Veritas Storage Foundation™ Cluster File System Release Notes

HP-UX 11i v3

5.1 Service Pack 1

Symantec

Veritas Storage Foundation™ Cluster File System 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: 5.1SP1.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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 Web site.

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

About Symantec Connect

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

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

Storage Foundation Cluster File System Release Notes

This document includes the following topics:

About this document

This document provides important information about Veritas Storage Foundation Cluster File System (SFCFS) version 5.1 SP1 for HP-UX 11i v3. Review this entire document before you install or upgrade SFCFS.

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

This is Document version: 5.1SP1.0 of the *Veritas Storage Foundation Cluster File System 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 Release Notes* (5.1 SP1)
- *Veritas Cluster Server Release Notes* (5.1 SP1)

For information regarding software features, limitations, fixed issues, and known issues of component products:

- Veritas Cluster Server (VCS)
  See *Veritas Cluster Server Release Notes (5.1 SP1)*.
- Storage Foundation (SF)
  See *Veritas Storage Foundation 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://www.symantec.com/docs/TECH144835
- For the latest patches available for this release, go to: http://sort.symantec.com/

Changes introduced in 5.1 SP1

This section lists the changes for Veritas Storage Foundation Cluster File System.

Installation and upgrade

SFCFS installation and upgrade includes the following changes in 5.1 SP1:

Install options for Storage Foundation High Availability and the Storage Foundation Cluster File System High Availability

The product installation programs now prompt you for whether you want to install the high availability packages when you are installing Storage Foundation or Storage Foundation Cluster File System. This change enables you to explicitly choose which functionality is installed. In previous releases, the installed license key determined what functionality was installed.

The product installer displays Storage Foundation High Availability and Storage Foundation Cluster File System High Availability on the product selection menu.

Installation simulator

The product installer includes the `-makeresponsefile` option to simulate installing the selected Veritas product. The simulation option steps through the installation script, including all of the pre-installation checks on the systems. However, the simulation does not actually install the packages, uninstall previously installed packages, or start or stop any processes.

The simulation process enables you to create a response file, that can be used as a template for performing an installation. You can also use the simulator to view the installation questions and check the installation requirements without disrupting your existing installation.

To use the installation simulator, specify the `-makeresponsefile` argument to the installer or product installation script at the command line.

You can also use the `-makeresponsefile` option to simulate an uninstallation or product configuration.

For more information about response files or the installation simulator, see the *Installation Guide* for your Veritas product.

Improved response file generation

You can now create response files without performing a live installation, using the installation simulator.

Option to install only the minimal packages

The product installer now provides several options for which packages to install. For each product, you can install the minimal packages, the recommended packages or all of the packages.

When you install with the product installer, you can select from one of the following install options:

- Minimal depots: installs only the basic functionality for the selected product.
- Recommended depots: installs the full feature set without optional depots.
- All depots: installs all available depots.

When you install with a product install script, such as `installsf`, you can use the following options to the install script:

- -minpkgs: displays the depots and patches required for basic functionality.
- -recpkgs: displays the recommended depots and patches.
- -allpkgs: displays all available depots and patches.

Veritas extension for Oracle Disk Manager depot is installed by default for Storage Foundation and Storage Foundation Cluster File System

The Veritas extension for Oracle Disk Manager package is supported with a Storage Foundation license.

In this release, the product installer now installs the required depot for ODM by default when Storage Foundation is installed.

Changes related to Storage Foundation Cluster File System

Storage Foundation Cluster File System includes the following changes in 5.1 SP1:

Clustered NFS Support

This new Clustered NFS (CNFS) feature is expected to gracefully handle failure of any node and reclaim the advisory locks taken by NFS clients in such a way as to not accidentally lose any existing lock grants without notification.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide* for more information.

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

Common Internet File System

This new Common Internet File System (CIFS) feature lets you share CFS file systems using CIFS protocol that can be accessed by Window clients. Upon node failure or service group failover, the CIFS shares continue to be served by other cluster nodes.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide* for more information.

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

Cluster File System agents and Asynchronous Monitoring Framework support

The Cluster File System (CFS) agents (CFSMount and CFSfsckd) are Asynchronous Monitoring Framework (AMF) aware.

See the *Veritas Storage Foundation Cluster File System Installation Guide* for more information.

CVMVolDg agent changes

This section describes the changes in the CVMVolDg agent.

Support for importing shared disk groups

The CVMVolDg agent now imports the shared disk group from the CVM master node, if the disk group is not already imported, when the corresponding CVMVolDg resource is brought online.

Support for deporting shared disk groups

When the last online CVMVolDg resource for a shared disk group is taken offline, the CVMVolDg agent now deports the disk group if the `CVMDeportOnOffline` attribute is set to 1.

Review the following notes before setting the attribute value:

- If multiple CVMVolDg resources are configured for a shared disk group, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources. The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the `CVMDeportOnOffline` attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.
- The shared disk group is not deported if it contains open volumes.

