Symantec™ Storage Foundation Release Notes

Solaris

6.2

October 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: 6.2

Document version: 6.2 Rev 0

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 6.2 for Solaris. 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: 6.2 Rev 0" 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 storage devices that are configured on the system. DMP creates DMP metadevices across all of the paths to each LUN. DMP uses the DMP metadevices to manage path failover and I/O 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 storage abstraction layer or storage management 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.
Volume Replicator (VVR)	Enables you to maintain a consistent copy of application data at one or more remote locations for disaster recovery. VVR replicates data to remote locations over any standard IP network to provide continuous data availability. VVR can replicate existing VxVM configurations, and can be transparently configured while the application is active. VVR option is a separately-licensable feature of Symantec Storage Foundation.

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

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

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

Changes introduced in 6.2

This section lists the changes in Symantec Storage Foundation 6.2.

Changes related to installation and upgrades

The product installer includes the following changes in 6.2.

Connecting to the SORT website through a proxy server

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

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

Support for centralized installations using the Deployment Server

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

The Deployment Server lets you do the following as described in Table 1-2.

Table 1-2 Deployment Server functionality

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

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

See the *Symantec Storage Foundation Installation Guide* for more information.

Support for upgrading SF using the web-based installer for Solaris 10 Live Upgrade

You can use the Symantec web-based installer to upgrade SF as part of the Live Upgrade.

Run the web-based installer on the DVD to upgrade SF.

The program uninstalls the existing version of SF on the alternate boot disk during the process. At the end of the process, SF 6.2 is installed on the alternate boot disk.

Support for upgrading SF using the web-based installer for Boot Environment on Solaris 11

You can use the Symantec product installer to upgrade SF on a BE.

Run the web-based installer on the DVD to upgrade SF.

At the end of the process, the SF 6.2 is installed on the alternate BE.

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

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

Changes related to Symantec Storage Foundation (SF)

Symantec Storage Foundation includes the following changes in 6.2:

SmartIO: Support for caching on Solid-State Drives

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 solid-state devices 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 6.2:

Layered volume enhancements for recovery and snapshots

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

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

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

Read policy enhancement

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

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

Changes related to Veritas File System

Veritas File System (VxFS) includes the following changes in 6.2:

Changes related to replication

Symantec Storage Foundation and High Availability Solutions includes the following changes related to replication in 6.2:

Changes related to SFDB tools

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

Support for multitenant databases

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

Release level terminology changes

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

Table 1-3 Release level terminology changes

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

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

No longer supported

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

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

System requirements

This section describes the system requirements for this release.

Supported Solaris 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	Chipsets
Solaris 10	Update 9, 10, and 11	SPARC
Solaris 11	Solaris 11.1 and up to Support Repository Update (SRU) 11.1.21.0.4.1 Solaris 11.2 and up to Support Repository Update (SRU) 11.2.1.0.2.0	SPARC

This release is not supported on the x86-64 architecture.

This release supports native brand zones on Solaris 10 operating system and solaris brand and solaris10 brand zones on the Solaris 11 operating system. This release does not support the Kernel Zones feature of Solaris 11 Update 2.

Supported Oracle VM Server for SPARC

Supported Oracle VM Server for SPARC versions are OVM 2.0, OVM 2.1, OVM 2.2, OVM 3.0, and OVM 3.1.

For supported OS version for Oracle VM Server for SPARC, refer to *Oracle VM server for SPARC Release Notes*.

The version of the Oracle Solaris operating system (OS) that runs on a guest domain is independent of the Oracle Solaris OS version that runs on the primary domain. Therefore, if you run the Oracle Solaris 10 OS in the primary domain, you can still run the Oracle Solaris 11 OS in a guest domain. Likewise if you run the Oracle Solaris 11 OS in the primary domain, you can still run the Oracle Solaris 10 OS in a guest domain.

The only difference between running the Oracle Solaris 10 OS or the Oracle Solaris 11 OS on the primary domain is the feature difference in each OS.

Storage Foundation for Databases features supported in database environments

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

Table 1-5 SFDB features supported in database environments

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

Notes:

- SmartTier is an expanded and renamed version of Dynamic Storage Tiering (DST).
- Storage Foundation for Databases (SFDB) tools Database 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

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

Table 1-6 Fixed issues related to installation and upgrades

Incident	Description
3325954	On Solaris 10 `xprtld` will not be started if user use jumpstart to install product
3326196	Rolling upgrade may encounter a problem if open volumes from different disk groups have the same name.
3343592	`installer -license sys1` command does not show the check result if the system already has the SF Basic license key registered.
3442070	If you select rolling upgrade task from the Install Bundles menu, the CPI exits with an error.

Symantec Storage Foundation fixed issues

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

See "Veritas File System fixed issues" on page 20.

See "Veritas Volume Manager fixed issues" on page 22.

