

Veritas Storage Foundation™ Cluster File System Release Notes

Linux

5.1 Service Pack 1



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: 5.1SP1.5

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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](#)
- [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 Cluster File System (SFCFS) version for Linux. Review this entire document before you install SFCFS.

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

This is Document version: 5.1SP1.5 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>

For the latest information on updates, patches, and known issues regarding this release, see the following TechNote on the Symantec Technical Support website:

<http://www.symantec.com/docs/TECH141448>

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/TECH75506>
- 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 Cluster File System.

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 *product installation guide*.

Using the installer for Veritas Dynamic Multi-pathing (DMP)

You can use the script- or Web-based installer to install, configure, and uninstall Veritas Dynamic Multi-pathing. You can enable DMP using the DMP license or using any Storage Foundation license key.

Using the installer for Symantec Virtual Store (SVS)

You can use the script- or Web-based installer to install, configure, and uninstall Symantec VirtualStore. You can enable SVS using an SVS license.

Unencapsulation not required for some upgrade paths

Unencapsulation is no longer required for certain upgrade paths.

See the *product installation guide*.

The new VRTSamf RPM is now included in all high availability products

The new VRTSamf RPM is now included in all high availability products. The asynchronous monitoring framework (AMF) allows the more intelligent monitoring of resources, lower resource consumption, and increased availability across clusters.

See the *product installation guide*.

The VRTScutil and VRTSacclib RPMs are no longer in use

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

See the *product 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 Cluster File System 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 *product installation guide*.

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

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

See the *product installation guide*.

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

You can now configure disk-based fencing for the Storage Foundation Cluster File System HA cluster using the Web-based installer.

See the *product installation guide*.

The installer can automatically detect and configure LLT links

The installer detects link connection status among all cluster nodes and chooses the most suitable links for LLT communication. It then can set the priority of the LLT private heartbeat links based on their media speed. Aggregated and bonded NICs are supported.

See the *product installation guide*.

The Web-based installer supports adding nodes

The Web-based installer has increased parity with the script-based installer. It now supports the ability to add nodes to a cluster. It also supports configuring secure clusters and fencing configuration.

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.

Packaging updates

The following lists package changes in this release.

Changes related to Storage Foundation Cluster File System

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

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
```

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 master based 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 to Thin Provisioning and Thin Reclamation features

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

SmartMove default changed

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

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

New initialization options for the `vxassist grow` command

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

<code>init=sync</code>	If the volume has multiple plexes, VxVM synchronizes the data between the plexes during initialization.
<code>init=zero</code>	Initializes the volume or grown region, and initializes the associated data plexes to zeroes. If the volume resides on thin reclaimable LUNs, VxVM also reclaims the space within the storage array
<code>init=active</code>	Initializes the volume or grown region without modifying the existing data on the plexes.
<code>init=default</code>	Performs the default operation.

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

Relayout operations on VxFS mounted volumes now use SmartMove

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

Reclamation writes are not counted in write statistics

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