Support for I/O polling on volume sets

You can enable the CVMVolDg agent to perform periodic I/O polling on volume sets by specifying their names in the `CVMVolumeIoTest` attribute of the resource. This enables the CVMVolDg agent to proactively check the availability of the volume sets by reading 4 KB blocks from its component volumes every monitor cycle. Errors, if any, are reported to the log file `/var/VRTSvcs/log/engine_A.log`.

Note: The CVMVolDg agent takes a volume set offline if the file system metadata volume in a volume set is discovered to be offline in a monitor cycle. However, if the CFSMount resource goes offline and the file system on the volume set is unmounted, the agent retains the online state of the volume set even if the metadata volume in the volume set is offline. This is because the CVMVolDg agent is unable to determine whether or not the volumes that are offline are metadata volumes.

New attribute `CVMDeportOnOffline`

The `CVMDeportOnOffline` attribute setting enables the CVMVolDg agent to determine whether or not a shared disk group must be deported when the corresponding CVMVolDg resource is taken offline. Set the value of this attribute to 1 if you want the agent to deport the disk group when the CVMVolDg resource is taken offline. The default value is set to 0.

You can set the attribute by running the following command:

```
# haconf -makerw
# hares -modify cvmvoldg_res CVMDeportOnOffline 1
# haconf -dump -makero
```

Verify the value of the attribute:

```
# hares -display cvmvoldg_res | grep CVMDeportOnOffline
```

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 5.1 SP1:

Issuing Cluster Volume Manager (CVM) commands from the slave node

In previous releases, Cluster Volume Manager (CVM) required that you issue configuration commands for shared disk groups from the master node of the cluster. Configuration commands change the object configuration of a CVM shared disk group. Examples of configuration changes include creating disk groups,

importing disk groups, deporting disk groups, and creating volumes. In this release, you can issue commands from any node, even when the command changes the configuration of the shared disk group. You do not need to know which node is the master to issue the command. If you issue the command on the slave node, CVM ships the commands from the slave node to the master node. CVM then executes the command on the master node.

Note the following limitations for issuing CVM commands from the slave node:

- The CVM protocol version must be at least 100.
- CVM does not support executing all commands on the slave node. You must issue the following commands only on the master node:
  - Commands that specify a controller name. For example:

    ```
    # vxassist -g shareddg make sharedvol 20M ctlr:fscsi0
    ```

  - Commands that specify both a shared disk group and a private disk group. For example:

    ```
    # vxdg destroy privatedg shareddg
    ```

  - Commands that include the defaults file as an argument. For example:

    ```
    # vxassist -d defaults_file
    ```

  - Veritas Volume Replicator (VVR) commands including `vxibc`, `vxrlink`, `vxrsync`, `vxrvg`, `vrport`, `vrstat`, and `vradmin`.
  - The `vxdisk` command.

Changing the CVM master online

Cluster Volume Manager (CVM) now supports changing the CVM master from one node in the cluster to another node, while the cluster is online. CVM migrates the master node, and reconfigures the cluster.

Symantec recommends that you switch the master when the cluster is not handling VxVM configuration changes or cluster reconfiguration operations. In most cases, CVM aborts the operation to change the master, if CVM detects that any configuration changes are occurring in the VxVM or the cluster. After the master change operation starts reconfiguring the cluster, other commands that require configuration changes will fail.

To change the master online, the cluster must be cluster protocol version 100 or greater.

Changes related to Veritas Volume Replicator

Veritas Volume Replicator includes the following changes in 5.1 SP1:

Default network protocol is now TCP/IP

The default network protocol has changed from UDP/IP to TCP/IP.

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 SP1, if you specify TCP as your protocol, then by default, 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 SP1, 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.

SmartMove for VVR

The SmartMove for VVR feature enables VVR to leverage information from VxFS knowledge of the file system blocks in use to optimize the time and network bandwidth required for initial resync of replicated volumes.

See the *Veritas Volume Replicator Administrator's Guide* for more information on SmartMove for VVR.

Veritas Volume Replicator supports IPv6

Veritas Volume Replicator supports IPv6 in this release. IPv6 is supported only with disk group version 150 or later.

The Internet Protocol version 6 (IPv6) is the next-generation Internet Layer protocol for packet-switched networks and the Internet. IPv4 is the first version of the protocol to be widely deployed. IPv6 has a much larger address space than IPv4. This results from the use of a 128-bit address, whereas IPv4 uses only 32 bits. This expansion provides flexibility in allocating addresses and routing traffic and eliminates the primary need for network address translation (NAT). IPv6 also

implemented new features that simplify aspects of address assignment and network renumbering when changing Internet connectivity providers. Network security is integrated into the design of the IPv6 architecture.

See "IPv6 software limitations" on page 52.

See the *Veritas Volume Replicator Administrator's Guide* for more information on VVR IP terminology.

See the *Veritas Storage Foundation Installation Guide* for more information on planning and upgrading VVR from a pervious version of IPv4 to IPv6.