Veritas File System fixed issues

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

Table 1-7 Veritas File System fixed issues

Incident	Description
3597482	The `pwrite(2)` function fails with the `EOPNOTSUPP` error.
3563796	The file system `fullfsck` flag is set when the inode table overflows.
3560968	The `delicache_enable` tunable is not persistent in the Cluster File System (CFS) environment.
3560187	The kernel may panic when the buffer is freed in the `vx_dexh_preadd_space()` function with the message `Data Key Miss Fault in kernel mode`.
3550103	After you upgrade or restart the system, mismatch in SSD cache usage may occur.
3520349	When there is a huge number of dirty pages in the memory, and a sparse write is performed at a large offset of 4 TB or above on an existing file that is not null, the file system hangs.
3486726	Veritas File Replicator (VFR) logs too much data on the target node.
3484336	The `fidtovp()` system call can panic in the `vx_itryhold_locked()` function.
3478017	Internal tests found assert in voprwunlock
3469644	The system panics in the `vx_logbuf_clean()` function.
3466020	The file system is corrupted with the error message `vx_direrr: vx_dexh_keycheck_1`.
3463464	Internal kernel functionality conformance tests hit a kernel panic due to null pointer dereference.
3457803	The file system gets disabled intermittently with metadata I/O errors.
3451284	While allocating extents during the write operation, if summary and bitmap data for file system allocation units get mismatched, a full `fsck` flag get set on the file system.

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

Incident	Description
3449150	The `vxtunefs`(1M) command accepts garbage values for certain tunables.
3448702	Checkpoint creation of different file systems may get serialized.
3444775	Internal noise testing on cluster file system results in a kernel panic in function `vx_fsadm_query()` with an error message.
3434811	The `vxfsconvert`(1M) command in VxFS 6.1 hangs.
3424564	`fsppadm` fails with `ENODEV` and `file is encrypted or is not a database` errors.
3417076	The `vxtunefs`(1M) command fails to set tunables when the file contains blank lines or white spaces.
3415639	The type of the `fsdedupadm`(1M) command always shows as MANUAL even when it is launched by the `fsdedupschd` daemon.
3413926	Internal testing hangs due to high memory consumption, resulting in fork failures.
3394803	A panic is observed in the VxFS routine `vx_upgrade7`() function while running the `vxupgrade` command (1M).
3340286	After a file system is resized, the tunable setting `dalloc_enable` is reset to a default value.
3335272	The `mkfs` (make file system) command dumps core when the log size provided is not aligned.
3332902	While shutting down, the system running the `fsclustadm`(1M) command panics.
3323866	Some Object Data Manager (ODM) operations fail with `ODM ERROR V-41-4-1-328-22 Invalid argument.`
3317368	File system operations needing a file system freeze may take longer in the presence of file level snapshots and when there is a heavy I/O load.

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

Incident	Description
3301716	In a Veritas File System (VxFS), if a file system has compression enabled, the file system gets disabled due to the ENOSPC error. This occurs because of a defect in the delayed allocation feature.
3297840	VxFS corruption is detected during a dynamic LUN resize operation.
3264062	The allocated space may be leaked from the free space while you are unsharing shared extents.
3228646	The NFSv4 server in unlock path.
3121933	There is a DB2 crash or corruption when `EOPNOTSUPP` is returned from VxFS.

Veritas Volume Manager fixed issues

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

Table 1-8 Veritas Volume Manager fixed issues

Incident	Description
3554608	Mirroring a volume creates a larger plex than the original one on a Cross-platform Data Sharing (CDS) disk.
3553407	The SmartMove feature cannot work on layered volumes that do not have thin disks.
3542272	The `vxconfigbackup` process takes a long time to complete.
3446415	A `vxassist` command volume shrink operation ignores specified disk sites.
3440790	After the 6.0.3 upgrade, the `vxassist mirror` command hangs with `vxtask list` in the auto-throttled state.
3339195	The `vxdiskadm` command cannot exclude or suppress devices.

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

Incident	Description
3317430	After an upgrade from 5.1 SP1 RP4, the `vxdiskunsetup` command displays the following message: `ERROR Device unexport failed: Operation is not supported.`
3287940	LUNs from any EMC CLARiiON arrays that have Not Ready state are shown in the "online invalid" state by Veritas Volume Manager (VxVM)
2049952	The `vxrootadm` command shows incorrect messages in Japanese with localization for Solaris.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues

Table 1-9 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-9 SFDB tools fixed issues

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

Known issues

This section covers the known issues in this release.

Installation known issues

This section describes the known issues during installation and upgrade.

On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail (2424410)

On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail with the following error:

```
Generating file list.
Copying data from PBE <source.24429> to ABE <dest.24429>.
99% of filenames transferredERROR: Data duplication process terminated
unexpectedly.
ERROR: The output is </tmp/lucreate.13165.29314/lucopy.errors.29314>.

29794 Killed
Fixing zonepaths in ABE.
Unmounting ABE <dest.24429>.
100% of filenames transferredReverting state of zones in PBE
<source.24429>.
ERROR: Unable to copy file systems from boot environment <source.24429>
to BE <dest.24429>.
ERROR: Unable to populate file systems on boot environment <dest.24429>.
Removing incomplete BE <dest.24429>.
ERROR: Cannot make file systems for boot environment <dest.24429>.
```

This is a known issue with the Solaris `lucreate` command.

Workaround: Check with Oracle for possible workarounds for this issue.

vxlustart failed due to lumount error when performing Live Upgrade to Solaris 10 Update 11 (3035982)

Live Upgrade (LU) to Solaris 10 Update 11 using `vxlustart` fails with following error:

```
# lumount -n dest.7667 /altroot.5.10
ERROR: mount point directory </altroot.5.10> is not empty
ERROR: failed to create mount point </altroot.5.10> for file system
</dev/dsk/c1t1d0s0>
ERROR: cannot mount boot environment by name <dest.7667>
ERROR: vxlustart: Failed: lumount -n dest.7667 /altroot.5.10
```

Workaround: To perform Live Upgrade to Solaris 10 Update 11, use one of the following procedures for your operating system version.

