Veritas Storage Foundation™ Release Notes

AIX

5.1 Service Pack 1

Symantec™

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: 5.1SP1.6

Legal Notice

Technical Support

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

Symantec's support offerings include the following:

■ A range of support options that give you the flexibility to select the right amount of service for any size organization

■ Telephone and/or Web-based support that provides rapid response and up-to-the-minute information

■ Upgrade assurance that delivers software upgrades

■ Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis

■ Premium service offerings that include Account Management Services

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

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

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

Contacting Technical Support

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

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

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

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

■ Product release level

■ Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

docs@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
- 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 AIX. 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: 5.1SP1.6 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

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 in PDF formats:

*/product_name/*docs

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

This release includes the following component product release notes:

- *Veritas Storage Foundation Cluster File System Release Notes* (5.1 SP1)

- *Veritas Cluster Server Release Notes* (5.1 SP1)

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:
  http://entsupport.symantec.com/docs/334998

- 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.

Rolling upgrade support

To reduce downtime, the installer supports rolling upgrades. A rolling upgrade requires little or no downtime. A rolling upgrade has two main phases. In phase 1, the installer upgrades kernel packages on a subcluster. In phase 2, non-kernel packages are upgraded.

All high availability products support a rolling upgrade. You can perform a rolling upgrade from 5.1 or from any RPs to the current release.

You can perform a rolling upgrade using the script-based or Web-based installer.

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

The installsfha script and the uninstallsfha script are now available

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

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

The VRTScutil and VRTSacclib filesets are no longer in use

For all high availability products, the VRTScutil and VRTSacclib filesets 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*.

Web-based installer supports configuring disk-based fencing for Storage Foundation HA

You can now configure disk-based fencing for the Storage Foundation HA cluster 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*.

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

Changes related to Veritas Storage Foundation

Veritas Storage Foundation includes the following changes:

Supports Active Memory Sharing feature of IBM PowerVM

The Veritas Storage Foundation supports VIO clients that use memory from the Active Memory Sharing (AMS) pool.

See the *Veritas Storage Foundation and High Availability Solutions Virtualization Guide* for more information.

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 make` commands.

init=sync	If the volume has multiple plexes, Veritas Volume Manager (VxVM) synchronizes the data between the plexes during initialization.

init=zero	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.
init=active	Initializes the volume or grown region without modifying the existing data on the plexes.
init=default	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*.

Online migration of a native file system to VxFS file system

The online migration feature provides a method to migrate a native file system to the VxFS file system. The migration takes minimum amounts of clearly bounded, easy to schedule downtime. Online migration is not an in-place conversion and requires a separate storage. During online migration the application remains online and the native file system data is copied over to the VxFS 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:

Changes to DMP coexistence with native multi-pathing

The following limitations apply when using DMP with native multi-pathing:

- DMP does not display extended attributes for devices under the control of the native multi-pathing driver, MPIO. Extended attributes include the AVID, TP, TP_RECLAIM, SSD, RAID levels, snapshots, and hardware mirrors.

- If an array of any class other than Active/Active (A/A) is under the control of MPIO, then DMP claims the devices in A/A mode.
  DMP does not store path-specific attributes such as primary/secondary paths, port serial number, and the array controller ID.

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 attibutes. 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 to virtual SCSI device support

Storage Foundation 5.0MP1 introduced support for virtual SCSI (vSCSI) devices. This release provides enhancements to the vSCSI support.

The following enhancements are provided:

■ DMP handles multiple VIO servers as an Active/Standby configuration. The vSCSI ASL claims the devices in A/A-S mode instead of A/P mode.

■ DMP now recognizes enclosure names and AVID attributes. Previously, all vSCSI devices were grouped in a single vSCSI enclosure.

■ DMP now enforces that the I/O policy for vSCSI devices must be Single-Active. If you attempt to change the I/O policy, you will get an error.

■ DMP now recognizes extended device attribute information for vSCSI devices that are backed by DMP devices.

Veritas Volume Manager volumes can now be used as AIX paging devices

Starting with this release, you can use Veritas Volume Manager (VxVM) volumes as secondary swap devices in AIX. This functionality is supported starting with AIX 5.3 TL 10 and AIX 6.1 TL 3.

The native LVM paging commands now take a suffix of `-t vxvm` to work with VxVM volumes. To operate VxVM volumes as secondary paging devices, specify the `-t vxvm` option with the LVM commands such as `mkps`, `chps`, and `lsps`.

The volume name should not exceed 15 char limit. This restriction comes from IBM allowing only 15 chars in name field of CuAt class.

Only simple volumes (concat) are allowed to be used as page spaces. No other layouts are allowed. Therefore, the relayout operation is also not allowed.

The `lsps -a` command sometimes prints warning or error messages for VxVM devices. You can safely ignore these messages. IBM plans to fix this defect soon.

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 previous Veritas releases, DMP was only available as a feature of Veritas Volume Manager (VxVM). DMP supported VxVM volumes on DMP metadevices, and Veritas File System (VxFS) file systems on those volumes. This release extends DMP metadevices to support OS native logical volume managers (LVM). You can create LVM volumes and volume groups on DMP metadevices.