Planning and upgrading VVR to use IPv6 as connection protocol

Storage Foundation High Availability supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

- VVR continues to support replication between IPv4-only nodes with IPv4 as the internet protocol
- VVR supports replication between IPv4-only nodes and IPv4/IPv6 dual-stack nodes with IPv4 as the internet protocol
- VVR supports replication between IPv6-only nodes and IPv4/IPv6 dual-stack nodes with IPv6 as the internet protocol
- VVR supports replication between IPv6 only nodes
- VVR supports replication to one or more IPv6 only nodes and one or more IPv4 only nodes from a IPv4/IPv6 dual-stack node
- VVR supports replication of a shared disk group only when all the nodes in the cluster that share the disk group are at IPv4 or IPv6

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.*

Partitioned directories

Normally, a large volume of parallel threads performing access and updates on a directory that commonly exist in an file system suffers from exponentially longer wait times for the threads. This feature creates partitioned directories to improve the directory performance of file systems. When any directory crosses the tunable threshold, this feature takes an exclusive lock on the directory inode and redistributes the entries into various respective hash directories. These hash directories are not visible in the name-space view of the user or operating system. For every new create, delete, or lookup thread, this feature performs a lookup for the respective hashed directory (depending on the target name) and performs the operation in that directory. This leaves the parent directory inode and its other hash directories unobstructed for access, which vastly improves file system performance.

See the *Veritas File System 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 Storage Foundation for Databases (SFDB) tools

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

- Storage Foundation for Cluster File (HA) System support
- Cached Oracle Disk Manager (ODM) support, including Cached ODM for clusters
- Cached ODM Manager support
- Multiple disk group support for FlashSnap
- Sub-file storage tiering is supported for SmartTier for Oracle (previously known as Database Dynamic Storage Tiering)
- SQLite repository
- Oracle Enterprise Manager (OEM) Plugin
- Oracle 11gR2 support
- Oracle Dataguard support
- HP-UX Service Guard support for replicated environments

Commands which have changed:

- `sfua_db_config` functionality is changed: this command is no longer needed to create a SFDB repository. The functionality of `sfua_db_config` is now used to set user and group access to various SFDB directories.
- Use the `dbed_update` command to create a new SQLite SFDB repository.
- `sfua_rept_adm` was used in 5.0 to perform repository backup and restore and this command will be obsolete in 5.1 SP1.

- The `sfua_rept_util` command is used to perform SQLite repository backup and restore.
- The Database Dynamic Storage Tiering (DBDST) feature is rebranded as SmartTier for Oracle and includes expanded functionality to support management of sub-file objects.
- Cached ODM: functionality is expanded with Cache ODM Manager

Commands which continue to be supported:

- `dbed_update`, `dbed_rept_node`
- Database Storage Checkpoint commands: `dbed_ckptcreate`, `dbed_ckptdisplay`, `dbed_ckptmount`, `dbed_ckptquota`, `dbed_ckptremove`, `dbed_ckptrollback`, `dbed_clonedb`, `dbed_ckptumount`
- Quick I/O commands: `qio_getdbfiles`, `qio_recreate`, `qio_convertdbfiles qio_recreate`, `qioadmin`,`qiostat`,`qio_convertdbfiles_ora`,`qio_getdbfiles_ora`,`qio_recreate_ora`, `qiomkfile`
- Database Flashsnap commands: `dbed_vmchecksnap`, `dbed_vmclonedb`, `dbed_vmsnap`
- SmartTier for Oracle commands: `dbdst_addvol`, `dbdst_admin`, `dbdst_classify`, `dbdst_convert`, `dbdst_file_move`, `dbdst_partition_move`, `dbdst_preset_policy`, `dbdst_rmvol`, `dbdst_show_fs`, `dbdst_tbs_move`, `dbdst_report`

New commands:

- 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 SFCFS products:

- Disk layout Version 4 and 5 are deprecated. A disk layout Version 4 or 5 file system can still be mounted, but this will be disallowed in future releases. Symantec recommends that you upgrade the disk layout to the latest default disk layout version.
- 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 HP-UX operating systems

This release of Veritas products can only be installed on a system running HP-UX B.11.31.1009, HP-UX 11i Version 3 September 2010 Operating Environments Update Release or later on the PA-RISC or Itanium platforms.

To verify the operating system version use the `swlist` command as follows:

```
# swlist | grep HPUX11i
HPUX11i-DC-OE     B.11.31.1009   HP-UX Data Center Operating Environment
```

JFS must be installed on your system prior to installing any Veritas software.

To verify that JFS is installed use the `swlist` command as follows:

```
# swlist -l product JFS
JFS              B.11.31        Base VxFS File System 4.1 for HP-UX
```

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://entsupport.symantec.com/docs/330441

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

Memory requirements

2 GB of memory is required for Veritas Storage Foundation Cluster File System.

CPU requirements

A minimum of 2 CPUs is required for Veritas Storage Foundation Cluster File System.

Node requirements

All nodes in a Cluster File System must have the same operating system version and update level.

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

Table 1-1 *(continued)*

Veritas Storage Foundations feature	DB2	Oracle	Sybase
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.

Disk space requirements

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

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

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

Number of nodes supported

SFCFS is capable of supporting cluster configurations with up to 64 nodes. Symantec has tested and qualified configurations of up to 32 nodes on IA-64 (Itanium) 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://www.symantec.com/docs/TECH144835

Fixed issues

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

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

See "Documentation" on page 54.