To perform Live Upgrade from Solaris 10 Update 10 to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 10 Update 9 or below to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 9 to Solaris 10 Update 11

1. Upgrade from Solaris 9 to Solaris 10 Update 10.
2. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
3. Use `vxlustart` to upgrade to Solaris 10 Update 11.

On Solaris 11 non-default ODM mount options will not be preserved across package upgrade (2745100)

On Solaris 11, before the package upgrade if Oracle Disk Manager (ODM) is mounted with non-default mount options such as nocluster, nosmartsync etc, these mount options will not get preserved after package upgrade.

There is no workaround at this time.

Live Upgrade to Solaris 10 Update 10 fails in the presence of zones (2521348)

SFCFSHA Live Upgrade from Solaris 10 Update 7 5.1SP1 to Solaris 10 Update 10 using the `vxlustart` commands fails in the presence of zones with the following error message:

```
ERROR: Installation of the packages from this media of the media failed;
pfinstall returned these diagnostics:
Processing default locales
        - Specifying default locale (en_US.ISO8859-1)
Processing profile
ERROR: This slice can't be upgraded because of missing usr packages for
the following zones:
ERROR:    zone1
ERROR:    zone1
```

```
ERROR: This slice cannot be upgraded because of missing usr packages for
one or more zones.
The Solaris upgrade of the boot environment <dest.27152> failed.
```

This is a known issue with the Solaris `luupgrade` command.

Workaround: Check with Oracle for possible workarounds for this issue.

During upgrade from 5.1SP1 to 6.2 with an encapsulated root disk, splitting the root mirror fails if the target disk group name is used by a deported disk group (2280560)

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

Workaround:

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

On Solaris 10, a flash archive installed through JumpStart may cause a new system to go into maintenance mode on reboot (2379123)

If a Flash archive is created on a golden host with encapsulated root disks, when this Flash archive is installed onto another host through JumpStart, the new system may go to maintenance mode when you initially reboot it.

This problem is caused by the predefined root disk mirror in the Flash archive. When the archive is applied to a clone system, which may have different hard drives, the newly cloned system may get stuck at root disk mirroring during reboot.

Workaround: Create the Flash archive on a golden host with no encapsulated root disks. Run `vxunroot` to clean up the mirrored root disks before you create the Flash archive.

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 Live Upgrade to Solaris 10 Update 10/Update 11, boot from an alternate boot environment fails [2370250]

If your setup involves volumes in a shared disk group that are mounted as CFS in a cluster, then during Live Upgrade using the `vxlustart` command from any supported Solaris version to Solaris 10 Update 10/11, boot from an alternate boot environment may fail.

Workaround:

1. Run the `vxlufinish` command. Enter:

   ```
   # vxlufinish
   ```

2. Manually delete the entries of all the volumes of shared disks that are mounted as CFS in the `/altroot.5.10/etc/vfstab` directory. Enter:

   ```
   rm -rf /altroot.5.10/etc/vfstab
   ```

3. Restart the system.

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

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

Workaround:

The problem has no critical functionality impact. However, for performance considerations, it is recommended that you verify that the value of the

vol_min_lowmem_sz and vol_max_nmpool_sz tunables are set to at least the default value. Use the vxtune command to modify the tunable value.

Installing VRTSvlic package during live upgrade on Solaris system non-global zones displays error messages [3623525]

While installing VRTSvlic package during live upgrade on Solaris system with non-global zones following error messages are displayed:

```
cp: cannot create /a/sbin/vxlicinst: Read-only file system
cp: cannot create /a/sbin/vxlicrep: Read-only file system
cp: cannot create /a/sbin/vxlictest: Read-only file system
```

Workaround: This message can be ignored. The vxlicinst, vxlicrep, vxlictest utilities are present in /opt/VRTSvlic/sbin/ inside a non-global zone.

Symantec Storage Foundation known issues

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

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.

Veritas Volume Manager known issues

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

The system may panic during shutdown (3107699)

Due to Oracles aggressive driver loading policy, any probe to the driver device causes concerned driver and drivers in dependency hierarchy. Hence a race condition occurs sometimes leading to system panic when you shut it down.

Workaround:

Before you shut down or restart the system, enter the following command:

```
vxdmpadm iostat stop
```

The `vxdisk resize` command does not claim the correct LUN size on Solaris 11 during expansion of the LUN from array side (2858900)

The `vxdisk resize` command fails on Solaris 11 during expansion of the LUN from array side. The `vxdisk resize` command does not claim correct LUN size on Solaris 11 during expansion of the LUN from array side. This is because of Oracle bug -19603615. On Solaris 11, the `vxdisk resize` command may exit without any error, returning incorrect LUN size or failing with similar error as follows:

```
bash# vxdisk -g testdg resize disk01 length=8g
VxVM vxdisk ERROR V-5-1-8643 Device disk01: resize failed:Operation would blo
```

Workaround:

There is no workaround available which can work in all the configuration. In some specific configurations, the following workaround works:

After expansion of LUN from array side, run `format -d` command and then run `vxdisk resize` command.

Root disk encapsulation fails for root volume and swap volume configured on thin LUNs (3538594)

Root disk encapsulation fails if the root disk configuration on a thin LUN includes volumes such as `var`, `usr`, or `home`, in addition to the root volumes and the swap volumes. Root disk encapsulation is not supported in this configuration.

Workaround:

There is no workaround.

SmartIO VxVM cache invalidated after relayout operation (3492350)

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

Workaround:

This behavior is expected. There is no workaround.