DMP supports LVM volume devices that are used as the paging devices.

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

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

DMP support on Virtual I/O Server (VIOS)

In this release, you can install DMP in the Virtual I/O server (VIOS). This enables the VIO server to export DMP devices as back-end devices for the VIO clients. The VIO clients access the DMP devices in the same way as any other vSCSI devices. DMP handles the I/O and multi-pathing for the vSCSI disks backed by the DMP devices.

Using DMP in the VIO server enables the DMP in the VIO Client to receive the extended attributes for the LUN. This enables the client LPAR to view back-end LUN attributes such as Thin, SSD, and RAID levels associated with the vSCSI devices in the VIO client.

Changes to DMP rootability on AIX

The upgrade of Storage Foundation with AIX DMP rootability enabled has been enhanced to allow the following upgrades with a single reboot of the system:

- Storage Foundation 5.1 to Storage Foundation 5.1 SP1 with AIX DMP rootability enabled.

- Operating system upgrade with AIX DMP rootability enabled on Storage Foundation 5.1 SP1.

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 vxdmpadm(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.

Changes related to Storage Foundation for Databases (SFDB) tools

New features in the Storage Foundation for Databases tools package for database storage management:

- Cached ODM support for clusters

- Cached ODM Manager support

- The Database Dynamic Storage Tiering (DBDST) feature is rebranded as SmartTier for Oracle and includes expanded functionality to support management of sub-file objects.

- Oracle 11gR2 support

New commands for 5.1 SP1:

- SmartTier for Oracle: commands added to support storage tiering of sub-file objects: `dbdst_obj_view`, `dbdst_obj_move`

- Cached ODM: command added to support Cached ODM Manager: `dbed_codm_adm`

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`
- `qio_convertdbfiles -f` option which is used to check for file fragmentation
- `dbed_scheduler`
- `sfua_rept_migrate` with `-r` and `-f` options

System requirements

This section describes the system requirements for this release.

Supported AIX operating systems

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

Product installation scripts verify the required update levels. The installation process terminates if the target systems do not meet the maintenance level requirements.

The minimum system requirements for this release are as follows:

For Power 6 or earlier processors at one of the following levels:

- AIX 6.1 TL2
- AIX 5.3 at one of the following levels:
  - TL7 with SP6 or later
  - TL8 with SP4 or later

For Power 7 processors at one of the following levels:

- AIX 6.1 TL5 with Service Pack 1 or later
- AIX Version 5.3 executing in POWER6 or POWER6+ compatibility at the following levels:
  - TL11 with Service Pack 2 or later
  - TL10 with Service Pack 4 or later

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

http://entsupport.symantec.com/docs/334998

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*.

Database requirements

Veritas Storage Foundations product features are supported for the following database environments:

Table 1-1

Veritas Storage Foundations feature	DB2	Oracle	Sybase
Oracle Disk Manager, Cached Oracle Disk Manager	No	Yes	No
Quick I/O, Cached Quick I/O	Yes	Yes	Yes
Concurrant I/O	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes
Database Storage Checkpoints	No	Yes	No
Database Flashsnap	No	Yes	No
SmartTier for Oracle	No	Yes	No

Storage Foundation for Databases (SFDB) tools Database Checkpoints, Database Flashsnap, and SmartTier for Oracle are supported only for Oracle database environments.

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

http://entsupport.symantec.com/docs/331625

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

Veritas Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Number of nodes supported

Storage Foundation is capable of supporting cluster configurations with up to 64 nodes. Symantec has tested and qualified configurations of up to 32 nodes at the time of the release.

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

http://entsupport.symantec.com/docs/334998

Fixed issues

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

This release includes fixed issues from the 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 release. For the list of fixed issues in the 5.1 SP1 RP2 release, see the Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 Release Notes.

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

Veritas Storage Foundation fixed issues

Veritas Storage Foundation: Issues fixed in 5.1 SP1

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

Table 1-2   Veritas Storage Foundation fixed issues

Incident	Description
2080659	5.1 and 5.1rp1 VRTSdbed.bff is not TCB (Trusted Computing Base) Compliant
1854459	db2exp may frequently dump core
1982674	dbed_clonedb of offline checkpoint fails with ORA-00600 with Oracle 11gR2 when ODM is enabled

Table 1-2   Veritas Storage Foundation fixed issues *(continued)*

Incident	Description
1959400	Dbed_ckptrollback fails for -F datafile option for Oracle database version 11gr2
1873738	The dbed_vmchecksnap command may fail
1736516	Clone command fails for instant checkpoint on Logical Standby database
1789290	dbed_vmclonedb -o recoverdb for offhost fails for Oracle 10gr2 and prior versions
1847850	dbed_vmchecksnap sometimes displays error message while validating snapplan
1810711	Flashsnap reverse resync command fails on offhost flashsnap cloning

Veritas Storage Foundation: Issues fixed in 5.1 RP2

Table 1-3   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.
2049381	Chkpoint Rollback failure likely due to incorrect NLS_LANG setting.
1976928	dbed_clonedb of offline checkpoint fails with ORA-00600

Veritas Storage Foundation: Issues fixed in 5.1 RP1

Table 1-4   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.