```
# vxstat -fm
```

Changes related to Veritas File System

Veritas File System includes the following changes:

Autolog replay on mount

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

Dynamic Storage Tiering is rebranded as SmartTier

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

FileSnap

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

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

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:

Veritas Volume Manager persisted attributes

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

Automatic recovery of volumes during disk group import

After a disk group is imported, disabled volumes are enabled and started by default. To control the recovery behavior, use the `vxdefault` command to turn on or off the tunable `autostartvolumes`. If you turn off the automatic recovery, the recovery behaves the same as in previous releases. This behavior is useful if you want to perform some maintenance after importing the disk group, and then start the volumes. To turn on the automatic recovery of volumes, specify `autostartvolume=on`.

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

Enhancements to the vxrootadm command

The `vxrootadm` command has the following new options:

- `vxrootadm split`
Splits the root disk mirror into a new root disk group.
- `vxrootadm join`
Reattaches mirrors from an alternate root disk group to the current (booted) root disk group.
- `vxrootadm addmirror`
Adds a mirror of the root disk to the root disk group, for redundancy in case the current root disk fails.
- `vxrootadm rmmirror`
Deletes a root disk mirror from the current (booted) root disk group.

See the `vxrootadm(1m)` man page.

In-place upgrade of Veritas Volume Manager in presence of root disk encapsulation

When you upgrade from Veritas Volume Manager (VxVM) 5.1 to VxVM 5.1 SP1, you can upgrade the `VRTSvxvm` RPM without having to unencapsulate the root disk. The Veritas installer uses the `rpm -U` command to perform an in-place upgrade. Previously, the installer uninstalled the old `VRTSvxvm` RPM and installed the new one.

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

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

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

Default format for auto-configured disk has changed

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

Changes related to Veritas Dynamic Multi-Pathing (DMP)

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

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

In 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.

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 `vxndmpadm(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:

- Storage Foundation for Oracle RAC is supported
- 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`

Changes to LLT

This release includes the following new features and changes to LLT:

- **LLT startup time** through the LLT init script is now optimized to use a constant time. LLT takes less than 16 seconds to start irrespective of the number of links specified in `/etc/llttab` file.
In the previous releases, LLT took around $(5 * \text{number_of_links_specified_in_the_}/etc/llttab_file)$ seconds to start.
- **The lltstat command** includes the following new options:
 - **lltstat -nv active**
This command filters the output of `lltstat -nv` to display the status of only the active nodes in the cluster.
 - **lltstat -nv configured**
This command filters the output of `lltstat -nv` to display the status of only the configured nodes in the cluster. Configured nodes include active nodes and any additional nodes which are listed in the `/etc/llthosts` file.

See the `lltstat` manual page for more information.

- **Support for different link speeds for LLT links**
LLT now removes the restriction to use private NICs with same media speed. You can now use different media speed for the private NICs and configure the NICs with lesser speed as low-priority links to enhance LLT performance.
- **Support for destination-based load balancing**
LLT now also provides destination-based load balancing where the LLT link is chosen based on the destination node id and the port. With destination-based load balancing, LLT sends all the packets of a particular destination on a link.

See the *product installation guide* and the *product administration guide* for more details.

Changes to GAB

This section lists the new features and changes related to GAB in this release.

- **GAB logging daemon**
GAB implements a distributed network protocol. For situations when GAB decides to take the drastic action of killing its userland client process or panicking a node to resolve an issue, data from the affected node alone may not suffice for a meaningful support analysis. The new `gablogd` daemon

attempts to address this issue. GAB starts this daemon by default at GAB configuration time.

See the *Veritas Cluster Server Administrator's Guide* for more information.

Changes to I/O fencing

This section covers the new features and changes related to I/O fencing in this release.

Support for preferred fencing

Traditional fencing prevents a split-brain condition by allowing only one of multiple sub-clusters to continue its operation in case a network partition disrupts regular communication between nodes. The preferred fencing feature gives preference to one sub-cluster over other sub-clusters in determining the surviving sub-cluster. This preference is based on factors such as which of the sub-clusters is running higher priority applications or the total importance of nodes which form that sub-cluster or both.

