

Veritas Storage Foundation™ Release Notes

Linux

5.1 Service Pack 1



Veritas 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: 5.1 SP1

Document version: 5.1SP1.5

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

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 Web site 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

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:

docs@symantec.com

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](#)
- [Component product release notes](#)
- [About Symantec Operations Readiness Tools](#)
- [Important release information](#)
- [Changes in version 5.1 SP1](#)
- [No longer supported](#)
- [System requirements](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation errata](#)
- [Documentation](#)

About this document

This document provides important information about Veritas Storage Foundation (Storage Foundation) version for Linux. Review this entire document before you install Storage Foundation.

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

This is Document version: 5.1SP1.5 of the *Veritas Storage Foundation Release Notes*. Before you start, ensure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

<http://www.symantec.com/business/support/overview.jsp?pid=15107>

Component product release notes

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

`/product_name/docs`

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

About Symantec Operations Readiness Tools

Symantec™ Operations Readiness Tools (SORT) is a set of Web-based tools and services that lets you proactively manage your Symantec enterprise products. SORT automates and simplifies administration tasks, so you can manage your data center more efficiently and get the most out of your Symantec products. SORT lets you do the following:

- Collect, analyze, and report on server configurations across UNIX or Windows environments. You can use this data to do the following:
 - Assess whether your systems are ready to install or upgrade Symantec enterprise products
 - Tune environmental parameters so you can increase performance, availability, and use
 - Analyze your current deployment and identify the Symantec products and licenses you are using
- Upload configuration data to the SORT Web site, so you can share information with coworkers, managers, and Symantec Technical Support
- Compare your configurations to one another or to a standard build, so you can determine if a configuration has "drifted"
- Search for and download the latest product patches
- Get notifications about the latest updates for:
 - Patches

- Hardware compatibility lists (HCLs)
- Array Support Libraries (ASLs)
- Array Policy Modules (APMs)
- High availability agents
- Determine whether your Symantec enterprise product configurations conform to best practices
- Search and browse the latest product documentation
- Look up error code descriptions and solutions

Note: Certain features of SORT are not available for all products.

To access SORT, go to:

<http://sort.symantec.com>

Important release information

- The latest product documentation is available on the Symantec Web site at: <http://www.symantec.com/business/support/overview.jsp?pid=15107>
- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:
- For the latest patches available for this release, go to: <http://sort.symantec.com/>

Changes in version 5.1 SP1

This section lists the changes for Veritas Storage Foundation.

Changes related to the installation

The product installer includes the following changes.

The VRTScutil and VRTSacclib RPMs are no longer in use

For all high availability products, the VRTScutil and VRTSacclib RPMs are no longer required.

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

Installer-related changes to configure LLT private links, detect aggregated links, and configure LLT over UDP

For all high availability products, the installer provides the following new features in this release to configure LLT private links during the Storage Foundation HA configuration:

- The installer detects and lists the aggregated links that you can choose to configure as private heartbeat links.
- The installer provides an option to detect NICs on each system and network links, and sets link priority to configure LLT over Ethernet.
- The installer provides an option to configure LLT over UDP.

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

Installer supports configuration of non-SCSI3 based fencing

You can now configure non-SCSI3 based fencing for VCS cluster using the installer.

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

The installer can copy CPI scripts to any given location using `-copyinstallscripts` option

The installer can copy CPI scripts to given location using `-copyinstallscripts` option. This option is used when customers install SFHA products manually and require CPI scripts stored on the system to perform product configuration, uninstallation, and licensing tasks without the product media.

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

Web-based installer supports configuring Storage Foundation HA cluster in secure mode

You can now configure the Storage Foundation HA cluster in secure mode using the Web-based installer.

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

The installer provides automated, password-less SSH configuration

When you use the installer, it enables SSH or RSH communication among nodes. It creates SSH keys and adds them to the authorization files. After a successful completion, the installer removes the keys and system names from the appropriate files.

When you use the installer for SSH communications, meet the following prerequisites:

- The SSH (or RSH) daemon must be running for auto-detection.
- You need the superuser passwords for the systems where you plan to install VCS.

The installer can check product versions

You can use the installer to identify the version (to the MP/RP/SP level depending on the product) on all platforms. Activate the version checker with `./installer -version system_name`.

Changes related to Veritas Storage Foundation

Veritas Storage Foundation includes the following changes:

Changes to Thin Provisioning and Thin Reclamation features

The following sections describe the changes related to Thin Provisioning and Thin Reclamation features.

SmartMove default changed

The default value of the system tunable `usefssmartmove` is now set to `all`. The change results in taking advantage of SmartMove feature during operations involving all types of disks – not just thin disks. It requires SmartMove feature support from VxFS. If required, you can change the default using the `vxdefault` command.

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

New initialization options for the `vxassist grow` command

The `vxassist grow` operation has new options for the initialization type. These changes align the initialization types for `vxassist grow` and `vxassist create` commands.