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

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

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

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

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

Workaround:

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

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

For Solaris 11.1 or later, the system can panic when system is rebooted after turning dmp_native_support to on (3341674)

For Solaris 11.1 or later when more than 512 LUNs are configured, the system can panic when the system is rebooted after setting the tunable parameter `dmp_native_support` to `on`.

Workaround:

For Solaris 11.1 or later, DMP native support for ZFS is restricted to set-ups with no more than 512 LUNs.

Disk greater than 1TB goes into error state (3269099)

If a path of a device having multiple paths is labelled with the EFI format using an operating system command such as `format`, the `vxdisk` list command output shows the device in error state.

Workaround:

This issue is a Solaris OS issue. There is no workaround for this issue.

Importing an exported zpool can fail when DMP native support is on (3133500)

On Solaris, when the tunable `dmp_native_support` is set to `on`, importing an exported zpool using the command `zpool import` *`poolname`* can fail with following error:

```
Assertion failed: rn->rn_nozpool == B_FALSE, file
../common/libzfs_import.c,
line 1084, function zpool_open_func
Abort (core dumped)
```

Workaround:

Import the zpool using the following command, specifying the DMP device directory:

```
# zpool import -d /dev/vx/dmp poolname
```

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
```

Probing vxio with DTrace fails on Sparc machines. (2180635)

This issue exists because of inability of DTrace to load a module whose text size is greater than 2MB on Sparc machines. While trying to load `vxio` with DTrace you may see following warning messages on console:

```
dtrace: WARNING: couldn't allocate SDT table for module vxio
fbt: WARNING: couldn't allocate FBT table for module vxio
```

There is no workaround for this issue.

Disks on the Oracle VM Server for SPARC guest are claimed under other_disks category (2354005)

The disks on the Oracle VM Server for SPARC guest are claimed under "other_disks" enclosure, because these disks are not capable of being multi-pathed by DMP. This is expected because these devices represent VxVM volumes in the host. By design, devices under other_disks enclosure have their name based on underlying OS path regardless of the DDL naming scheme.

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 any path-level attributes that were defined in the /etc/vx/dmppolicy.info file.

Workaround:

To configure path-level attributes

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

2 Reset the path attributes.

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

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In SF 6.0, if this command is run while the volumes are online and being actively used, the shown %dirty may lag from actual percentage dirty for instant snap data cache object (DCO) volumes. That is, the command output may show less %dirty than actual.

After changing the preferred path from the array side, the secondary path becomes active (2490012)

For EVA arrays, DMP requires that the prefer bit is static. If the prefer bit is not static, issues like the following may occur. After changing the prefer path of LUN from the array side, and performing a disk discovery (`vxdisk scandisks`) from the host, the secondary path becomes active for the LUN.

Workaround:

To work around this issue

1 Set the pref bit for the LUN.

2 Perform disk discovery again:

```
# vxdisk scandisks
```

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 6.2 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 6.2 release, XIV LUNs may go into an error state. Note that the latest RPs on 5.1/5.1SP1 are already modified to use the same logic for enclosure naming.

Workaround:

After the upgrade, run `vxddladm assign names`.

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

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

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

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

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

Workaround:

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

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

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

The vxdisksetup command fails to initialize disks in cdsdisk format for disks in logical domains greater than 1 TB (2557072)

The `vxdisksetup` command fails to initialize disks in `cdsdisk` format for disks in logical domains greater than 1 TB. This issue is due to an Oracle VM Server command which fails when the number of partitions in the GUID partition table (GPT) label is greater than 9. The `cdsdisk` format requires at least 128 partitions to be compatible with Linux systems.

Workaround: There is no workaround for this issue.

1 TB luns goes in error state with Solaris x86 (2706776)

If you label a disk device as EFI using format on a subset of the paths or on the DMP device, Solaris will not be able to propagate the label to all the other paths of the LUN. This will lead the device to appear in the error state under 'vxdisk list'.

Workaround: There is no workaround for this issue.

The vxrecover command does not handle RAID5 volumes correctly (2715124)

On Solaris 11, the `vxrecover` command calls the recovery process for the top-level volume, which internally takes care of recovering its subvolumes. The `vxrecover` command does not handle RAID5 volumes correctly. The recovery process fails to recover the subvolumes, which remain in the NEEDSYNC state.

Workaround:

Manually recover the RAID5 volumes using the `vxvol` utility, as follows:

```
# vxvol -g diskgroup resync volume
```

vxmirror to SAN destination failing when 5 partition layout is present: for example, root, swap, home, var, usr (2815311)

The `vxmirror` command may fail with following error on a Solaris 10 host, for a thin LUN, if more than one partition excluding root and swap is present.

```
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
```

Example

```
# /etc/vx/bin/vxmirror" -f -g rootdg_17_23_49 rootdisk01 rootdisk02
        ! vxassist -g rootdg_17_23_49 mirror swapvol rootdisk02
```