Table 1-4 Veritas Storage Foundation fixed issues in 5.1 RP1 *(continued)*

Fixed issues	Description
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_vmsnap 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 1-5 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
1789290	dbed_vmclonedb -o recoverdb for offhost fails for Oracle 10gr2 and prior versions
1847850	dbed_vmchecksnap sometimes displays error message while validating snapplan
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 1-6 Veritas File System fixed issues

Incident	Description
2026603	Added quota support for the user "nobody".
2026625	The `sar -v` command now properly reports VxFS inode table overflows.

Table 1-6   Veritas File System fixed issues *(continued)*

Incident	Description
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 1-7   Veritas File System fixed issues in 5.1 RP2

Fixed issues	Description
1995399	Fixed a panic due to null i_fsext pointer de-reference in vx_inode structure
2016373	Fixed a warning message V-3-26685 during freeze operation without nested mount points
2036841	Fixed a panic in vx_set_tunefs
2081441	Fixed an issue in vxedquota regarding setting quota more than 1TB
2018481	Fixed an issue in fsppadm(1M) when volume did not have placement tags
2066175	Fixed panic in vx_inode_mem_deinit
2025155	Fixed an issue in fsck(1m) which was trying to free memory which was not allocated.
2043634	Fixed an issue in quotas API
1933844	Fixed a panic due to race condition in vx_logbuf_clean()
1960836	Fixed an issue in Thin Reclaim Operation
2026570	Fixed a hang issue in vx_dopreamble () due to ENOSPC error.
2026622	Fixed a runqueue contention issue for vx_worklists_thr threads
2030889	Fixed a hang issue during fsppadm(1m) enforce operation with FCL
2036214	Fixed a core dump issue in ncheck(1m) in function printname().
2076284	Optimized some VxMS api for contiguous extents.
1995418	Fixed a panic in vx_timeout_free function.
2085395	Fixed a hang issue in vxfsckd.
2059621	Fixed a panic due to null pointer de-reference in vx_unlockmap()

Table 1-7 Veritas File System fixed issues in 5.1 RP2 *(continued)*

Fixed issues	Description
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

Table 1-8 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number)

Fixed issues	Description
1897458, 1805046	Fixed issue in alert generation from vxfs when file system usage threshold is set.
1933635, 1914625	Fixed issues in fs pattern assignment policy of the file system.
1933975, 1844833	Fixed VX_EBMAPMAX error during filesystem shrinking using fsadm..

Table 1-8   Veritas File System 5.1 RP1 fixed issues (listed incident
                       number/parent number) *(continued)*

Fixed issues	Description
1934085, 1871935	We now update ilist on secondary even if error received from primary for a VX_GETIAS_MSG is EIO.
1934095, 1838468	Fixed a race in qiostat update which was resulting in data page fault.
1934096, 1746491	Fix to avoid core dump while running fsvmap by initializing a local pointer.
1934098, 1860701	Moved drop of active level and reaquire to top of loop to stop resize from being locked out during clone removal.
1934107, 1891400	Fixed incorrect ACL inheritance issue by changing the way it cached permission data.
1947356, 1883938	Added utility mkdstfs to create DST policies.
1934094, 1846461	Fixed an issue with vxfsstat(1M) counters.

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 1-9   Veritas Volume Manager fixed issues

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

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

Incident	Description
1097258	vxconfigd hung when an array is disconnected.
1239188	Enhance vxprivutil to enable, disable, and display config+log copies state.
1301991	When vxconfig 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.
1407429	AIX: volslabd consuming memory.
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
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.

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

Incident	Description
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
1863155	Unable to boot system restored with NIM mksysb image
1869995	VVR: Improve Replication performance in presence of SO snapshots on secondary.
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

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

Incident	Description
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	vxdg 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
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.
1989057	Cannot restore root file system that was backed up by mksysb
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout

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

Incident	Description
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
1998686	vxconfigd memory allocation failure occurs in large configurations
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.
2022148	Hang on VxFS filesystem which are part of RVG.
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
2038137	System panics if volrdmirbreakup() is called recursively.
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.