See the *product installation guide* and the *product administration guide* for more details.

Support for Non-SCSI3 fencing

In environments that do not support SCSI-3, non-SCSI-3 fencing provides reasonable data protection by causing the winning side to delay by a configurable amount (`loser_exit_delay`, default 55). Additionally, Symantec has enhanced the fencing component to help panic the losing side quickly. Together, these enhancements help narrow down the window of potential data corruption drastically.

See the *product installation guide* and the *product administration guide* for more details.

Enhancements to server-based fencing

This release includes the following enhancements and new features related to server-based fencing:

- Single CP-server based fencing
Support to use a single highly available CP server that is configured on an SFHA cluster to provide server-based fencing support for multiple application clusters

Support to migrate between fencing modes when the cluster is running

The `vxfsnwap` utility now supports migrating between disk-based and server-based fencing configurations in a cluster that is running.

See the *product administration guide* for more details.

No longer supported

The following features are not supported in this release of SFCFS 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`, `vxdbs_fstatsummary`, `dbdst_fiostat_collector`, `vxdbs_get_datafile_stats`
- `dbed_saveconfig`, `dbed_checkconfig`
- `dbed_ckptplan`, `dbed_ckptpolicy`
- `dbed_scheduler`
- `sfua_rept_migrate` with `-r` and `-f` options

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems

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

The Veritas 5.1 SP1 release supports the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)

Note: Symantec VirtualStore is only supported on RHEL 5 x86_64 U3 or higher.

- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel) or SP3 on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5-default kernel) or SUSE Linux Enterprise Server 11 (SLES 11) with SP1 on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux 5 (OEL 5) with Update 3 or later (Red Hat compatible kernel mode only)

Note: 64-bit operating systems are only supported.

If your system is running an older version of either Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Oracle Enterprise Linux, you must upgrade it before attempting to install the Veritas software. Consult the Red Hat, SUSE, or Oracle documentation for more information on upgrading or reinstalling your system.

Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported Red Hat errata and updates and SUSE service packs is available in the following TechNote. Read this TechNote before you install Symantec products.

<http://entsupport.symantec.com/docs/335001>

Required Linux RPMs for VCS

Make sure you installed the following operating system-specific RPMs on the systems where you want to install or upgrade VCS. VCS will support any updates made to the following RPMs, provided the RPMs maintain the ABI compatibility.

[Table 1-1](#) lists the RPMs that VCS requires for a given Linux operating system.

Table 1-1 Required RPMs

Operating system	Required RPMs
RHEL 5	compat-libstdc++-33-3.2.3-61.x86_64 rpm glibc-2.5-42.i686 rpm glibc-2.5-42.x86_64 rpm ksh-20080202-14.el5.x86_64 rpm libgcc-4.1.2-46.el5.i386 rpm libgcc-4.1.2-46.el5.x86_64 rpm libstdc++-4.1.2-46.el5.i386 rpm pam-0.99.6.2-6.el5.x86_64 rpm
SLES 10	compat-libstdc++-5.0.7-22.2.x86_64 rpm glibc-2.4-31.54.x86_64 rpm glibc-32bit-2.4-31.54.x86_64 rpm ksh-93s-59.7.x86_64 rpm libgcc-4.1.2_20070115-0.21.x86_64 rpm libstdc++-4.1.2_20070115-0.21.x86_64rpm pam-0.99.6.3-28.13.x86_64 rpm
SLES 11	glibc-2.9-13.2.x86_64 glibc-32bit-2.9-13.2.x86_64 rpm ksh-93t-9.4.x86_64 rpm libgcc43-32bit-4.3.3_20081022-11.18.x86_64 rpm libgcc43-4.3.3_20081022-11.18.x86_64 rpm libstdc++33-3.3.3-11.9.x86_64 rpm libstdc++43-32bit-4.3.3_20081022-11.18.x86_64 rpm

Mandatory patch required for Oracle Bug 4130116

If you are running Oracle versions 9.2.0.6 or 9.2.0.7, you must apply the Oracle patch for Oracle Bug 4130116. Contact Oracle to obtain this patch, and for details on how to apply it.

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.

Supported Oracle versions

Oracle versions 10g Release 2 and 11g Release 1 are supported for use with Storage Foundation Cluster File System for Oracle RAC.

Database requirements

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

Table 1-2

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

Table 1-2 (continued)

Veritas Storage Foundations feature	DB2	Oracle	Sybase
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.

Cross-Platform Data Sharing licensing

The Cross-Platform Data Sharing (CDS) feature is also referred to as Portable Data Containers.

The ability to import a CDS disk group on a platform that is different from the platform on which the disk group was last imported is controlled by a CDS license. CDS licenses are included as part of the Veritas Storage Foundation license.

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 supports cluster configurations with up to 64 nodes. Symantec has tested and qualified configurations of up to 32 nodes at the time of release.

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

<http://entsupport.symantec.com/docs/335001>

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 Cluster File System fixed issues

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

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 RP2