<code>init=sync</code>	If the volume has multiple plexes, VxVM synchronizes the data between the plexes during initialization.
<code>init=zero</code>	Initializes the volume or grown region, and initializes the associated data plexes to zeroes. If the volume resides on thin reclaimable LUNs, VxVM also reclaims the space within the storage array

<code>init=active</code>	Initializes the volume or grown region without modifying the existing data on the plexes.
<code>init=default</code>	Performs the default operation.

For more information, see the `vxassist(1m)` manual page.

Relayout operations on VxFS mounted volumes now use SmartMove

This is a performance related enhancement. Relayout operations on VxFS mounted volumes take advantage of its SmartMove capability. The change results in faster relayout of the volume.

Reclamation writes are not counted in write statistics

When you issue a reclamation command on a LUN, a disk group, or an enclosure, the request is passed down as writes to the Volume Manager from VXFS. This feature differentiates the writes generated by reclamation from the writes generated by normal application IO in the stats. By default, the reclamation writes are not shown with the `vxstat` command. To display the reclamation writes, use the command:

```
# vxstat -fm
```

Changes related to Veritas File System

Veritas File System includes the following changes:

Autolog replay on mount

The `mount` command automatically runs the VxFS `fsck` command to clean up the intent log if the `mount` command detects a dirty log in the file system. This functionality is only supported on file systems mounted on a Veritas Volume Manager (VxVM) volume.

Dynamic Storage Tiering is rebranded as SmartTier

In this release, the Dynamic Storage Tiering (DST) feature is rebranded as SmartTier.

FileSnap

FileSnaps provide an ability to snapshot objects that are smaller in granularity than a file system or a volume. The ability to snapshot parts of a file system name space is required for application-based or user-based management of data stored in a file system. This is useful when a file system is shared by a set of users or

applications or the data is classified into different levels of importance in the same file system.

See the *Veritas Storage Foundation Advanced Features Administrator's Guide*.

SmartTier sub-file movement

In this release, the Dynamic Storage Tiering (DST) feature is rebranded as SmartTier. With the SmartTier feature, you can now manage the placement of file objects as well as entire files on individual volumes.

See the *Veritas Storage Foundation Advanced Features Administrator's Guide* and the `fsppadm(1M)` manual page.

Tuning performance optimization of inode allocation

You can now set the `delicache_enable` tunable parameter, which specifies whether performance optimization of inode allocation and reuse during a new file creation is turned on or not.

See the *Veritas File System Administrator's Guide* and the `vxtunefs(1M)` manual page.

Veritas File System is more thin friendly

You can now tune Veritas File System (VxFS) to enable or disable thin-friendly allocations.

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes:

Veritas Volume Manager persisted attributes

The `vxassist` command now allows you to define a set of named volume allocation rules, which can be referenced in volume allocation requests. The `vxassist` command also allows you to record certain volume allocation attributes for a volume. These attributes are called persisted attributes. You can record the persisted attributes and use them in later allocation operations on the volume, such as growing the volume.

Automatic recovery of volumes during disk group import

After a disk group is imported, disabled volumes are enabled and started by default. To control the recovery behavior, use the `vxdefault` command to turn on or off the tunable `autostartvolumes`. If you turn off the automatic recovery, the recovery

behaves the same as in previous releases. This behavior is useful if you want to perform some maintenance after importing the disk group, and then start the volumes. To turn on the automatic recovery of volumes, specify `autostartvolume=on`.

After a disk group split, join, or move operation, Veritas Volume Manager (VxVM) enables and starts the volumes by default.

Cross-platform data sharing support for disks greater than 1 TB

Previous to this release, the `cdsdisk` format was supported only on disks up to 1 TB in size. Therefore, cross-platform disk sharing (CDS) was limited to disks of size up to 1 TB. Veritas Volume Manager (VxVM) 5.1 SP1 removes this restriction. VxVM 5.1 SP1 introduces CDS support for disks of size greater than 1 TB as well.

Note: The disk group version must be at least 160 to create and use the `cdsdisk` format on disks of size greater than 1 TB.

Default format for auto-configured disk has changed

By default, VxVM initializes all auto-configured disks with the `cdsdisk` format. To change the default format, use the `vxdiskadm` command to update the `/etc/default/vxdisk` file.

Changes related to Veritas Dynamic Multi-Pathing (DMP)

The following sections describe changes in this release related to DMP.

Veritas Dynamic Multi-Pathing (DMP) support for native logical volumes

In this release, Veritas Dynamic Multi-Pathing does not support Veritas File System (VxFS) on DMP devices.

DMP also supports creating single instance Oracle ASM or migrating an existing single instance of Oracle ASM onto DMP devices.

See the *Veritas Dynamic Multi-Pathing Administrator's Guide* for details.