```
! vxassist -g rootdg_17_23_49 mirror rootvol rootdisk02
! vxassist -g rootdg_17_23_49 mirror usr rootdisk02
! vxassist -g rootdg_17_23_49 mirror var rootdisk02
! vxassist -g rootdg_17_23_49 mirror home rootdisk02
! vxbootsetup -g rootdg_17_23_49
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'home_dcl'
 because no free partitions are available on disk 'disk_0'.
 Either remove the volume or make a partition available
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'usr_dcl'
 because no free partitions are available on disk 'disk_0'.
 Either remove the volume or make a partition available
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'var_dcl' because
 no free partitions are available on disk 'disk_0'.
 Either remove the volume or make a partition available
/usr/lib/vxvm/bin/vxmksdpart: 3pardata0_2492: is not an identifier
```

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.

Workaround:

There is no workaround available.

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

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

Workaround:

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

For Solaris 11.1 or later, system hangs when both QLogic and Emulex HBAs are present and dmp_native_support is turned on (3138703)

For Solaris 11.1 or later, the system may hang when both QLogic and Emulex HBAs are present, and **dmp_native_support** is turned on.

Workaround:

The system hang is not seen if all of the HBAs are either from Emulex or from QLogic. Do not combine both HBAs on the same system.

For Solaris 11.1 or later, enabling DMP native support requires steps to enable booting from alternate root pools (3133514)

For Solaris 11.1 or later, if the tunable parameter dmp_native_support is set to on, using the following command causes alternate root pools on OS devices to migrate to DMP devices:

```
# zpool import -d /dev/vx/dmp
```

After the above command is run, the system cannot boot using these alternate root pools because the DMP driver is not configured for these root pools. This scenario is shown by the following output.

```
# zpool status

pool: crpool
 state: ONLINE
  scan: none requested
config:
```

```
NAME                      STATE   READ WRITE CKSUM
crpool                    ONLINE  0     0     0
  /dev/vx/dmp/disk_0s0    ONLINE  0     0     0
```

Workaround:

To boot using the alternate root pools, export and re-import the root pools using the OS device.

To boot using the alternate root pools

1 Export the root pool:

```
# zpool export crpool
```

2 Display the OS path name for the device:

```
# vxdmpadm getsubpaths dmpnodename=disk_0
NAME     STATE[A]   PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS
=====================================================================
c3t2d0s2 ENABLED(A) -            c3        Disk       disk       -
```

3 Re-import the root pools using the OS device.

```
# zpool import crpool -d /dev/dsk/c3t2d0s0
```

The system is now bootable using the alternate root pools.

For Solaris 11.1 or later, uninstalling DMP or disabling DMP native support requires steps to enable booting from alternate root pools (3178642)

For Solaris 11.1 or later, after you uninstall the VxVM package or after you turn off DMP native support, you may see this issue. After reboot, the root pool containing the active boot environment is migrated to the OS device but alternate root pools continue to show DMP device. The status of the alternate root pools and their DMP devices is shown as "UNAVAIL".

```
pool: crpool
 state: UNAVAIL
status: One or more devices are unavailable in response to persistent
        errors. There are insufficient replicas for the pool to continue
        functioning.
action: Destroy and re-create the pool from a backup source. Manually
        marking the device repaired using 'zpool clear' or 'fmadm repaired'
        may allow some data to be recovered.
        Run 'zpool status -v' to see device specific details.
```

```
  scan: none requested
config:

        NAME                    STATE     READ WRITE CKSUM
        crpool                  UNAVAIL      0     0     0
          emc_clariion1_82s0    UNAVAIL      0     0     0
```

The tunable parameter `dmp_native_support` only unconfigures DMP for the single root pool containing the active boot environment. If the setup has any alternate root pools, for which DMP native support was enabled, then the alternate root pools continue to show the DMP device. If the alternate root pool is configured in the current boot environment and DMP support is removed, the DMP devices required for ZFS are not found. The DMP devices and the root pools display the state as "UNAVAIL".

Workaround:

Even though the status of alternate root pool is "UNAVAIL", the system is bootable using the disk containing the alternate root pool. Reboot the system with the disk containing the alternate root pool. The system comes up with the root pool using the DMP device.

For Solaris 11.1 or later, after enabling DMP native support for ZFS, only the current boot environment is bootable (3157394)

After enabling DMP native support for ZFS on Solaris 11.1 or later, only the current boot environment (BE) is bootable. Any alternate BEs in the same root pool are not bootable. This situation occurs because the DMP native support configures the ZFS root pool so that only DMP can import the root pool. If you attempt to boot the system from the alternate BE, the system panics with the following message:

```
NOTICE: zfs_parse_bootfs: error 19
Cannot mount root on rpool/193 fstype zfs

panic[cpu0]/thread=10012000: vfs_mountroot: cannot mount root

Warning - stack not written to the dumpbuf
000000001000fa00 genunix:main+17c (1, 100dc958, 12d5c00, 124702c, 0, 10828000)
  %l0-3: 0000000010010000 0000000000000000 00000000100dc800 0000000000000000
  %l4-7: 0000000010012000 0000000000000000 000000001038f7c0 000000000104c800
```

Workaround:

To enable booting from another BE, configure the ZFS root pool so that it can be imported without DMP.

To configure ZFS root pool to enable booting from all the BEs

1 At the OBP PROM, run the following command to list all the BEs:

```
ok> boot -L
```

2 Use the following command to boot from the BE for which DMP native support for ZFS is enabled.