Table 1-3 Veritas Storage Foundation Cluster File System 5.1 RP2 fixed issues (listed incident number, parent number)

Fixed issues	Description
1982730, 1952484	Fixed a panic in <code>vx_recv_getemapmsg()</code> due to an alignment fault.
2043651, 1991446	Changing the nodes in a cluster from <code>largefiles</code> to <code>no largefiles</code> with the <code>fsadm</code> command no longer results in the following error when you re-mount the nodes: <pre>UX:vxfs mount: ERROR: V-3-21272: mount option(s) incompatible with file system</pre>
1933839, 1807536	Added support for <code>VX_FREEZE_ALL</code> ioctl in a cluster environment.
2069672, 2069059	Fixed a hang issue in a cluster environment.
2049381, 2049378	Fixed an issue that caused database checkpoint rollback to fail on a non-English locale setup.

Table 1-3 Veritas Storage Foundation Cluster File System 5.1 RP2 fixed issues (listed incident number, parent number) *(continued)*

Fixed issues	Description
2092088, 2030289	Fixed a file system corruption issue in a cluster environment that occurred when mounting the file system on the secondary node while the primary node was 100% full.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 RP1

Table 1-4 Veritas Storage Foundation Cluster File System 5.1 RP1 fixed issues (listed incident number, parent number)

Fixed issues	Description
1878583, 1544221	getattr call optimization to speedup the case when binaries are being mmapped from many nodes on CFS.

Veritas File System fixed issues

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

Table 1-5 Veritas File System fixed issues

Incident	Description
2026603	Added quota support for the user "nobody".
2026625	The <code>sar -v</code> command now properly reports VxFS inode table overflows.
2050070	Fixed an issue in which the volume manager area was destroyed when spinlock was held.
2080276	Fixed the cause of a panic in <code>vx_naio_worker()</code> .

Veritas File System: Issues fixed in 5.1 RP2

Table 1-6 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
1939772	Fixed an issue in vxrepquota(1m) where username and groupname were truncated to 8 characters
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.
2085395	Fixed a hang issue in vxfsckd.
2072162	Fixed the issue of writing zero length to null buffer
2059621	Fixed a panic due to null pointer de-reference in vx_unlockmap()
2016345	Fixed an error EINVAL issue with O_CREATE while creating more than 1 million files.
1976402	Fixed the issue in fsck replay where it used to double fault for 2TB luns.

Table 1-6 Veritas File System fixed issues in 5.1 RP2 (*continued*)

Fixed issues	Description
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.
2072161	Fixed a hang issue in vx_traninit()
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
2026637	Support for kernel extended attributes
2072165	Fixed an active level leak issue while fsadm resize operation.
2059008	Fixed an issue with quotas where hard limit was not enforced in CFS environment
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 1-7 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 1-7 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) (*continued*)

Fixed issues	Description
1934085, 1871935	We now update ilst 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 1-8 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 1-8 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 vxconfigd is restarted with -k option, all log messages are sent to stdout. syslog should be the default location.
1321475	Join Failure Panic Loop on axe76 cluster.
1441406	'vxdisk -x list' displays wrong DGID.
1458792	After upgrade from SF5.0mp1 to SF5.0mp3, *unit_io and *pref_io was set to 32m.
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.
1485075	DMP sending I/O on an unopened path causing I/O to hang
1504466	VxVM: All partitions aren't created after failing original root disk and restoring from mirror.
1513385	VVR:Primary panic during autosync or dcm replay.
1528121	FMR: wrong volpagemod_max_memsz tunable value cause buffer overrun
1528160	An ioctl interrupted with EINTR causes frequent vxconfigd exits.
1586207	"vxsnap refresh" operations fail occasionally while data is replicating to secondary.
1589022	Infinite looping in DMP error handling code path because of CLARIION APM, leading to I/O hang.
1594928	Avoid unnecessary retries on error buffers when disk partition is nullified.
1662744	RVG offline hung due to I/Os pending in TCP layer
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.

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

Incident	Description
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
1848722	VOL_NOTE_MSG definition needs to be revisited
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.
1873220	LoP Root disk encapsulation failed on RHEL5_U4, system goes into the panic state
1874034	Race between modunload and an incoming IO leading to panic
1875054	After upgrade to 5.0MP3, CDS disks are presented as LVM disks.

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

Incident	Description
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 core dumps 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

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

Incident	Description
1982715	vxclustadm dumps core during memory re-allocation.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout
1992872	vxresize fails after DLE.
1993953	CVM Node unable to join in Sun Cluster environment due to wrong coordinator selection
1998447	Vxconfigd dumps core due to incorrect handling of signal
1999004	I/Os hang in VxVM on linked-based snapshot
2002703	Misleading message while opening the write protected device.