Enhancements to DMP I/O retries

Veritas Dynamic Multi-Pathing (DMP) has a new tunable parameter, `dmp_lun_retry_timeout`. This tunable specifies a retry period for handling transient errors.

When all paths to a disk fail, there may be certain paths that have a temporary failure and are likely to be restored soon. If I/Os are not retried for a period of time, the I/Os may be failed to the application layer even though some paths are experiencing a transient failure. The DMP tunable `dmp_lun_retry_timeout` can be used for more robust handling of such transient errors by retrying the I/O for the specified period of time in spite of losing access to all the paths.

The DMP tunable `dmp_failed_io_threshold` has been deprecated.

See the `vxddmpadm(1m)` man page for more information.

Changes related to Veritas Volume Replicator

Veritas Volume Replicator includes the following changes:

vvrcheck configuration utility

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

Default network protocol is now TCP/IP

TCP/IP is now the default transport protocol for communicating between the Primary and Secondary sites. However, you have the option to set the protocol to UDP.

For information on setting the network protocol, see the *Veritas™ Volume Replicator Administrator's Guide*.

Checksum is disabled by default for the TCP/IP protocol

Beginning with Storage Foundation 5.1 with TCP as the default network protocol, VVR does not calculate the checksum for each data packet it replicates. VVR relies on the TCP checksum mechanism. However, if a node in a replicated data set is using a version of VVR earlier than 5.1 SP1PR4, VVR calculates the checksum regardless of the network protocol.

If you are using UDP/IP, checksum is enabled by default.

Improved replication performance in the presence of snapshots on the Secondary site

The effect of snapshots on the Secondary site is less drastic on replication performance.

No longer supported

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

- Bunker replication is not supported in a Cluster Volume Manager (CVM) environment.

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

Commands which are no longer supported as of version 5.1:

- ORAMAP (`libvxoramap`)
- Storage mapping commands `dbed_analyzer`, `vxstorage_stats`
- DBED providers (DBEDAgent), Java GUI, and `dbed_dbprocli`.
The SFDB tools features can only be accessed through the command line interface. However, Veritas Operations Manager (a separately licensed product) can display Oracle database information such as tablespaces, database to LUN mapping, and tablespace to LUN mapping.
- Storage statistics: commands `dbdst_makelbfs`, `vxdbts_fstatsummary`, `dbdst_fiostat_collector`, `vxdbts_get_datafile_stats`
- `dbed_saveconfig`, `dbed_checkconfig`
- `dbed_ckptplan`, `dbed_ckptpolicy`
- `dbed_scheduler`
- `sfua_rept_migrate` with `-r` and `-f` options

System requirements

This section describes the system requirements for this release.

Hardware compatibility list (HCL)

The hardware compatibility list contains information about supported hardware and is updated regularly. Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

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

<http://www.symantec.com/docs/TECH74012>

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

Veritas Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Number of nodes supported

For more updates on this support, see the Late-Breaking News TechNote on the Symantec Technical Support website:

Fixed issues

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

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

See “[Documentation](#)” on page 43.

Veritas Storage Foundation fixed issues

Veritas Storage Foundation: Issues fixed in 5.1 RP2

Table 1-1 Veritas Storage Foundation fixed issues in 5.1 RP2

Fixed issues	Description
2088355	dbed_ckptrollback fails for -F datafile option for 11gr2
2080633	Fixed the issue with vxdbd dumping core during system reboot.
1976928	dbed_clonedb of offline checkpoint fails with ORA-00600

Veritas Storage Foundation: Issues fixed in 5.1 RP1

Table 1-2 Veritas Storage Foundation fixed issues in 5.1 RP1

Fixed issues	Description
1974086	reverse_resync_begin fails after successfully unmount of clone database on same node when primary and secondary host names do not exactly match.
1940409, 471276	Enhanced support for cached ODM
1901367, 1902312	dbed_vmclonedb failed to umount on secondary server after a successful VM cloning in RAC when the primary SID string is part of the snapplan name.
1896097	5.1 GA Patch:dbed_vmclonedb -o recoverdb for offhost get failed
1873738, 1874926	dbed_vmchecksnap fails on standby database, if not all redologs from primary db are present.
1810711, 1874931	dbed_ymsnap reverse_resync_begin failed with server errors.

Storage Foundation for Databases (SFDB) tools fixed issues

This section describes the incidents that are fixed in Veritas Storage Foundation for Databases tools in this release.

Table 1-3 Veritas Storage Foundation for Databases tools fixed issues

Incident	Description
1873738	The dbed_vmchecksnap command may fail
1736516	Clone command fails for instant checkpoint on Logical Standby database
1810711	Flashsnap reverse resync command fails on offhost flashsnap cloning

Veritas File System fixed issues

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

Table 1-4 Veritas File System fixed issues

Incident	Description
2026603	Added quota support for the user "nobody".
2026625	The <code>sar -v</code> command now properly reports VxFS inode table overflows.
2050070	Fixed an issue in which the volume manager area was destroyed when spinlock was held.