```
ok> boot -Z rpool/ROOT/BE_name
```

3 After booting through new BE, disable the DMP native support using the following command:

```
# vxdmpadm settune dmp_native_support=off
```

The system is now bootable from any BEs in the ZFS root pool.

Limitation to DMP support for ZFS root in the Oracle VM Server for SPARC guest (3221944)

DMP support for ZFS root is not supported in the Oracle VM Server for SPARC guest. If DMP meta devices are exported to Oracle VM Server for SPARC and used for root pool, then enabling the dmp_native_support tunable fails with the following error:

```
root@swsx39-v05#vxdmpadm settune dmp_native_support=on
VxVM vxdmpadm ERROR V-5-1-15690 Operation failed for one or more
zpools

VxVM vxdmpadm ERROR V-5-1-15686 The following zpool(s) could not
be migrated as failed to obtain root pool information -

        rpool
```

Where *rpool* specifies the root pool name on the Oracle VM Server for SPARC guest:

DMP supports non-root ZFS in the Oracle VM Server for SPARC guest. You can use DMP devices for non-root ZFS.

When dmp_native_support is set to on, commands hang for a long time on SAN failures (3084656)

When dmp_native_support is set to on, on SAN failures, commands that do I/O operations to the root file system or I/O to disks that contain the root pool may hang for about 1-5 minutes. The commands include commands like "zpool status", or telnet initiated to connect the system. The hang is seen because the drivers below the DMP layer take more time to report the I/O failure when some of the paths to the disk containing the root pool are disconnected. This situation should not lead to any root pool data corruption.

Workaround:

This hang cannot be avoided but the hang time can be reduced by tuning the following parameters

To tune the parameters

1 In the `/kernel/drv/fp.conf` file, set

```
fp_offline_ticker=15
```

2 In the `/kernel/drv/fcp.conf` file, set

```
fcp_offline_dely=10
```

3 Reboot the sytem to apply the changes.

These steps reduce the hang time to a maximum of 1 minute.

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

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

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

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

Because of events that are generated as part of the online process, the `vxattachd` script triggers the vxdisk scandisks command again. This cycle causes continuous I/O error messages. This problem can also cause other commands to run slowly

because the VxVM configuration daemon (`vxconfigd`) is busy servicing `vxdisk scandisks`.

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

1 Disable the `vxattachd` process.

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

2 Set the following EMC CLARiiON values:

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

Enter:

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

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

The Veritas Volume Manager (VxVM) 6.2 includes several array names that differ from the array names in releases 5.1SP1 or prior. Therefore, if you upgrade to VxVM 6.2 from a release 5.1SP1 or earlier, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.2.

Workaround:

Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-10 shows the Hitachi arrays that have new array names.

Table 1-10 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

Table 1-10 Hitachi arrays with new array names *(continued)*

Previous name	New name
Hitachi AMS2300 Series arrays	New array names are based on the Model Number 8x. For example, AMS_100, AMS_2100, AMS_2300, AMS_2500, etc.

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.2. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

MPxIO device names shown in error state (3169587)

In this release, DMP does not support extended attributes like AVID for Solaris MPxIO devices. Up until the 5.1SP1 release, DMP used to support AVID for the MPxIO devices. When you upgrade from 5.1SP1 or prior release to 6.0 or later release, DMP assigns new names to the MPxIO devices.

The MPxIO device may go into an error state after the upgrade, if a persistent disk access record (entry in `/etc/vx/darecs`) exists with the old name, and the device was assigned a new name.

The same issue may occur if the MPxIO device name changes for another reason, such as the changed cabinet serial numbers for 3PAR or XIV devices from 6.0.

Workaround:

Use the following procedure to remove the persistent disk access record and resolve the issue.

To resolve the issue with MPxIO devices in error state

1 Remove the following file:

```
# rm /etc/vx/darecs
```

2 Reset the `vxconfigd` daemon:

```
# vxconfigd -kr reset
```

When all Primary/Optimized paths between the server and the storage array are disconnected, ASM disk group dismounts and the Oracle database may go down (3289311)

The Oracle database shows an I/O error on the control file, but there was no I/O error seen on any DMP device. When all Primary/Optimized paths are disconnected, DMP fails over to other available paths but the failover takes time. In the meantime, the application (ASM/Oracle database) times out the I/O.

The ASM alert log file displays messages such as the following:

```
Errors in file /u01/app/oracle/diag/rdbms/orcl/orcl2/trace/orcl2_ckpt_6955.trc:
ORA-00221: error on write to control file
ORA-00206: error in writing (block 4, # blocks 1) of control file
ORA-00202: control file: '+DATA_P6/ORCL/CONTROLFILE/current.261.826783133'
ORA-15081: failed to submit an I/O operation to a disk
ORA-15081: failed to submit an I/O operation to a disk
Wed Oct 09 14:16:07 2013
WARNING: group 2 dismounted: failed to read virtual extent 0 of file 261
Wed Oct 09 14:16:07 2013
USER (ospid: 6955): terminating the instance due to error 221
Wed Oct 09 14:16:07 2013
WARNING: requested mirror side 2 of virtual extent 0 logical extent 1 offset
16384
is not allocated; I/O request failed
WARNING: requested mirror side 3 of virtual extent 0 logical extent 2 offset
16384
is not allocated; I/O request failed
```

The above issue may occur when the server is configured as follows:

DB: Oracle 12c

Volume Manager: ASM

Multi-pathing Solutions: DMP

OS: Solaris

Disk Array : HP EVA in ALUA mode

Workaround:

The following workaround can reduce the probability of this issue, and when you see this issue, you could use Oracle commands to start the database manually.

Increase the application time out and make the following changes to reduce the time taken to mark the path as offline:

- In the `/kernel/drv/fp.conf` file, add `fp_offline_ticker=15`.

- In the `/kernel/drv/fcp.conf` file, add `fcp_offline_delay=10`.

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

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

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

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

For example:

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

The administrator must explicitly enable and disable support for a clone device created from an existing root pool (3152984)

A non-rpool is a clone of the existing root pool. When native support is enabled, DMP does not touch the clone root pool because the clone may or may not have the VxVM package.

Workaround: To add or remove DMP support for a clone boot device, the administrator must boot through the clone and turn on/off `dmp_native_support`.

Veritas File System known issues

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

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

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

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

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

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

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

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

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

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

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

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

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.

Warning message sometimes appear in the console during system startup (2354829)

During system startup, following messages sometimes appear in system console:

```
WARNING: couldn't allocate SDT table for module vxfs
WARNING: couldn't allocate FBT table for module vxfs
Loading smf(5) service descriptions: 2/2
```

These warnings indicate that the SDT and FBT DTrace probes might not be available for the VxFS module. The VxFS module still loads and works correctly. Dtrace SDT/FBT has limits on the size of module that it can support. Since the VxFS module exceeds the size that Dtrace can support, SDT and FBT Dtrace probes might not work for VxFS.

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.

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.

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.

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.

On a system that has Solaris 11 Update 1, certain driver modules such as "fdd" may not be removed properly (3348829)

On a system that has Solaris 11 Update 1, certain driver modules such as “fdd" may not be removed properly during the uninstallation of the SF or SFCFS stack.

Workaround: Prior to uninstallation of the stack, this can be mitigated by following the workaround indicated below:

```
# rm /usr/kernel/drv/sparcv9/fdd
```

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

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

For example:

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

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

Issue: After upgrading VVR to an IPv6-only environment in 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 stop
# /etc/init.d/vras-vradmind.sh start
```

vxassist relayout removes the DCM (145413)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

1 Pause or stop the applications.

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

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

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

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

Workaround:

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

1 Pause or stop the applications.

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

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

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

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

vradmin verifydata operation fails if the RVG contains a volume set (2808902)

In a VVR environment, the `vradmin verifydata` command fails with the following error if the replicated volume group (RVG) contains any volume set:

```
Message from Primary:
VxVM VVR vxrsync ERROR V-5-52-2009 Could not open device
/dev/vx/dsk/vvrdg/<volname> due to: stat of raw character volume path
failed
```

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.

SRL resize followed by a CVM slave node join causes the RLINK to detach (3259732)

In a CVR environment, performing a CVM slave node join after an SRL resize may stop replication due to a detached RLINK.

Workaround:

There is no workaround for this issue.

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

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

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

Workaround:

To resolve this issue:

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

   ```
   export IPM_HEARTBEAT_TIMEOUT
   IPM_HEARTBEAT_TIMEOUT=120
   ```

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

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

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

In a Volume Replicator (VVR) environment, after you start the initial synchronization with the `vradmin -a startrep` command with file system mounted on the primary

data volumes, the `vradmin repstatus` command does not show that the SmartSync feature is running. This is an only issue with the output of the `vradmin repstatus` command.

Workaround:

To confirm that SmartSync is running, enter:

```
vxrlink status rlink
```

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

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

Workaround:

There is no workaround for this issue.

Bunker replay does not occur with volume sets (3329970)

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

Workaround:

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

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.

Workaround:

There is no workaround at this point of time.

Attempt to use SmartTier commands fails (2332973)

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

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

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

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

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

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

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

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

- BALANCE
- CHECKPOINT
- METADATA

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

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

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

Workaround:

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

- For FlashSnap, resync the snapshot and try the clone operation again.
- For FileSnap and Database Storage Checkpoint, destroy the clone and create the clone again.
- For space-optimized snapshots, destroy the snapshot and create a new snapshot.

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

Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.2 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.2.

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

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

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

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

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

Workaround: Edit the PFILE to arrange the text so that the parameter values are on a single line. If the database uses a spfile and some parameter values are spread

across multiple lines, then use the Oracle commands to edit the parameter values such as they fit in a single line.

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

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

There is no workround at this point of time.

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

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

Workaround: There is no workaround for this issue.

vxdbd process is online after Flash archive installation (2869269)

After a Flash archive installation of the SF stack, the `vxdbd` process is up, even if the stack is not configured.

Workaround: You can ignore, or stop the `vxdbd` process using the `/opt/VRTSdbed/common/bin/vxdbdctrl stop` command.

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

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

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

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

As per the Reverse Resync design, the first RRBegin operation relocates the SFDB repository to a backup location, and the ReverseResyncAbort and ReverseResyncCommit operations restore it to the original location. When the

second RRBegin operation attempts to relocate the same repository which is already relocated, SFDB reports the error message.

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

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

The ReverseResyncBegin (RRBegin) operation 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 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.

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'
```

On Solaris 11.1 SPARC, setting up the user-authentication process using the sfae_auth_op command fails with an error message (3556996)

The debug logs display the missing `ps` utility as the 'ucb' package was absent in the default operating system installation. Due to which, the user-authentication process fails and the following error message is reported:

```
#/opt/VRTS/bin/sfae_auth_op -o setup
Setting up AT
Starting SFAE AT broker

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

Workaround: Install the `pkg:/compatibility/ucb` package such that the `ps` utility is available in `/usr/ucb/ps`.

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

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

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

Workaround: There is no workaround for this issue.