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

Incident	Description
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
2065669	After upgrading to 5.1, reinitializaing the disk makes public region size smaller than
2067473	SF 5.1SP1 Beta - failure to register disk group with cluster.
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
2081515	vxresize command hund due to insufficient memory
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
2119531	In an NPIV client configuration with 2 controllers, only 1 is seen by VxVM.
2122009	vxddladm list shows incorrect hba information after running vxconfigd -k
2126731	vxdisk -p list output is not consistent with previous versions

Table 1-9 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

Table 1-10 Veritas Volume Manager 5.1 RP2 fixed issues

Fixed issues	Description
2126869	failed to enable DMP support for LVM bootability on extended rootvg
2126731	VxVM 5.1: vxdisk -p list output is not consistent with previous versions
2118721	vxconfigd core dumps after disabling and reenabling rootability
2113831	vxconfigd core dumps while including the previously excluded controller
2112568	System panics while attaching back two Campus Cluster sites due to incorrect DCO offset calculation
2081515	vxresize command hung due to insufficient memory
2078111	When the IOs are large and need to be split, DRL for linked volumes cause I/Os to hang
2065669	After upgrading to 5.1, reinitalizing the disk makes public region size smaller than the actual size.
2061066	vxisforeign command fails on internal cciss devices
2059046	FMR:TP: snap vol data gets corrupted if vxdisk reclaim is run while sync is in progress
2055609	Allocation specifications not being propagated for DCO during a grow operation
2053975	Snapback operation panicked the system
2052459	CFS mount failed on slave node due to registration failure on one of the paths
2040150	Existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.

Table 1-10   Veritas Volume Manager 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
2036929	renaming a volume with link object attached causes inconsistencies in the disk group configuration
2034564	I/Os hung in serialization after one of the disk which formed the raid5 volume was pulled out
2034104	Unable to initialize a disk using vxdiskadm
2031462	Node idle events are generated every second for idle paths controlled by Third Party drivers.
2029735	System panic while trying to create snapshot
2029480	Diskgroup join failure renders source diskgroup into inconsistent state
2025313	Storage keys support not getting enabled on Power 7
2021737	vxdisk list shows HDB TrueCopy S-VOL read only devices in error state.
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
2015570	File System read failure seen on space optimized snapshot after cache recovery
2012016	Slave node panics while vxrecovery is in progress on master
2011316	VVR: After rebooting 4 nodes and try recovering RVG will panic all the slave nodes.
2010426	Tag setting and removal do not handle wrong enclosure name
2006454	AxRT5.1P1: vxsnap prepare is displaying vague error message
2002703	Misleading message while opening the write protected device.
1999004	I/Os hang in VxVM on linked-based snapshot
1998606	vxconfigd memory allocation failure occurs in large configurations
1998447	Vxconfigd dumped core due to incorrect handling of signal
1992872	Vxresize fails after DLE.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout

Table 1-10   Veritas Volume Manager 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
1989662	/opt/VRTSsfmh/bin/vxlist causes panic.
1989057	Cannot restore root file system that was backed up by mksysb
1983768	IO hung on linked volumes while carrying out third mirror breakoff operation.
1982715	vxclustadm dumping core while memory re-allocation.
1974393	Avoiding cluster hang when the transaction client timed out
1972755	TP/ETERNUS:No reclaim seen with Stripe-Mirror volume.
1960341	Toggling of naming scheme is not properly updating the daname in the vxvm records.
1956777	CVR: Cluster reconfiguration in primary site caused master node to panic due to queue corruption
1952177	Machine panics after creating RVG
1946941	vxsnap print shows incorrect year
1946939	CVM: Panic during master takeover, when there are cache object I/Os being started on the new master
1946936	CVM: IO hangs during master takeover waiting for a cache object to quiesce
1942985	Improve locking mechanism while updating mediatype on vxvm objects
1936611	vxconfigd core dump while splitting a diskgroup
1933528	During Dynamic reconfiguration vxvm disk ends up in error state after replacing physical LUN.
1933375	Tunable value of 'voliomem_chunk_size' is not aligned to page-size granularity
1929083	Vxattachd fails to reattach site in absence of vxnotify events
1923906	CVM: Master should not initiate detaches while leaving the cluster due to complete storage failure
1920894	vxcheckhbaapi can loop forever
1920761	I/O hang observed after connecting the storage back to master node incase of local detach policy
1911546	Vxrecover hung with layered volumes

Table 1-10   Veritas Volume Manager 5.1 RP2 fixed issues *(continued)*

Fixed issues	Description
1899943	CPS based fencing disks used along with CPS servers does not have coordinator flag set
1897007	vxesd coredumps on startup when the system is connected to a switch which has more than 64 ports
1880279	Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment.
1874034	Race between modunload and an incoming IO leading to panic
1831969	VxVM: ddl log files are created with world write permission
1829337	Array firmware reversal led to disk failure and offlined all VCS resources
1665094	Snapshot refresh causing the snapshot plex to be detached.
1513385	VVR:Primary panic during autosync or dcm replay.
1485075	DMP sending I/O on an unopened path causing I/O to hang
1441406	'vxdisk -x list' displays wrong DGID
1097258	vxconfigd hung when an array is disconnected

Veritas Volume Manager: Issues fixed in 5.1 RP1

Table 1-11   Veritas Volume Manager 5.1 RP1 fixed issues

Fixed issues	Description
1948412	64 bit VxMS plugins of VxVM are missing in the path /opt/VRTSvxms/lib/map/aix64
1948167	vxdisk resize cannot handle over 1TB gpt labeled disk as expected
1940166	Problem while creating more than 30 shared DGs
1938484	EFI: Prevent multipathing don't work for EFI disk
1915356	I/O stuck in vxvm caused cluster node panic
1935332	NASGW:vxdisk updateudid marks the disk as clone_disk
1935230	Panic in voldsio_timeout() function.