Veritas File System: Issues fixed in 5.1 RP2

Table 1-5 Veritas File System fixed issues in 5.1 RP2

Fixed issues	Description
1995399	Fixed a panic due to null <code>i_fsext</code> pointer de-reference in <code>vx_inode</code> structure
2016373	Fixed a warning message V-3-26685 during freeze operation without nested mount points
2036841	Fixed a panic in <code>vx_set_tunefs</code>
2081441	Fixed an issue in <code>vxquota</code> regarding setting quota more than 1TB
2018481	Fixed an issue in <code>fsppadm(1M)</code> when volume did not have placement tags
2066175	Fixed panic in <code>vx_inode_mem_deinit</code>
2025155	Fixed an issue in <code>fsck(1m)</code> which was trying to free memory which was not allocated.
2043634	Fixed an issue in <code>quotas API</code>
1933844	Fixed a panic due to race condition in <code>vx_logbuf_clean()</code>
1960836	Fixed an issue in Thin Reclaim Operation
2026570	Fixed a hang issue in <code>vx_dopreamble ()</code> due to <code>ENOSPC</code> error.
2026622	Fixed a runqueue contention issue for <code>vx_worklists_thr</code> threads
2030889	Fixed a hang issue during <code>fsppadm(1m)</code> enforce operation with FCL
2036214	Fixed a core dump issue in <code>ncheck(1m)</code> in function <code>printname()</code> .
2076284	Optimized some VxMS api for contiguous extents.

Table 1-5 Veritas File System fixed issues in 5.1 RP2 (*continued*)

Fixed issues	Description
2085395	Fixed a hang issue in vxfsckd.
2059621	Fixed a panic due to null pointer de-reference in vx_unlockmap()
2016345	Fixed an error EINVAL issue with O_CREATE while creating more than 1 million files.
1976402	Fixed the issue in fsck replay where it used to double fault for 2TB luns.
1954692	Fixed a panic due to NULL pointer de-reference in vx_free()
2026599	Fixed a corruption issue when Direct IO write was used with buffered read.
2030773	Fixed issue with fsppadm(1m) where it used to generate core when an incorrectly formatted XML file was used.
2026524	Fixed a panic in vx_mkimtran()
2080413	Fixed an issue with storage quotas
2084071	Fixed an issue in fcladm(1m) where it used to generate core when no savefile was specified
2072165	Fixed an active level leak issue while fsadm resize operation.
1959374	Fixed a resize issue when IFDEV is corrupt
2098385	Fixed a performance issue related to 'nodatainlog' mount option.
2112358	Fixed an issue with file-system I/O statistics.

Veritas File System: Issues fixed in 5.1 RP1

Veritas Volume Manager fixed issues

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

Table 1-6 Veritas Volume Manager fixed issues

Incident	Description
150476	Add T for terabyte as a suffix for volume manager numbers

Table 1-6 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
248925	If vxdg import returns error, parse it
311664	vxconfigd/dmp hang due to a problem in the dmp_reconfig_update_cur_pri() function's logic
321733	Need test case to deport a disabled dg.
339282	Failed to create more than 256 config copies in one DG.
597517	Tunable to initialize EFI labeled >1tb PP devices.
1097258	vxconfigd hung when an array is disconnected.
1239188	Enhance vxprivutil to enable, disable, and display config+log copies state.
1301991	When vxconfigd is restarted with -k option, all log messages are sent to stdout. syslog should be the default location.
1321475	Join Failure Panic Loop on axe76 cluster.
1441406	'vxdisk -x list' displays wrong DGID.
1458792	After upgrade from SF5.0mp1 to SF5.0mp3, *unit_io and *pref_io was set to 32m.
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.
1485075	DMP sending I/O on an unopened path causing I/O to hang
1504466	VxVM: All partitions aren't created after failing original root disk and restoring from mirror.
1513385	VVR:Primary panic during autosync or dcm replay.
1528121	FMR: wrong volpagemod_max_memsz tunable value cause buffer overrun
1528160	An ioctl interrupted with EINTR causes frequent vxconfigd exits.
1586207	"vxsnap refresh" operations fail occasionally while data is replicating to secondary.
1589022	Infinite looping in DMP error handling code path because of CLARIION APM, leading to I/O hang.
1594928	Avoid unnecessary retries on error buffers when disk partition is nullified.
1662744	RVG offline hung due to I/Os pending in TCP layer