Cloning of a CDB fails for point-in-time copies when one of the PDBs is in the read-only mode (3513432)

For Oracle version 12.1.0.1, cloning a container database (CDB) fails if one of the pluggable databases (PDBs) is in the read-only mode. The failure occurs with the following error message:

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

Reason: ORA-00376: file 9 cannot be read at this time
ORA-01111: name for data file 9 is unknown - rename to correct file
ORA-01110: data file 9: '/ora_base/db_home/dbs/MISSING00009'...
```

Workaround: There is no workaround for this issue.

If a CDB has a tablespace in the read-only mode, then the cloning fails (3512370)

For Oracle version 12.1.0.1, when a container database (CDB) has a tablespace in the read-only mode for all point-in-time copies, cloning of that CDB fails with the following error message:

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

Reason: ORA-01122: database file 15 failed verification check
ORA-01110: data file 15: '/tmp/test1/data/sfaedb/newtbs1.dbf'
```

```
ORA-01202: wrong incarnation of this file - wrong creation time
...
```

Workaround: There is no workaround for this issue.

Virtualization known issues

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

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 70.

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
```

SmartSync is not supported for Oracle databases running on raw VxVM volumes

SmartSync is not supported for Oracle databases that are configured on raw volumes, because Oracle does not support the raw volume interface.

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.

When an I/O domain fails, the vxdisk scandisks or vxdctl enable command take a long time to complete (2791127)

When an I/O domain fails, the vxdisk scandisks or vxdctl enable from the Oracle VM Server for SPARC guest take a long time to complete. `vdc_ioctls` like `DKIOCGGEOM` and `DKIOCINFO` also take more time to return. These issues seem to be due to retry operations performed at the Solaris operating system layer.

Reducing the vdc_timeout value to lower value might help to bring down time. Dynamic multi-pathing (DMP) code is optimized to avoid making such `vdc_ioctl`

calls in an Oracle VM Server for SPARC guest environment as much possible. This change considerably reduces delays.

A complete resolution to this issue may require changes at the Solaris operating system level.

A 1 TB disk that is not labeled using operating system commands goes into an error state after the vxconfigd daemon is restarted

A 1 TB disk that is not labeled using operating system commands goes into an error state after the vxconfigd daemon is restarted.`vxconfigd` daemon is restarted.

Currently, a solution from the vendor is not available.

Veritas File System software limitations

The following are software limitations in this release of Veritas File System.

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.

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.

SmartIO software limitations

The following are the SmartIO software limitations in this release.

Cache is not online after a reboot

Generally, the SmartIO cache is automatically brought online after a reboot of the system.

If the SSD driver module is not loaded automatically after the reboot, you need to load the driver 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
```

The sfcache operations may display error messages in the caching log when the operation completed successfully (3611158)

The `sfcache` command calls other commands to perform the caching operations. If a command fails, additional commands may be called to complete the operation. For debugging purposes, the caching log includes all of the success messages and failure messages for the commands that are called.

If the `sfcache` command has completed successfully, you can safely ignore the error messages in the log file.

Documentation

Product guides are available in the PDF format on the software media in the `/docs/product_name` directory. Additional documentation is available online.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The latest product documentation is available on the Symantec website.

http://sort.symantec.com/documents

Documentation set

Each product in the Storage Foundation and High Availability Solutions product line includes release notes, an installation guide, and additional documents such as administration and agent guides. In most cases, you may also need to refer to the documentation for the product's components.

The SFHA Solutions documents describe functionality and solutions that apply across the product line. These documents are relevant whichever SFHA Solutions product you use.

Note: The GNOME PDF Viewer is unable to view Symantec documentation. You must use Adobe Acrobat to view the documentation.

Symantec Storage Foundation documentation

Table 1-11 lists the documentation for Symantec Storage Foundation.

Table 1-11 Symantec Storage Foundation documentation

Document title	File name	Description
Symantec Storage Foundation Release Notes	sf_notes_62_sol.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_62_sol.pdf	Provides information required to install the product.
Symantec Storage Foundation Administrator's Guide	sf_admin_62_sol.pdf	Provides information required for administering the product.
Symantec Storage Foundation: Storage and Availability Management for Oracle Databases	sfhas_oracle_admin_62_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.
Veritas File System Programmer's Reference Guide (This document is available online only.)	vxfs_ref_62_sol.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 1-12 lists the documentation for Symantec Storage Foundation and High Availability Solutions products.

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

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions—What's new in this release (This document is available online.)	sfhas_whats_new_62_unix.pdf	Provides information about the new features and enhancements in the release.
Symantec Storage Foundation and High Availability Solutions Getting Started Guide	getting_started.pdf	Provides a high-level overview of installing Symantec products using the script-based installer. The guide is useful for new users and returning users that want a quick refresher.
Symantec Storage Foundation and High Availability Solutions Solutions Guide	sfhas_solutions_62_sol.pdf	Provides information about how SFHA Solutions product components and features can be used individually and in concert to improve performance, resilience and ease of management for storage and applications.
Symantec Storage Foundation and High Availability Solutions Virtualization Guide (This document is available online.)	sfhas_virtualization_62_sol.pdf	Provides information about Symantec Storage Foundation and High Availability support for virtualization technologies. Review this entire document before you install virtualization software on systems running SFHA products.
Symantec Storage Foundation and High Availability Solutions SmartIO for Solid State Drives Solutions Guide	sfhas_smartio_solutions_62_sol.pdf	Provides information on using and administering SmartIO with SFHA solutions. Also includes troubleshooting and command reference sheet for SmartIO.
Symantec Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide (This document is available online.)	sfhas_dr_impl_62_sol.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_62_sol.pdf	Provides information on using Volume Replicator (VVR) for setting up an effective disaster recovery plan by maintaining a consistent copy of application data at one or more remote locations.

Table 1-12 Symantec Storage Foundation and High Availability Solutions products documentation *(continued)*

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions Troubleshooting Guide	sfhas_tshoot_62_sol.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.

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

https://sort.symantec.com/documents