Table 1-11 Veritas Volume Manager 5.1 RP1 fixed issues *(continued)*

Fixed issues	Description
1932091	Need for dmp_revive_paths() in dmp reconfiguration/restore_demon code path.
1907796	Corrupted Blocks in Oracle after Dynamic LUN expansion and vxconfigd core dump
1901827	vxdg move failed silently and drops disks.
1899688	[VVR] Every I/O on smartsync enabled volume under VVR leaks memory
1892387	VM commands getting hung on master node with 32-node cluster.
1892372	vxdisk reclaim should skip cache object update
1884070	When running iotest on volume, primary node runs out of memory
1881336	VVR: Primary Panic in vol_ru_replica_sent()
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
1852212	When vxesd is enabled, dmp/dr procedure(CLAB CCT test case #1805-phase 2) with PowerPath panics the system
1846165	Data corruption seen on cdsdisks on Solaris-x86 in several customer cases
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
1766452	VVR: VRAS: AIX: vradmind dumps core during collection of memory stats.
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 68.

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 fileset not installed on system_name
```

Workaround:

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

Reboot required when upgrading from 4.0x to 5.1 SP1 or 5.1 SP1 PR1 (2207134)

If you are upgrading the product from the version 4.0x to 5.1 SP1 or 5.1 SP1 PR1 then you need to upgrade the TL version to TL7 or later. A reboot is mandatory after upgrading the TL version.

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 (\").

Incorrect error messages: error: failed to stat, etc. (2120567)

During installation, you may receive errors such as, "error: failed to stat /net: No such file or directory." Ignore this message. You are most likely to see this message on a node that has a mount record of /net/x.x.x.x. The /net directory, however, is unavailable at the time of installation.

EULA changes (2161557)

The locations for all EULAs have changed.

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

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

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

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

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

NetBackup 6.5 or older version is installed on a VxFS file system. Before upgrading to Veritas Storage Foundation (SF) 5.1, the user umounts all VxFS file systems including the one which hosts NetBackup binaries (`/usr/openv`). While upgrading SF 5.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure packages `VRTSpbx`, `VRTSat`, and `VRTSicsco`, which causes NetBackup to stop working.

Workaround: Before you umount the VxFS file system which hosts NetBackup, copy the two files `/usr/openv/netbackup/bin/version` and `/usr/openv/netbackup/version` to the `/tmp` directory. After you umount the NetBackup file system, manually copy these two version files from the `/tmp` directory to their original path. If the path does not exist, make the same directory path with the command: `mkdir -p /usr/openv/netbackup/bin` and `mkdir -p /usr/openv/netbackup/bin`. Run the `installer` to finish the upgrade process. After upgrade process is done, remove the two version files and their directory paths.

How to recover from systems that are already affected by this issue: Manually install `VRTSpbx`, `VRTSat`, and `VRTSicsco` filesets after the upgrade process is done.

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

The installer displays the space that all the product filesets and patches needs. During migration some filesets are already installed and during migration some filesets 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 fileset is deprecated (2032052)

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

- Fresh installs: Do not install VRTSacclib.

- Upgrade: Uninstall old VRTSacclib and install new VRTSacclib.

- Uninstall: Ignore VRTSacclib.

The `-help` option for certain commands prints an erroneous argument list (2138046)

For installsf, installat, and the installdmp scripts , although the `-help` option prints the `-security`, `-fencing`, `-addnode` options as supported, they are in fact not supported. These options are only applicable for high availability products.

Web installation looks hung when `-tmppath` option is used (2160878)

If you select the `-tmppath` option on the first page of the webinstaller after installing or uninstalling is finished on the last page of webinstaller, when you click the **Finish** button, the webpage hangs. Despite the hang, the installation or the uninstallation finishes properly and you can safely close the page.

Ignore VRTSgms request to boot during installation (2143672)

During installation, you may see this error which you can ignore.

```
VRTSgms: old driver is still loaded...
VRTSgms: You must reboot the system after installation...
```

Veritas Storage Foundation known issues

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

The web installer gives warning of time mismatch when none exists (2172292)

In situations where network connectivity is very bad, you may encounter time
mismatch warnings when you use the web installer. You may get warningss such
as: nodes have difference in clock by more than 5 sec. This can happen even these
nodes keep the same time. You can ignore these warnings.

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 nebula`

Or

`127.0.0.1 nebula.nebuladomain.com`

`127.0.0.1` is the IPv4 loopback address.

`nebula` and `nebula.nebuladomain.com` are the IPv6 hostnames.

AT Server crashes when authenticating unixpwd user multiple times (1705860)

There is a known issue in the AIX kernel code that causes '`getgrent_r`' function to corrupt the heap. This issue is present in AIX 5.3 and AIX 6.1 Refer to IBM's Web site for more information:

http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52585

AT uses `getgrent_r` function to get the groups of the authenticated user.

IBM has released the fix as a patch to fileset bos.rte.libc. There are different patches available for different version of bos.rte.libc. You need to check the version of `bos.rte.libc` (For example: `lslpp -l | grep bos.rte.libc`) and apply the appropriate IBM patch:

■ For version 6.1.3.1:
http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52959
For the fix:
ftp://ftp.software.ibm.com/aix/efixes/iz52959/

■ For version 6.1.2.4:
http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52720
For the fix:
ftp://ftp.software.ibm.com/aix/efixes/iz52720/

■ For version 6.1.2.5 :
http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52975
For the fix:
ftp://ftp.software.ibm.com/aix/efixes/iz52975/

There are IBM patches for only certain version of `bos.rte.libc` that are available. If your system has a different `bos.rte.libc` version, you may have to upgrade to

a higher version where the fix is available. If your version is not available, you may have to contact IBM.

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

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

Running AIX 6.1, you may receive the following error message when using sqlplus:

```
$ sqlplus " / as sysdba"
SQL> startup nomount
SQL> ORA 0-0-0-0
```

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.

Upgrading operating system Technology Levels along with Storage Foundation using an alternate disk fails (2162945)

Upgrading the operating system Technology Levels (TL) along with Storage Foundation using an alternate disk fails occasionally with the following error:

```
alt_disk_copy: 0505-224 ATTENTION:
An error occurred during installation of
one or more software components.
Modifying ODM on cloned disk.
Building boot image on cloned disk.
forced unmount of /alt_inst/var/adm/ras/platform
forced unmount of /alt_inst/var
umount: error unmounting /dev/alt_hd2: Device busy
0505-144 alt_disk_install: Unable to unmount alt_inst filesystems.
```

No issues have been observed with Storage Foundation in the cause of the failure.

Veritas Volume Manager known issues

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

The cluster may hang if a node goes down (1835718)

The cluster may hang if a node goes down while one array is disabled or offline in a mirror=enclosure configuration.

This may occur, if a node panics or loses power while one array of a mirror=enclosure configuration is offline or disabled, then the cluster, fencing, I/O loads, and VxVM transactions hang.

Workaround: There is no workaround for this issue.

After installing Volume Manager, you may be prompted to reinstall it (1704161)

If you remove pre-5.1 Volume Manager packages and then install 5.1 Volume Manager without using the product installer, the following message is displayed:

```
The Volume Manager appears to be installed already. You should use
vxdiskadm to add more disks to the system. Installation with vxinstall
will attempt to reinstall the Volume Manager from the beginning.
Depending upon how your system is currently configured, a
reinstallation may fail and could leave your system unusable.

Are you sure you want to reinstall [y,n,q,?] (default: n)
```

Workaround

When you are prompted to reinstall, enter y.

Note: This message is not displayed if you install Volume Manager with the product installer.

vxconvert failures if PowerPath disks are formatted as simple disks (857504)

If a PowerPath disk is formatted as a simple disk (a foreign device), then the vxconvert utility may fail during conversion of LVM to VxVM. To view the format of the disk, use the vxdisk list command. This issue may also occur if the /etc/vx/darecs file contains an hdiskpower disk entry. This entry may be present if PowerPath disks were configured as foreign disks in Storage Foundation 4.0, and the entry was not changed after subsequent upgrades.

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***

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

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

Workaround:

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

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

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

```
Cannot setup space
```

Enabling tagmeta=on on a disk group causes delay in disk group split/join operations (2105547)

When vxdg set tagmeta=on is run on a diskgroup, multiple iterations of disk group split/join operations on the disk group causes huge delay in split/join operations.

Co-existence check might fail for CDS disks

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

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

If a path loses connectivity to the array, the path is marked 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.

Node is not able to join the cluster with high I/O load on the array with Veritas Cluster Server (2124595)

When the array has a high I/O load, the DMP database exchange between master node and joining node takes a longer time. This situation results in VCS resource online timeout, and then VCS stops the join operation.

Workaround:

Increase the online timeout value for the HA resource to 600 seconds. The default value is 300 seconds.

To set the OnlineTimeout attribute for the HA resource type CVMCluster

1   Make the VCS configuration to be read/write:

```
# haconf -makerw
```

2   Change the OnlineTimeout attribute value of CVMCluster:

```
# hatype -modify CVMCluster OnlineTimeout 600
```

3 Display the current value of OnlineTimeout attribute of CVMCluster:

```
# hatype -display CVMCluster -attribute OnlineTimeout
```

4 Save and close the VCS configuration:

```
# haconf -dump -makero
```

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 1-12 shows the Hitachi arrays that have new array names.

Table 1-12   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

DS4K series array limitations

In case of DS4K array series connected to AIX host(s), when all the paths to the storage are disconnected and reconnected back, the storage does not get discovered automatically. To discover the storage, run the cfgmgr OS command on all the affected hosts. After the cfgmgr command is run, the DMP restore daemon brings the paths back online automatically in the next path restore cycle. The time of next path restore cycle depends on the restore daemon interval specifed (in seconds) by the tunable dmp_restore_interval.

