

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 SF. 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.

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 of 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 6.1 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.1.

When upgrading from SF version 5.0 to SF 6.1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

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

Workaround

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

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

If you have a 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.

Workaround

There is no workaround for this issue.

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

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

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

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


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

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

Workaround

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

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

1. Shut down the database.
2. Unmount the `/dev/shm` file system:

   ```
   # umount /dev/shm
   ```

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

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

4 Start the database.

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

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

You might encounter an Oracle error such as the following:

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

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

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

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

Workaround:

Retry the cloning operation until it succeeds.

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

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

There is no workround at this point of time.

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

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

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

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

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

Workaround: For the 6.1 release, create distinct archive and datafile mounts for the checkpoint service.

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

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

Workaround: There is no workaround for this issue.

`sfua_rept_migrate` fails after phased SF Oracle RAC upgrade from 5.0MP3RP5 to 6.0.1 (2874322)

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

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

```
failed.
SFORA sfua_rept_migrate ERROR V-81-9162 Failed to umount repository.
```

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

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

For more information, see TECH64812.

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:

```
[oracle@dblxx64-3-vip3 ~]$ 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 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 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.

The SmartIO options are not restored after the Reverse Resync Commit operation is performed (3313775)

The RRCommit operation mounts file systems with a default file system option. However, the non-default configuration options for VxFS SmartIO are lost when a Reverse Resync Commit operation is performed.

Workaround: Remount the file systems with the required configuration options after a successful completion of the RRCommit operation.

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.

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.

The information file that is generated after a DBED data collector operation reports an error (2795490)

When the VRTSexplorer DBED scripts use the old VRTSdbms3-specific scripts that are removed from the products, the information file reports the following error:

```
/opt/VRTSdbms3/vxdbms_env.sh: cannot open [No such file or directory]
```

Workaround:

1. Run the `cd /opt/VRTSspt/DataCollector/sort` command. If this directory does not exist, run `sh /opt/VRTSspt/DataCollector/*.sh`.
2. Run the `cd advanced/lib/VOS/v10/Collector/VxExpCollector/explorer_scripts` command.
3. In `dbed_rept_sql`, comment

   `$`**`VXDBMS_DIR/vxdbms_env.sh`**

   Or

   Replace **$VXDBMS_DIR/vxdbms_env.sh** with

   ```
   [[ -f $VXDBMS_DIR/vxdbms_env.sh ]] &&
       {
           . $VXDBMS_DIR/vxdbms_env.sh
       }
   ```

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'
```

Instant mode clone fails in RAC environment for all FSMs with data loading (3517782)

When you use the instant clone mode for RAC databases, the clone operation may fail during Oracle recovery. The issue is more likely to be seen when there is load activity on some of the RAC nodes.

Workaround: Use either online or offline snapshot mode.

Virtualization known issues

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

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.

Disks exported using the VirtIO-disk interface from an SLES11 SP2 or SLES11 SP3 host are not visible to Veritas Volume Manager running inside KVM guests (3325022)

Devices exported from a host running SLES11 SP2 or SLES11 SP3, to guest using VirtIO-disk interfaces are not supported by ASL, which claims the devices during device discovery. Therefore the devices are not visible to Veritas Volume Manager running inside a KVM guest.

For example, if disks vda and vdb are the only exported disks to guest “guest1” using the VirtIO-disk interface, then they are not visible in the `vxdisk list` output:

```
guest1:~ # vxdisk list
DEVICE       TYPE            DISK         GROUP        STATUS
```

As none of the ASLs have claimed the devices:

```
guest1:~ # vxddladm list devices
DEVICE               TARGET-ID    STATE   DDL-STATUS (ASL)
============================================================
vdb                  -            Online  -
vda                  -            Online  -
```

Workaround: Devices exported using VirtIO-LUN interfaces to a KVM guest, are supported by ASLs. As SLES11 SP2 does not support the VirtIO-LUN interface, upgrade the host and guest (if running SLES11 SP2) to SLES11 SP3 along with the SF/SFHA/SFCFS stack. After performing the upgrade, export the devices using VirtIO-LUN interface to the KVM guest. In the guest, devices are claimed by the respective vendor ASL. Refer to SUSE documentation and SORT for the latest patches and upgrade best practices.

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 94.

Veritas Volume Manager software limitations

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

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

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

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

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

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

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

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

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

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

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

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.

Symantec Storage Foundation for Databases (SFDB) tools software limitations

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

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

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

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

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

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

Virtualization software limitations

This section describes the virtualization software limitations in this release of Symantec Storage Foundation (SF).

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.

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.

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 documentation

Table 1-12 lists the documentation for Symantec Storage Foundation.

Table 1-12 Symantec Storage Foundation documentation

Document title	File name	Description
Symantec Storage Foundation Release Notes	sf_notes_61_lin.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.
Symantec Storage Foundation Installation Guide	sf_install_61_lin.pdf	Provides information required to install the product.
Symantec Storage Foundation Administrator's Guide	sf_admin_61_lin.pdf	Provides information required for administering the product.
Symantec Storage Foundation: Storage and Availability Management for DB2 Databases	sfhas_db2_admin_61_unix.pdf	Provides information about the deployment and key use cases of the SFDB tools with Storage Foundation High Availability (SFHA) Solutions products in DB2 database environments. It is a supplemental guide to be used in conjunction with SFHA Solutions product guides.
Symantec Storage Foundation: Storage and Availability Management for Oracle Databases	sfhas_oracle_admin_61_unix.pdf	Provides information about the deployment and key use cases of the SFDB tools with Storage Foundation High Availability (SFHA) Solutions products in Oracle database environments. It is a supplemental guide to be used in conjunction with SFHA Solutions product guides.

Table 1-12 Symantec Storage Foundation documentation *(continued)*

Document title	File name	Description
Veritas File System Programmer's Reference Guide (This document is available online only.)	vxfs_ref_61_lin.pdf	Provides developers with the information necessary to use the application programming interfaces (APIs) to modify and tune various features and components of the Veritas File System.

Symantec Storage Foundation and High Availability Solutions products documentation

Table 1-13 lists the documentation for Symantec Storage Foundation and High Availability Solutions products.

Table 1-13 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_61_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 Veritas 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_61_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_61_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.

Table 1-13 Symantec Storage Foundation and High Availability Solutions products documentation *(continued)*

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions SmartIO for Solid State Drives Solutions Guide	sfhas_smartio_solutions_61_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_61_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.
Symantec Storage Foundation and High Availability Solutions Replication Administrator's Guide	sfhas_replication_admin_61_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_61_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