Table 1-6 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1665094	Snapshot refresh causing the snapshot plex to be detached.
1713670	'vxassist -g <dg-name> maxsize' doesn't report no free space when applicable
1715204	Failure of vxsnap operations leads to orphan snap object which cannot be removed.
1766452	vradmind dumps core during collection of memory stats.
1792795	Supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
1825270	I/O failure causes VCS resources to fault, as dmpnode get disabled when storage processors of array are rebooted in succession
1825516	Unable to initialize and use ramdisk for VxVM use.
1826088	After pulling out the Fibre Channel cables of a local site array, plex becomes DETACHED/ACTIVE.
1829337	Array firmware reversal led to disk failure and offlined all VCS resources
1831634	CVR: Sending incorrect sibling count causes replication hang, which can result in I/O hang.
1831969	VxVM: ddl log files are created with world write permission
1835139	I/Os hung after giveback of NetApp array filer
1840673	After adding new LUNs, one of the nodes in 3 node CFS cluster hangs
1846165	Data corruption seen on cdsdisks on Solaris-x86 in several customer cases
1857558	Need to ignore jeopardy notification from GAB for SFCFS/RAC, since oracle CRS takes care of fencing in this stack
1857729	CVM master in the VVR Primary cluster panicked when rebooting the slave during VVR testing
1860892	Cache Object corruption when replaying the CRECs during recovery
1869995	VVR: Improve Replication performance in presence of SO snapshots on secondary.

Table 1-6 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1874034	Race between modunload and an incoming IO leading to panic
1880279	Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment.
1881336	VVR: Primary node panicked due to race condition during replication
1884070	When running iotest on a volume, the primary node runs out of memory
1897007	vxesd coredumps on startup when the system is connected to a switch which has more than 64 ports
1899688	VVR: Every I/O on smartsync enabled volume under VVR leaks memory
1899943	CPS based fencing disks used along with CPS servers does not have coordinator flag set
1901827	vx dg move fails silently and drops disks.
1907796	Corrupted Blocks in Oracle after Dynamic LUN expansion and vxconfigd core dump
1915356	I/O stuck in vxvm causes a cluster node panic.
1933375	Tunable value of 'voliomem_chunk_size' is not aligned to page-size granularity
1933528	During Dynamic reconfiguration vxvm disk ends up in error state after replacing physical LUN.
1936611	vxconfigd core dump while splitting a diskgroup
1938907	WWN information is not displayed due to incorrect device information returned by HBA APIs
1946941	vxsnap print shows incorrect year
1954062	vxrecover results in os crash
1956777	CVR: Cluster reconfiguration in primary site caused master node to panic due to queue corruption
1969526	Panic in voldiodone when a hung priv region I/O comes back
1972848	vxconfigd dumps core during upgradation of VxVM

Table 1-6 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1974393	Cluster hangs when the transaction client times out
1982178	vxdiskadm option "6" should not list available devices outside of source diskgroup
1982715	vxclustadm dumps core during memory re-allocation.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout
1992872	vxresize fails after DLE.
1993953	CVM Node unable to join in Sun Cluster environment due to wrong coordinator selection
1998447	Vxconfigd dumps core due to incorrect handling of signal
1999004	I/Os hang in VxVM on linked-based snapshot
2002703	Misleading message while opening the write protected device.
2009439	CVR: Primary cluster node panicked due to queue corruption
2010426	Tag setting and removal do not handle wrong enclosure name
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
2016129	Tunable to disable OS event monitoring by vxesd
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1
2021737	vxdisk list shows HDS TrueCopy S-VOL read only devices in error state.
2025593	vxdg join hang/failure due to presence of non-allocator inforecords and when tagmeta=on
2027831	vxdg free not reporting free space correctly on CVM master. vxprint not printing DEVICE column for subdisks.
2029480	Diskgroup join failure renders source diskgroup into inconsistent state
2029735	System panic while trying to create snapshot
2034564	I/Os hung in serialization after one of the disks which formed the raid5 volume was pulled out
2036929	Renaming a volume with link object attached causes inconsistencies in the disk group configuration

Table 1-6 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2038137	System panics if volrdmirbreakup() is called recursively.
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.
2040150	Existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration
2052203	Master vold restart can lead to DG disabled and abort of pending transactions.
2052459	CFS mount failed on slave node due to registration failure on one of the paths
2055609	Allocation specifications not being propagated for DCO during a grow operation
2060785	Primary panics while creating primary rvg
2061066	vxisforeign command fails on internal cciss devices
2061758	Need documentation on list of test suites available to evaluate CDS code path and verification of the code path.
2063348	Improve/modify error message to indicate its thin_reclaim specific
2070531	Campus cluster: Couldn't enable site consistency on a dcl volume, when trying to make the disk group and its volumes siteconsistent.
2075801	VVR: "vxnetd stop/start" panicked the system due to bad free memory
2076700	VVR: Primary panic due to NULL pointer dereference
2094685	Diskgroup corruption following an import of a cloned BCV image of a SRDF-R2 device
2097320	Events generated by dmp_update_status() are not notified to vxconfigd in all places.
2105722	VVR: I/O hang on Primary with link-breakoff snapshot
2112568	System panics while attaching back two Campus Cluster sites due to incorrect DCO offset calculation
2122009	vxddladm list shows incorrect hba information after running vxconfigd -k
2126731	vxdisk -p list output is not consistent with previous versions