Veritas Storage Foundation Cluster File System fixed issues

This section describes the incidents that are fixed in Veritas Storage Foundation Cluster File System.

Table 1-2 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
1856719	Fixed an smap update issue due to which 'df' and 'rm' commands can hang.
1544221	Optimize getattr call to speedup the case when binaries are mmaped from multiple nodes.

Veritas File System fixed issues

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

Table 1-3 Veritas File System fixed issues

Incident	Description
1537570	Fixed the cause of a hang when mounting a frozen file system.
1537588	Fixed the cause of an assertion panic.
1665456	Fixed the cause of a panic in the `volkio_to_kio_copy`() call.
1805046	Fixed an incorrect alert generation from VxFS when the file system usage threshold is set.
1828175	Fixed the cause of a system panic in `inctext`, in which `VTEXT` was not set and `tcount` was greater than 0.
1903977	Fixed an issue in which a panic could happen because of a mutex getting destroyed while the protected structure was still in use.
1922948	You can now mount deprecated disk layout versions so that they can be upgraded.

Table 1-3 Veritas File System fixed issues *(continued)*

Incident	Description
1926141	Direct I/O can now be enabled using the `mincache` and `convosync` mount options with the base VxFS license.
1935374	Storage Checkpoint creation now properly fails to assign a metadata allocation policy on a data-only volume.
1945171	Setting vxfs_bc_bufhwm higher than the available physical memory now reports an error.
1946063	Fixed a performance issue in `fsadm` where `fsadm` kept relocating and copying already reorganized regions of a file in subsequent passes.
1969314	Interchanged the order in which `fcache_vn_destroy`() and `VFS_TEARDOWN_STACK`() are called to avoid a panic.
2017776	The virtual memory area is no longer destroyed if there are active mappings on the vnode Fix to during a force unmount.
2018439	The `fsppadm` command no longer dumps core if a volume does not have placement tags.
2026603	Added quota support for the user "nobody".
2026625	The `sar -v` command now properly reports VxFS inode table overflows.
2026675	Fixed the cause of ENOTBLK being returned via the async driver due to fdd not being a block device.
2028782	Fixed an issue in which `Q_SETQUOTA` was not setting the current usage using the `quotactl`() API.
2050070	Fixed an issue in which the volume manager area was destroyed when spinlock was held.
2072121	Fixed the cause of a hang that was due to a disowned beta semaphore.
2098371	Fixed a performance issue in which write I/O performance degraded rapidly when the size of a file reached 64 MB and I/O size was not a multiple of 64 bytes.
2106668	Fixed an issue in which in some cases EFBIG was returned soon after resizing a file system through the `fsadm` command.

Table 1-3 Veritas File System fixed issues *(continued)*

Incident	Description
2111614	Fixed an issue in which file modification time was not updated when the `O_SYNC` and `nodatainlog` mount options were used.
2149407	Fixed an issue in which modification and access times were not getting updated through mmap and msync.
2163013	Fixed an issue in which the `odmstat` command was showing very high average IO time.

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 1-4 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.
1089875	Increasing vol_maxspecialio to 1 MB on HP-UX.
1097258	vxconfigd hung when an array is disconnected.
1239188	Enhance vxprivutil to enable, disable, and display config+log copies state.
1301991	When vxconfigd is restarted with -k option, all log messages are sent to stdout. syslog should be the default location.
1321475	Join Failure Panic Loop on axe76 cluster.
1405756	CVM: Add support to set PFTO values cluster-wide.

Table 1-4 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
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.
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.

Table 1-4 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
1826088	After pulling out the Fibre Channel cables of a local site array, plex becomes DETACHED/ACTIVE.
1829337	Array firmware reversal led to disk failure and offlined all VCS resources
1831634	CVR: Sending incorrect sibling count causes replication hang, which can result in I/O hang.
1831969	VxVM: ddl log files are created with world write permission
1835139	I/Os hung after giveback of NetApp array filer
1840673	After adding new LUNs, one of the nodes in 3 node CFS cluster hangs
1846165	Data corruption seen on cdsdisks on Solaris-x86 in several customer cases
1857558	Need to ignore jeopardy notification from GAB for SFCFS/RAC, since oracle CRS takes care of fencing in this stack
1857729	CVM master in the VVR Primary cluster panicked when rebooting the slave during VVR testing
1860892	Cache Object corruption when replaying the CRECs during recovery
1869995	VVR: Improve Replication performance in presence of SO snapshots on secondary.
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1874034	Race between modunload and an incoming IO leading to panic
1880279	Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment.
1881336	VVR: Primary node panicked due to race condition during replication
1884070	When running iotest on a volume, the primary node runs out of memory
1897007	vxesd coredumps on startup when the system is connected to a switch which has more than 64 ports
1899688	VVR: Every I/O on smartsync enabled volume under VVR leaks memory
1899943	CPS based fencing disks used along with CPS servers does not have coordinator flag set
1901827	vxdg move fails silently and drops disks.