```
# vxdmpadm gettune dmp_restore_interval
        Tunable           Current Value  Default Value
----------------------- -------------  -------------
dmp_restore_interval         300            300
```

On DS4K array series connected to AIX host(s) DMP is supported in conjunction with RDAC. DMP is not supported on DS4K series arrays connected to AIX hosts In MPIO environment.

vxconfigd hang with path removal operation while IO is in-progress (1932829)

In AIX with HBA firmware version SF240_320, vxdisk scandisks (device discovery) takes a long time when a path is disabled from the switch or from the array.

Workaround:

To resolve this issue, upgrade the HBA firmware version to SF240_382.

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.

Asynchronous cached ODM requests do not use the cache (2010139)

Asynchronous cached ODM requests do not use the cache on AIX, and you might observe some performance degradation only during async cached ODM request. However, synchronous reads will not be affected.

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 shrink 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.

Online migration can miss copying some files and directories if you issue the umount command during a migration (2097960)

When the `umount` command is issued on a file system, the namespace momentarily disappers even before the `umount` command executes.This is true even for a busy `umount`. Because of this, if you run the `umount` command on a file system while an online migration is in progress, the background copy process can miss copying some files and directories. Files and directories that are being copied in the window when the namespace disappears are skipped. The background copy will not be able to detect this because calls such as `opendir`() fail with the ENOENT error at that time. This results in giving a false migration completion.

Workaround: Do not unmount a file system that is being online migrated.

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.

When online migration is in progress, df command with no mount point or device argument fails with error 1 (2162822)

When online migration is in progress, the `df` command with no mount point or device argument fails with error 1.

The output of the `df` command also shows the hidden mounts of the file systems that are undergoing migration, with dashes for all the `statfs` values.

Workaround: To avoid the error, specify a mount point or device.

Tunable not enabling the lazy copy-on-write optimization for FileSnaps (2164580)

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
```

Enabling the D_REFUND parameter on AIX 6.1 causes a hang in some situations on a cluster file system (2166515)

In some situations, enabling the D_REFUND parameter on AIX 6.1 causes a hang on a cluster file system. Some example situations include creating a Storage Checkpoint, unmounting a file system after receiving an I/O error, and having a high GLM load.

Workaround: Disable the D_REFUND parameter.

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.

The dynamic vmm buffer allocation feature requires certain AIX APARs to be installed (1849083)

VxFS supports the use of the dynamic vmm buffer allocation (D_REFUND) feature, which IBM added to AIX 6.1 TL2 and later releases of AIX. However, IBM fixed some issues in the D_REFUND feature through certain APARs, which you must install to use the D_REFUND feature with VxFS. The TL of the operating system determines which APAR you must install:

Operating system	Required APAR
AIX 6.1 TL2	IZ41494, which is packaged in SP3
AIX 6.1 TL3	IZ37627
AIX 6.1 TL4	IZ38189

Possible write performance degradation with VxFS local mounts (1837394)

Some applications that allocate large files without explicit preallocation may exhibit reduced performance with the VxFS 5.1 release and later releases compared to the VxFS 5.0 MP3 release due to a change in the default setting for the tunable `max_seqio_extent_size`. One such application is DB2. Hosting DB2 data on a single file system extent maximizes the potential for sequential pre-fetch processing. When DB2 detects an application performing sequential reads against database data, DB2 begins to read ahead and pre-stage data in cache using efficient sequential physical I/Os. If a file contains many extents, then pre-fetch processing is continually interrupted, nullifying the benefits. A larger `max_seqio_extent_size` value reduces the number of extents for DB2 data when adding a data file into a tablespace without explicit preallocation.

The `max_seqio_extent_size` tunable controls the amount of space that VxFS automatically preallocates to files that are allocated by sequential writes. Prior to the 5.0 MP3 release, the default setting for this tunable was 2048 file system blocks. In the 5.0 MP3 release, the default was changed to the number of file system blocks equaling 1 GB. In the 5.1 release, the default value was restored to the original 2048 blocks.

The default value of `max_seqio_extent_size` was increased in 5.0 MP3 to increase the chance that VxFS will allocate the space for large files contiguously, which tends to reduce fragmentation and increase application performance. There are two separate benefits to having a larger `max_seqio_extent_size` value:

■ Initial allocation of the file is faster, since VxFS can allocate the file in larger chunks, which is more efficient.

- Later application access to the file is also faster, since accessing less fragmented files is also more efficient.

In the 5.1 release, the default value was changed back to its earlier setting because the larger 5.0 MP3 value can lead to applications experiencing "no space left on device" (ENOSPC) errors if the file system is close to being full and all remaining space is preallocated to files. VxFS attempts to reclaim any unused preallocated space if the space is needed to satisfy other allocation requests, but the current implementation can fail to reclaim such space in some situations.

Workaround: If your workload has lower performance with the VxFS 5.1 release and you believe that the above change could be the reason, you can use the `vxtunefs` command to increase this tunable to see if performance improves.