Table 1-6 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2131814	VVR: System panic due to corrupt sio in _VOLRPQ_REMOVE

Veritas Volume Manager: Issues fixed in 5.1 RP2

Veritas Volume Manager: Issues fixed in 5.1 RP1

Table 1-7 Veritas Volume Manager 5.1 RP1 fixed issues

Fixed issues	Description
1938484	EFI: Prevent multipathing don't work for EFI disk
1915356	I/O stuck in vxvm caused cluster node panic
1899688	[VVR] Every I/O on smartsync enabled volume under VVR leaks memory
1884070	When running iotest on volume, primary node runs out of memory
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1860892	Cache Object corruption when replaying the CRECs during recovery
1857729	CVM master in the VVR Primary cluster panic when rebooting the slave during VVR testing
1840673	After adding new luns one of the nodes in 3 node CFS cluster hangs
1835139	CERT : pnate test hang I/O greater than 200 seconds during the filer giveback
1826088	After pulling out FC cables of local site array, plex became DETACHED/ACTIVE
1792795	supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.

Known issues

This section covers the known issues in this release.

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

See “[Documentation](#)” on page 43.

Issues related to installation

This section describes the known issues during installation and upgrade.

Installation precheck can cause the installer to throw a license package warning (2320279)

If the installation precheck is attempted after another task completes (for example checking the description or requirements) the installer throws the license package warning. The warning reads:

```
VRTSvlic RPM not installed on system_name
```

Workaround:

The warning is due to a software error and can be safely ignored.

While configuring authentication passwords through the Veritas product installer, the double quote character is not accepted (1245237)

The Veritas product installer prompts you to configure authentication passwords when you configure Veritas Cluster Server (VCS) as a secure cluster, or when you configure Symantec Product Authentication Service (AT) in authentication broker (AB) mode. If you use the Veritas product installer to configure authentication passwords, the double quote character (") is not accepted. Even though this special character is accepted by authentication, the installer does not correctly pass the characters through to the nodes.

Workaround: There is no workaround for this issue. When entering authentication passwords, do not use the double quote character (").

EULA changes (2161557)

The locations for all EULAs have changed.

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

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

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

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

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

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

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

The VRTSacclib RPM is deprecated (2032052)

Veritas Storage Foundation known issues

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

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

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

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

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

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

```
DBI1070I Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

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

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

To communicate in a dual-stack environment

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

```
127.0.0.1 swlx20-v6
```

Or

```
127.0.0.1 swlx20-v6.punipv6.com
```

127.0.0.1 is the IPv4 loopback address.

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

Oracle 11gR1 may not work on pure IPv6 environment (1819585)

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

Workaround: There is no workaround for this, as Oracle 11gR1 does not fully support pure IPv6 environment. Oracle 11gR2 release may work on a pure IPv6 enviroment, but it has not been tested or released yet.

Not all the objects are visible in the SFM GUI (1821803)

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

Workaround:

To resolve this known issue

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

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

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

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