Table 1-4 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
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.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout
1992872	vxresize fails after DLE.
1993953	CVM Node unable to join in Sun Cluster environment due to wrong coordinator selection
1998447	Vxconfigd dumps core due to incorrect handling of signal
1999004	I/Os hang in VxVM on linked-based snapshot
2002703	Misleading message while opening the write protected device.
2009439	CVR: Primary cluster node panicked due to queue corruption

Table 1-4 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2010426	Tag setting and removal do not handle wrong enclosure name
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
2016129	Tunable to disable OS event monitoring by vxesd
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1
2021737	vxdisk list shows HDS TrueCopy S-VOL read only devices in error state.
2025593	vxdg join hang/failure due to presence of non-allocator inforecords and when tagmeta=on
2027831	vxdg free not reporting free space correctly on CVM master. vxprint not printing DEVICE column for subdisks.
2029480	Diskgroup join failure renders source diskgroup into inconsistent state
2029735	System panic while trying to create snapshot
2034564	I/Os hung in serialization after one of the disks which formed the raid5 volume was pulled out
2036929	Renaming a volume with link object attached causes inconsistencies in the disk group configuration
2038137	System panics if volrdmirbreakup() is called recursively.
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.
2040150	Existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration
2052203	Master vold restart can lead to DG disabled and abort of pending transactions.
2052459	CFS mount failed on slave node due to registration failure on one of the paths
2055609	Allocation specifications not being propagated for DCO during a grow operation
2060785	Primary panics while creating primary rvg
2061066	vxisforeign command fails on internal cciss devices

Table 1-4 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2061758	Need documentation on list of test suites available to evaluate CDS code path and verification of the code path.
2063348	Improve/modify error message to indicate its thin_reclaim specific
2070531	Campus cluster: Couldn't enable site consistency on a dcl volume, when trying to make the disk group and its volumes siteconsistent.
2075801	VVR: "vxnetd stop/start" panicked the system due to bad free memory
2076700	VVR: Primary panic due to NULL pointer dereference
2094685	Diskgroup corruption following an import of a cloned BCV image of a SRDF-R2 device
2097320	Events generated by dmp_update_status() are not notified to vxconfigd in all places.
2105547	Enabling tagmeta=on on a disk group no longer causes a delay in disk group split/join operations.
2105722	VVR: I/O hang on Primary with link-breakoff snapshot
2112568	System panics while attaching back two Campus Cluster sites due to incorrect DCO offset calculation
2122009	vxddladm list shows incorrect hba information after running vxconfigd -k
2126731	vxdisk -p list output is not consistent with previous versions
2131814	VVR: System panic due to corrupt sio in _VOLRPQ_REMOVE

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 54.

Issues related to installation

This section describes the known issues during installation and upgrade.

The Web-based installer does not work from the disc (2321818)

The Web-based installer fails to run.

Workarounds:

For this first workaround, you need to have about 1.7 GB of local storage available. Copy the disc to a local system and start the Web-based installer from the local copy. Symantec recommends that you use cpio for these operations.

If you have limited local disk space, use the second workaround.