To restore the benefits of the higher tunable value

1 Increase the tunable back to the 5.0 MP3 value, which is 1 GB divided by the file system block size.

   Increasing this tunable also increases the chance that an application may get a spurious ENOSPC error as described above, so change this tunable only for file systems that have plenty of free space.

2 Shut down any applications that are accessing any large files that were created using the smaller tunable setting.

3 Copy those large files to new files, which will be allocated using the higher tunable setting.

4 Rename the new files back to the original names.

5 Restart any applications that were shut down earlier.

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.

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

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

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

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

Workaround:

To resolve this issue

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

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

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

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

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

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

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname could not be
imported on bunker host hostname. Operation failed with error 256
```

```
and message VxVM VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote
server unreachable... Timestamp VCS ERROR V-16-2-13066 (hostname)
Agent is calling clean for resource(RVGPrimary) because the resource
is not up even after online completed.
```

Workaround:

To resolve this issue

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

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

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 -V 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 -V vxfs /dev/vx/dsk/dg/data_volume
```

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

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

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

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

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

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

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

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

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

While vradmin changeip is running, vradmind 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, `vradmind` 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 `vradmind` to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following:

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

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 relayout removes the DCM (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

1   Pause or stop the applications.

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

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

3   Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7   Start the RVG. Enter the following:

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

8   Resume or start the applications.

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.

VVR primary node panics while replicating in lossy and high latency network (2445066)

In a lossy and a high latency network, the connection between the VVR primary and secondary can get closed and re-established frequently because of heartbeat timeouts or data acknowledgement timeouts. When a connection is closed, the messages waiting for the acknowledgements are cleared. Hence the acknowledgments received for these messages over the new connection should be dropped. Processing the stale messages can lead to a system panic.

Workaround:

This issue is fixed in Storage Foundation 6.0 and later. The fix is not back-ported to older releases as it involves changes in the VVR protocol, which can only be applied to major releases. Customers can upgrade to Storage Foundation 6.0 or later releases for the fix.

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)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 or 5.0MP3 to SFHA or SF for Oracle RAC 5.1SP1. The error message is:

When upgrading from Storage Foundation version 5.0 or 5.0MP3 to Storage Foundation 5.1 SP1 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.

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, ues 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** *snapshop_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** *snapshop_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 68.

Veritas Storage Foundation software limitations

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

Upgrades on alternate disk supported only from version 5.1

Storage Foundation supports upgrade on an alternate disk only from version 5.1 to version 5.1 SP1. If you are running earlier versions of Storage Foundation, perform a full or phased upgrade to version 5.1 and then upgrade to version 5.1 SP1 using an alternate disk.

Veritas Volume Manager software limitations

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

Limitation with device renaming on AIX 6.1TL6

If you rename an operating system (OS) path with the `rendev` command on AIX 6.1TL6, the operation might remove the paths from DMP control. DMP cannot discover these paths.

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, change the default values for the DMP tunable parameters.

Table 1-13 describes the DMP tunable parameters and the new values.

Table 1-13   DMP settings for NetApp storage attached environment

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

The change is persistent across reboots.

To change the tunable parameters

1   Issue the following commands:

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

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

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

DMP support in AIX virtualization environment (2038475)

A single enclosure cannot have both NPIV and vSCSI LUNs. Each enclosure can have either vSCSI or NPIV LUNs. DMP does not support a mixed configuration.

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: 5.1SP1.6 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.

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

lists the documentation for Veritas Storage Foundation.

Table 1-14   Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_51sp1_aix.pdf
Veritas Storage Foundation and High Availability Installation Guide	sf_install_51sp1_aix.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sf_adv_ora_51sp1_aix.pdf
Veritas Storage Foundation Advanced Features Administrator's Guide	sf_adv_admin_51sp1_aix.pdf

lists the documentation for Veritas Volume Manager and Veritas File System.

Table 1-15   Veritas Volume Manager and Veritas File System documentation

Document title	File name
Veritas Volume Manager Administrator's Guide	vxvm_admin_51sp1_aix.pdf
Veritas Volume Manager Troubleshooting Guide	vxvm_tshoot_51sp1_aix.pdf
Veritas File System Administrator's Guide	vxfs_admin_51sp1_aix.pdf
Veritas File System Programmer's Reference Guide	vxfs_ref_51sp1_aix.pdf

Table 1-16 lists the documentation for Veritas Volume Replicator.

Table 1-16   Veritas Volume Replicator documentation

Document title	File name
Veritas Volume Replicator Administrator's Guide	vvr_admin_51sp1_aix.pdf
Veritas Volume Replicator Planning and Tuning Guide	vvr_planning_51sp1_aix.pdf
Veritas Volume Replicator Advisor User's Guide	vvr_advisor_users_51sp1_aix.pdf

Table 1-17 lists the documentation for Symantec Product Authentication Service (AT).

Table 1-17   Symantec Product Authentication Service documentation

Title	File name
Symantec Product Authentication Service Release Notes	vxat_notes.pdf
Symantec Product Authentication Service Administrator's Guide	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.