```
Cannot open file /etc/vx/vxdba/rac11g1/.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'rac11g1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdg ERROR V-5-1-582 Disk group SNAP_rac11dg1: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dg1.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dg1 failed.
```

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

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

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

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

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

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

To resolve this issue

- 1 In the Web GUI, go to **Settings > Deployment**.
- 2 Select **HF020008500-06 hotfix**.
- 3 Click **Install**.
- 4 Check the **force** option while reinstalling the hotfix.

A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880622)

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

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

Veritas Volume Manager known issues

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

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

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

Workaround:

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

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

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

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

The Veritas Volume Manager (VxVM) 5.1 SP1 includes several array names that differ from the array names in previous releases. Therefore, if you upgrade from

a previous release to VxVM 5.1 SP1, 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 5.1 SP1. Manually reconfigure the enclosure attributes to resolve the issue.

[Table 1-8](#) shows the Hitachi arrays that have new array names.

Table 1-8 Hitachi arrays with new array names

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

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. The persistence of the enclosure name is achieved with the `/etc/vx/array.info` file, which stores the mapping between cabinet serial number and array name. 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 5.1 SP1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

Veritas File System known issues

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

VxFS read ahead can cause stalled I/O on all write operations (1965647)

Changing the `read_ahead` parameter can lead to frozen I/O. Under heavy load, the system can take several minutes to recover from this state.

Workaround: There is no workaround for this issue.

Shrinking a file system that is larger than 1 TB takes a long time (2097673)

Shrinking a file system via either the `fsadm` command or `vxresize` command can take a long time to complete in some cases, such as if the shrink size is large and some large extent of a file is overlapping with the area to be shrunk.

Workaround: One possible workaround is to use the `vxtunefs` command and set `write_pref_io` and `write_nstream` to high values, such that `write_pref_io` multiplied by `write_nstream` is around 8 MB.

Storage Checkpoints can exceed the quota limit (2102201)

Under some circumstances, Storage Checkpoints can exceed the quota limit set by the `fsckptadm setquotalimit` command. This issue can arise if all of the following conditions are met:

- The Storage Checkpoint quota has been enabled.
- The Storage Checkpoint quota is not exceeded.
- A file content modification operation, including removing a file, needs to push some or all blocks of the file to the Storage Checkpoint.
- Number of blocks that need to be pushed to the Storage Checkpoint is enough to exceed Storage Checkpoint quota hard limit.

Workaround: There is no workaround for this issue.

vxfsconvert can only convert file systems that are less than 1 TB (2108929)

The `vxfsconvert` command can only convert file systems that are less than 1 TB. If the file system is greater than 1 TB, the `vxfsconvert` command fails with the "Out of Buffer cache" error.

Truncate operation of a file with a shared extent in the presence of a Storage Checkpoint containing FileSnaps results in an error (2149659)

This issue occurs when Storage Checkpoints are created in the presence of FileSnaps or space optimized copies, and one of the following conditions is also true:

- In certain cases, if a FileSnap is truncated in the presence of a Storage Checkpoint, the `i_nblocks` field of the inode, which tracks the total number of blocks used by the file, can be miscalculated, resulting in inode being marked bad on the disk.
- In certain cases, when more than one FileSnap is truncated simultaneously in the presence of a Storage Checkpoint, the file system can end up in a deadlock state.

This issue causes the following error to display:

```
f:xted_validate_cuttran:10 or f:vx_te_mklbtran:1b
```

Workaround: In the first case, run a full `fsck` to correct the inode. In the second case, restart the node that is mounting the file system that has this deadlock.

Tunable not enabling the lazy copy-on-write optimization for FileSnaps (2164568)

The lazy copy-on-write tunable does not enable the lazy copy-on-write optimization for FileSnaps.

Workaround: There is no workaround for this issue.

vxfilesnap fails to create the snapshot file when invoked with the following parameters: vxfilesnap source_file target_dir (2164744)

The `vxfilesnap` command fails to create the snapshot file when invoked with the following parameters:

```
# vxfilesnap source_file target_dir
```

Invoking the `vxfilesnap` command in this manner is supposed to create the snapshot with the same filename as the source file inside of the target directory.

Workaround: You must specify the source file name along with the target directory, as follows:

```
# vxfilesnap source_file target_dir/source_file
```

Panic due to null pointer de-reference in vx_unlockmap() (2059611)

A null pointer dereference in the `vx_unlockmap()` call can cause a panic. A fix for this issue will be released in a future patch.

Workaround: There is no workaround for this issue.

Veritas Volume Replicator known issues

This section describes the known issues in this release of Veritas Volume Replicator (VVR).

vradmin syncvol command compatibility with IPv6 addresses (2075307)

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

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

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

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

Interrupting the vradmin syncvol command may leave volumes open (2063307)

Interrupting the `vradmin syncvol` command may leave volumes on the Secondary site in an open state.

Workaround: On the Secondary site, restart the `in.vxrsyncd` daemon. Enter the following:

```
# /etc/init.d/vxrsyncd.sh stop
# /etc/init.d/vxrsyncd.sh start
```

Storage Foundation 5.0MP3 Rolling Patch 2 required for replication between 5.0 MP3 and 5.1 SP1 (1800600)

In order to replicate between Primary sites running Storage Foundation 5.0 MP3 and Secondary sites running Storage Foundation 5.1 SP1, or vice versa, you must install the Storage Foundation 5.0MP3 Rolling Patch 2 on the nodes using 5.0MP3. This patch resolves several outstanding issues for replicating between versions.

In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon

Issue: After upgrading VVR to an IPv6-only environment in 5.1 release, `vradmin` commands may not work when a colon is specified in the RVG, data volume(s) and/or SRL name. It is also possible that after upgrading VVR to an IPv6-only environment, `vradmin createpri` may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

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

While `vradmin changeip` is running, `vradmin` may temporarily lose heart beats (2162625)

This issue occurs when you use the `vradmin changeip` command to change the host name or IP address set in the Primary and Secondary RLINKs. While the `vradmin changeip` command runs, `vradmin` may temporarily lose heart beats, and the command terminates with an error message.

Workaround:

To resolve this issue

- 1 Depending on the application I/O workload, 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 `vradmin` to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following:

If using VEA to create a replicated data set fails, messages display corrupt strings in the Japanese locale (1726499, 1377599)

When using VEA to create a replicated data set, because the volumes do not have a DCM log on all nodes, the message window displays corrupt strings and unlocalized error messages.

Workaround: There is no workaround for this issue.

vxassist layout removes the DCM (2162522)

If you perform a layout 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:

```
# vxrvrg -g diskgroup stop rvrg
```
- 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 (2162537)

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

Workaround:

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

1 Pause or stop the applications.

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

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0.x to 5.1SP1 (2184482)

```
/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`.