To start the Web-based installer workaround

1. Create a mount point.

   ```
   # mkdir /mnt/dvd
   ```

2. Optionally to find the specific device path (/dev/dsk/cxtxdx), run this command:

   ```
   # /usr/sbin/ioscan -fnkC disk
   ```

3. Mount the disc to the mount point.

   ```
   # mount /dev/dsk/cxtxdx /mnt/dvd
   ```

4. Create a temporary installation directory.

   ```
   # mkdir /tmp/HXRT51SP1
   ```

5. Create a symbolic link from the disc to the temporary installation directory.

   ```
   # ln -s /mnt/dvd/* /tmp/HXRT51SP1/
   ```

6. Remove the installer link from the temporary installation directory.

   ```
   # rm -rf /tmp/HXRT51SP1/scripts
   ```

7. Copy the installer scripts from the disc to the temporary installation directory.

   ```
   # cp -rf /mnt/dvd/scripts/ /tmp/HXRT51SP1/
   ```

8. Start the Web-based installer from the temporary installation directory.

   ```
   # /tmp/HXRT51SP1/webinstaller start
   ```

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 depot not installed on system_name
```

Workaround:

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

Installer assigns duplicate node ID during `-addnode` procedure

While performing an `-addnode` using a CPI installer to a cluster where a node has failed, VCS appends the new node with a duplicate node ID of its last node. This happens only to the cluster in which any but the last node has failed. In this case, `/etc/llthost` displays two nodes with same node IDs. This is because VCS assigns the node ID by simply counting the number of node entries without checking the assigned node IDs.

Workaround: Instead of running the CPI command, add the new node manually as described in the Veritas Cluster Server Installation Guide.

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 version listed after upgrading (2121881)

When you upgrade from SFCFS 5.1 RP2 to SFCFS 5.1 SP1, the previous version is incorrectly listed as 5.1.001.000

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 SP1, the user umounts all VxFS file systems including the one which hosts NetBackup binaries (/usr/openv). While upgrading SF 5.1 SP1, 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 /tmp directory. After you umount the NetBackup file system, manually copy these two version files from /tmp to their original path. If the path doesn't 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 systems already affected by this issue: Manually install VRTSpbx, VRTSat, VRTSicsco packages 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 depots and patches needs. During migration some depots are already installed and during migration some

depots 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 depot is deprecated (2032052)

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

- Fresh installs: Do not install VRTSacclib.
- Upgrade: Ignore VRTSacclib.
- Uninstall: Ignore VRTSacclib.

SFCFSHA upgrade shows partial upgrade warning

When you install 5.1 SFCFSHA and try to upgrade to SFCFSHA 5.1SP1 using the `./installsfcfs` command, you may receive a partial upgrade error message.

Workaround: Use the `./installer -upgrade` command instead of the `./installsfcfs` command.

Veritas Storage Foundation Cluster File System known issues

This section describes the known issues in this release of Veritas Storage Foundation Cluster File System (SFCFS).

Miscalculated file set usage (2123429)

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

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

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

Workaround

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

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

Multiple CFSmount resources are in a single service group they may not all come online after a reboot (2164670)

In some cases when multiple CFSmount resources are in a single service group they may not all come online after a reboot. You will need to manually bring them online after a reboot.

Workaround

Create a resource dependency between the various CFSmount resources.

installer –makeresponsefile detects the wrong product (2044525)

If you generate a response file to upgrade SFCFS or SFCFSHA using the `./installer -makeresponsefile` command, and then choose `G` (Upgrade a Product) option, the installer detects it as SFCFS RAC.

You can safely ignore that the installer detects it as SFCFS RAC.

CVMVolDg agent may fail to deport CVM disk group

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

Workaround: If multiple CVMVolDg resources are configured for a shared disk group, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

Veritas File System known issues

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

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

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

Workaround: There is no workaround for this issue.

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

Shrinking a file system 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.

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.

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

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

Workaround: There is no workaround for this issue.

Veritas Volume Manager known issues

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

vxdg split or join operations can fail for disks with a disk media name greater than or equal to 27 characters (2063387)

If a disk's media name is greater than or equal to 27 characters, certain operations, such as diskgroup split or join, can fail with the following error:

```
VxVM vxdg ERROR : vxdg move/join dg1 dg2 failed subdisk_name : Record
already exists in disk group
```

VxVM uses disk media names to create subdisk names. If multiple subdisks are under the same disk, then the serial number, starting from 1, is generated and appended to the subdisk name so as to identify the given subdisk under the physical disk. The maximum length of the subdisk name is 31 characters. If the disk media name is long, then the name is truncated to make room for serial numbers. Therefore, two diskgroups can end up having same subdisk names due to this truncation logic, despite having unique disk media names across diskgroups. In such scenarios, the diskgroup split or join operation fails.

Workaround:

To avoid such problems, Symantec recommends that disk media name length should be less than 27 characters.

After initializing a disk for native LVM, the first instance of vxdisk list fails with a 'get_contents' error and errant flags are displayed (2074640)

After you initialize a disk that is under the operating system's native LVM control and not under Veritas Volume Manager (VxVM) control by using the `pvcreate` *`path_to_physical_disk`* command, the first time that you run the `vxdisk list` *`disk_name`* command, the command displays the following error:

```
VxVM vxdisk ERROR V-5-1-539 Device disk_name: get_contents failed:
          Disk device is offline
```

In addition, the `flags` field is incorrectly populated. However, in the next instantiation of the same command, VxVM does not produce an error and the flags are correctly populated with the LVM tag.

Workaround:

Issue the `vxdisk list` *`disk_name`* command a second time.

vxconfigd fails to allocate memory until the daemon is restarted (2112448)

Veritas Volume Manager (VxVM) utilities may fail with the following error message:

```
Memory allocation failure
```

This error implies that there is insufficient memory for the `vxconfigd` daemon. A program's data segment size is enforced by the operating system tunable `maxdsiz`. The default value of `maxdsiz` is 1 GB. With this default `maxdsiz` value, the `vxconfigd` daemon can allocate a maximum of 1 GB of memory.

Workaround:

You might need to increase the operating system `maxdsiz` tunable's value appropriately to increase the data storage segment for the programs.

See the `maxdsiz`(5) manual page for more information.

After increasing the value, you must stop and restart the `vxconfigd` daemon. Depending on the `maxdsiz` tunable value, `vxconfigd` can allocate a maximum up to 2 GB of memory on PA machines, and 4 GB of memory on IA machines.

The vxcdsconvert utility is not supported for EFI disks (2064490)

Pending decision about whether to include.

Node join can lead to hang if an upgrade of the cluster protocol version is in progress (2103567)

If you attempt to join a node to the cluster while Cluster Volume Manager (CVM) is upgrading the cluster protocol version, the system may hang. This issue occurs if the node is attempting to join the cluster after you issue the `vxdctl upgrade` command to upgrade the CVM cluster.

Work-around:

Avoid joining a new node to the cluster until the CVM cluster upgrade is completed.

Thin reclamation on disks with the hpdisk format is not supported (2136238)

Thin reclamation on disks with the hpdisk format is not supported. An attempt to perform reclamation on such disks automatically aborts.

Work-around:

There is no workaround for this issue.

vxdisksetup fails to give a LUN the cdsdisk format if the LUN is larger than 1 TB and the system is using Tachyon HBAs (2146340)

The `vxdisksetup` command fails to initialize a LUN to have the cdsdisk format if the LUN is larger than 1 TB and the system is using Tachyon HBAs. The `vxdisksetup` command displays the following error:

```
VxVM vxdisk ERROR V-5-1-5433 Device disk_name: init failed:
        Disk is not useable, bad format
```

Work-around:

There is no workaround for this issue.

Shared disk group creation on slave fails if the naming scheme on slave is operating system native scheme with the mode as the new name (2148981)

While creating shared disk groups on slaves using the command shipping feature, the disk group creation may fail if the naming scheme on the slave where the command was issued is the operating system's native scheme with the mode as the new name.

Workaround:

You can create the shared disk group from the slave by changing the naming scheme to the operating system's native scheme while in the "Legacy" mode.

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.

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

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

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

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

While vradmin changeip is running, 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:

```
# /sbin/init.d/vras-vradmind.sh stop
# /sbin/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.

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

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

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

Workaround:

To restore vradmin functionality after a master switch operation

1. Restart `vradmind` on all cluster nodes. Enter the following:
2. Re-enter the command that failed.

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.

Issues related to Symantec Product Authentication Service with VCS

This section covers the known issues related to Symantec Product Authentication Service (AT) in this release.

Verification for VRTSat package or patch returns errors

If you run swverify command on VRTSat package or patch, the command returns errors for missing files on VRTSat.CLIENT-PA32. [1244204]

Workaround: This message may be safely ignored.

Issues related to LLT

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

LLT port stats sometimes shows recvcnt larger than recvbytes

With each received packet, LLT increments the following variables:

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

Both these variables are integers. With constant traffic, recvbytes hits and rolls over MAX_INT quickly. This can cause the value of recvbytes to be less than the value of recvcnt. [1788315]

This does not impact the LLT functionality.

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

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This occurs in some corner cases even though a port is not even registered on the peer node. [1809827]

Issues related to I/O fencing

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

All nodes in a sub-cluster panic if the node that races for I/O fencing panics

At the time of a network partition the lowest node in each sub-cluster races for the coordination points on behalf of that sub-cluster. If the lowest node is unable to contact a majority of the coordination points or the lowest node itself unexpectedly panics during the race, then all the nodes in that sub-cluster will panic. [1965954]

Coordination Point agent does not provide detailed log message for inaccessible CP servers

The Coordination Point agent does not log detailed information of the CP servers that are inaccessible. When CP server is not accessible, the agent does not mention the UUID or the virtual IP of the CP server in the engine log. [1907648]

Preferred fencing does not work as expected for large clusters in certain cases

If you have configured system-based or group-based preferred fencing policy, preferred fencing does not work if all the following cases are true:

- The fencing setup uses customized mode with one or more CP servers.
- The application cluster has more than eight nodes.
- The node weight for a single node (say galaxy with node id 0) is more than the sum total of node weights for the rest of the nodes.
- A network fault occurs and the cluster partitions into two with the single node (galaxy) on one part and the rest of the nodes on the other part.

Under such circumstances, for group-based preferred fencing, the single node panics even though more high priority services are online on that node. For system-based preferred fencing, the single node panics even though more weight is assigned to the node. [2161816]

See the *Veritas Storage Foundation Cluster File System Administrator's Guide* for more information on preferred fencing.

Server-based I/O fencing fails to start after configuration on nodes with different locale settings

On each (application cluster) node, the vxfen module retrieves and stores the list of the UUIDs of coordination points. When different nodes have different locale settings, the list of UUIDs on one (application) node does not match with that of the other (application) nodes. Hence, I/O fencing does not start after configuration. [2112742]

Workaround: Start I/O fencing after fixing the locale settings to use the same values on all the (application) cluster nodes.

Reconfiguring Storage Foundation Cluster File System HA with I/O fencing fails if you use the same CP servers

When you reconfigure an application cluster that uses server-based I/O fencing (customized fencing mode), the installer does not remove the application cluster

information from the CP servers before the reconfiguration. As a result, if you reconfigure the application cluster and choose to configure I/O fencing in customized mode using the same CP servers, then reconfiguration of server-based fencing for the application cluster fails. [2076240]

Workaround: Manually remove the application cluster information from the CP servers after you reconfigure Storage Foundation Cluster File System HA but before you reconfigure server-based I/O fencing for the application cluster.

See the *Veritas Cluster Server Administrator's Guide* for instructions to remove the application cluster information from the CP servers.

CP server cannot bind to multiple IPs (2085941)

Coordination point server (CP server) binds only to a single virtual IP and listens on the same. Application clusters cannot access the CP server if it fails to establish connection to this virtual IP. Therefore, if the connection fails because of the subnet in which the virtual IP of the CP server exists, you cannot access the CP server even if there is another subnet through which the client can connect to the CP server over a different IP.

Resolution: No known resolution for this issue.

Installer is unable to split a cluster that is registered with one or more CP servers

Splitting a cluster that uses server-based fencing is currently not supported. [2110148]

You can split a cluster into two and reconfigure Storage Foundation Cluster File System HA on the two clusters using the installer. For example, you can split a cluster *clus1* into *clus1A* and *clus1B*.