Database fails over during Flashsnap operations (1469310)

In an Storage Foundation environment, if the database fails over during Flashsnap operations such as the `dbed_vmsnap -o resync` command and various error messages appear. This issue occurs because Flashsnap commands do not create a VCS resource for the SNAP disk group. As such, when the database fails over, only the primary disk group is moved to another node.

Workaround

There is no workaround for this issue.

The error messages depend on the timing of the database failover. To fix the problem, you need to bring the FlashSnap state to `SNAP_READY`. Depending on the failure, you may have to use base VxVM commands to reattach mirrors. After mirrors are attached, you need to wait until the mirrors are in `SNAPDONE` state. Re-validate the snapplan again.

Reattach command failure in a multiple disk group environment (1840672)

In a multiple disk group environment, if the snapshot operation fails then `dbed_vmsnap` fails to reattach all the volumes. This operation must be performed as root user.

Workaround

In case the reattach operation fails, use the following steps to reattach the volumes.

To reattach volumes in a multiple disk group environment if the snapshot operation fails

- 1 Join the snapshot disk groups to primary diskgroups. The snapshot disk group name is a concatenation of “SNAPSHOT_DG_PREFIX” parameter value in snapplan and primary disk group name. Use the following command to join the disk groups:

```
# vxdg join snapshot_disk_group_name  
           primary_disk_group_name
```

- 2 Start all the volumes in primary disk group.

```
# vxvol -g primary_disk_group_name startall
```

- 3 Reattach the snapshot volumes with primary volumes. The snapshot volume names is a concatenation of “SNAPSHOT_VOL_PREFIX” parameter value in snapplan and primary volume name. Use the following command to reattach the volumes.

```
# vxsnap -g primary_disk_group_name reattach snapshot_volume_name  
source=primary_volume_name
```

Repeat this step for all the volumes.

Clone command fails if archive entry is spread on multiple lines (1764885)

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

Workaround

There is no workaround for this issue.

VCS agent for Oracle: Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 (1985055)

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

Workaround: Set `MonitorOption` attribute for Oracle resource to 0.

Software limitations

This section covers the software limitations of this release.

See “[Documentation](#)” on page 43.

Veritas Storage Foundation software limitations

Veritas Volume Manager software limitations

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

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, set the following DMP tunables:

Table 1-9

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

The change is persistent across reboots.

To change the tunable parameters

- 1 Issue the following commands:

```
# vxdmpadm settune dmp_restore_interval=60
# vxdmpadm settune dmp_path_age=120
```

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

```
# vxdmpadm gettune dmp_restore_interval
# vxdmpadm gettune dmp_path_age
```

Veritas File System software limitations

The following are software limitations in the 5.1 SP1 release of Veritas Storage Foundation.

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.

Veritas Volume Replicator software limitations

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

Replication in a shared environment

Currently, replication support is limited to 4-node cluster applications.

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.
- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.
- A replication configuration in which an IPv6 address is specified for the `local_host` attribute of a primary RLINK and an IPv4 address is specified for the `remote_host` attribute of the same RLINK.
- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.
- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 5.1SP1 and the prior major releases of Storage Foundation (5.0 MP3 and 5.1). Replication between versions is supported for disk group versions 140, 150, and 160 only. Both the Primary and Secondary hosts must be using a supported disk group version.

Veritas Storage Foundation for Databases tools software limitations

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

Oracle Data Guard in an Oracle RAC environment

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

Upgrading if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 5.1SP1: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to SP1.

Documentation errata

The following sections, if present, cover additions or corrections for Document version: 5.1SP1.5 of the product documentation. These additions or corrections may be included in later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).

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

See [“Documentation”](#) on page 43.

See [“About Symantec Operations Readiness Tools”](#) on page 8.

Documentation

Product guides are available on the documentation disc in PDF formats. Symantec recommends copying pertinent information, such as installation guides and release notes, from the disc to your system's `/opt/VRTS/docs` directory for reference.

Documentation set

[Table 1-10](#) lists the documentation for Symantec Product Authentication Service (AT).

Table 1-10 Symantec Product Authentication Service documentation

Title	File name
<i>Symantec Product Authentication Service Release Notes</i>	vxat_notes.pdf

Table 1-10 Symantec Product Authentication Service documentation (*continued*)

Title	File name
<i>Symantec Product Authentication Service Administrator's Guide</i>	vxat_admin.pdf

Manual pages

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

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

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

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

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

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

See the `man(1)` manual page.