However, if you use the installer to reconfigure the Storage Foundation Cluster File System HA, the installer retains the same cluster UUID of *clus1* in both *clus1A* and *clus1B*. If both *clus1A* and *clus1B* use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: None.

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 5.0MP2 to 5.1SP1 (2003131)

While upgrading from 50mp2 to 51SP1 the following error message could be seen when running `sfua_rept_migrate`:

```
# /opt/VRTSdbed/migrate/sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
SFORA sfua_rept_migrate ERROR V-81-8903 Could not start repository database
/usr/lib/dld.sl: Can't find path for shared library: libcur_colr.1
/usr/lib/dld.sl: No such file or directory
sh: 3845 Abort(coredump)
Symantec DBMS 3.0.85.0 vxdbms_start_db utility
ASA failed. Sybase ASA error code: [134].
Sybase ASA Error text: {{{}}}

SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Using SFDB tools after upgrading Oracle to 11.2.0.2 (2203228)

The procedure which Oracle recommends for upgrading to Oracle 11.2.0.2 results in the database home changing. After you upgrade to Oracle 11.2.0.2, you must run the `dbed_update`command with the new Oracle home provided as an argument to the -H option before using any SFDB tools. After this step, the SFDB tools can be used normally.

Database fails over during Flashsnap operations (1469310)

In an SFCFS 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 54.

Veritas Storage Foundation Cluster File System software limitations

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

cfsmntadm command does not verify the mount options (2078634)

You must confirm if the mount options are correct which are then passed to the `cfsmntadm` command. If the mount options are incorrect, the mount fails and the CFSMount resource will not come online. You can check the VCS engine log file for any mount failure messages.

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 Manager software limitations

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

DMP settings for NetApp storage attached environment

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

Table 1-5

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

The change is persistent across reboots.

To change the tunable parameters

1 Issue the following commands:

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

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

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

Veritas 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.

Limitations related to I/O fencing

This section covers I/O fencing-related software limitations.

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or "split brain." See the *Veritas Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

I/O fencing uses SCSI-3 PR keys to implement data protection. Keys are placed on I/O fencing coordinator points and on data disks. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator points and data disks to prevent possible difficulties with subsequent cluster startup.

Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordinator points and data disks. Depending on the order of reboot and subsequent startup events, the cluster may warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown -r command on one node at a time and wait for each node to complete shutdown.

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

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

Documentation set

Table 1-6 lists the documentation for Veritas Storage Foundation Cluster File System.

Table 1-6 Veritas Storage Foundation Cluster File System documentation

Document title	File name
Veritas Storage Foundation Cluster File System Release Notes	sfcfs_notes_51sp1_hpux.pdf
Veritas Storage Foundation Cluster File System Installation Guide	sfcfs_install_51sp1_hpux.pdf
Veritas Storage Foundation Cluster File System Administrator's Guide	sfcfs_admin_51sp1_hpux.pdf

Table 1-7 lists the documents for Veritas Cluster Server.

Table 1-7 Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_51sp1_hpux.pdf
Veritas Cluster Server Release Notes	vcs_notes_51sp1_hpux.pdf
Veritas Cluster Server Administrator's Guide	vcs_admin_51sp1_hpux.pdf

Table 1-7 Veritas Cluster Server documentation *(continued)*

Title	File name
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_51sp1_hpux.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev_51sp1.pdf
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	vcs_vvr_agent_51sp1_hpux.pdf
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_51sp1_hpux.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_51sp1_hpux.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_51sp1_hpux.pdf

Table 1-8 lists the documentation for Veritas Storage Foundation.

Table 1-8 Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_51sp1_hpux.pdf
Veritas Storage Foundation and High Availability Installation Guide	sf_install_51sp1_hpux.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sf_adv_ora_51sp1_hpux.pdf
Veritas Storage Foundation Advanced Features Administrator's Guide	sf_adv_admin_51sp1_hpux.pdf

Table 1-9 lists the documentation for Veritas Volume Manager and Veritas File System.

Table 1-9 Veritas Volume Manager and Veritas File System documentation

Document title	File name
Veritas Volume Manager Administrator's Guide	vxvm_admin_51sp1_hpux.pdf
Veritas Volume Manager Troubleshooting Guide	vxvm_tshoot_51sp1_hpux.pdf
Veritas File System Administrator's Guide	vxfs_admin_51sp1_hpux.pdf

Table 1-9 Veritas Volume Manager and Veritas File System documentation *(continued)*

Document title	File name
Veritas File System Programmer's Reference Guide	vxfs_ref_51sp1_hpux.pdf

Table 1-10 lists the documentation for Veritas Volume Replicator.

Table 1-10 Veritas Volume Replicator documentation

Document title	File name
Veritas Volume Replicator Administrator's Guide	vvr_admin_51sp1_hpux.pdf
Veritas Volume Replicator Planning and Tuning Guide	vvr_planning_51sp1_hpux.pdf
Veritas Volume Replicator Advisor User's Guide	vvr_advisor_users_51sp1_hpux.pdf

Table 1-11 lists the documentation for Symantec Product Authentication Service (AT).

Table 1-11 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.