2009439	CVR: Primary cluster node panicked due to queue corruption
2010426	Tag setting and removal do not handle wrong enclosure name
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
2016129	Tunable to disable OS event monitoring by vxesd
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1
2021737	vxdisk list shows HDS TrueCopy S-VOL read only devices in error state.
2025593	vxdg join hang/failure due to presence of non-allocator inforecords and when tagmeta=on
2027831	vxdg free not reporting free space correctly on CVM master. vxprint not printing DEVICE column for subdisks.
2029480	Diskgroup join failure renders source diskgroup into inconsistent state
2029735	System panic while trying to create snapshot
2034564	I/Os hung in serialization after one of the disks which formed the raid5 volume was pulled out
2036929	Renaming a volume with link object attached causes inconsistencies in the disk group configuration
2038137	System panics if volrdmirbreakup() is called recursively.
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.

Table 1-8 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
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
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
2112547	VxVM/PowerPath Migration Enabler Interop issue. Host panics when running vxdisk
2122009	vxddladm list shows incorrect hba information after running vxconfigd -k
2126731	vxdisk -p list output is not consistent with previous versions

Table 1-8 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 1-9 Veritas Volume Manager 5.1 RP2 fixed issues

Fixed issues	Description
1973367	VxVM Support for Virtio Virtual Disks in KVM virtual Machines
1938907	RHEL5 U3: WWN information is not displayed due to incorrect device information returned by HBA APIs
2069022	Booting between Linux kernels results in stale APM key links.
2067568	EqualLogic iSCSI - Disabling array switch port leads to disk failure and disabling of path.
2015570	File System read failure seen on space optimized snapshot after cache recovery
1665094	Snapshot refresh causing the snapshot plex to be detached.
2015577	VVR init scripts need to exit gracefully if VVR license not installed.
1992537	Memory leak in vxconfigd causing DiskGroup Agent to timeout
1946936	CVM: IO hangs during master takeover waiting for a cache object to quiesce
1946939	CVM: Panic during master takeover, when there are cache object I/Os being started on the new master
2053975	Snapback operation panicked the system
1983768	IO hung on linked volumes while carrying out third mirror breakoff operation.
1513385	VVR:Primary panic during autosync or dcm replay.
2052459	CFS mount failed on slave node due to registration failure on one of the paths
1936611	vxconfigd core dump while splitting a diskgroup
1992872	Vxresize fails after DLE.

Table 1-9 Veritas Volume Manager 5.1 RP2 fixed issues (*continued*)

Fixed issues	Description
1960341	Toggling of naming scheme is not properly updating the daname in the vxvm records.
1933528	During Dynamic reconfiguration vxvm disk ends up in error state after replacing physical LUN.
2019525	License not present message is wrongly displayed during system boot with SF5.1 and SFM2.1
1933375	Tunable value of 'voliomem_chunk_size' is not aligned to page-size granularity
2040150	Existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration
1441406	'vxdisk -x list' displays wrong DGID
1956777	CVR: Cluster reconfiguration in primary site caused master node to panic due to queue corruption
1942985	Improve locking mechanism while updating mediatype on vxvm objects
1911546	Vxrecover hung with layered volumes
2012016	Slave node panics while vxrecovery is in progress on master
2078111	When the IOs are large and need to be split, DRL for linked volumes cause I/Os to hang
1880279	Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment.
1952177	Machine panics after creating RVG
1929083	Vxattachd fails to reattach site in absence of vxnotify events
1097258	vxconfigd hung when an array is disconnected
1972755	TP/ETERNUS:No reclaim seen with Stripe-Mirror volume.
2061066	vxisforeign command fails on internal cciss devices
2021737	vxdisk list shows HDB TrueCopy S-VOL read only devices in error state.
2065669	After upgrading to 5.1, reinitializing the disk makes public region size smaller than the actual size.

Table 1-9 Veritas Volume Manager 5.1 RP2 fixed issues (*continued*)

Fixed issues	Description
1974393	Avoiding cluster hang when the transaction client timed out
2038735	Incorrect handling of duplicate objects resulting in node join failure and subsequent panic.
2031462	Node idle events are generated every second for idle paths controlled by Third Party drivers.
1982715	vxclustadm dumping core while memory re-allocation.
1998447	Vxconfigd dumped core due to incorrect handling of signal
1999004	I/Os hang in VxVM on linked-based snapshot
1899943	CPS based fencing disks used along with CPS servers does not have coordinator flag set
1923906	CVM: Master should not initiate detaches while leaving the cluster due to complete storage failure
2006454	AxRT5.1P1: vxsnap prepare is displaying vague error message
1989662	/opt/VRTSsfmh/bin/vxlist causes panic.
2059046	FMR:TP: snap vol data gets corrupted if vxdisk reclaim is run while sync is in progress
2011316	VVR: After rebooting 4 nodes and try recovering RVG will panic all the slave nodes.
1485075	DMP sending I/O on an unopened path causing I/O to hang
1874034	Race between modunload and an incoming IO leading to panic
2055609	Allocation specifications not being propagated for DCO during a grow operation
2029480	Diskgroup join failure renders source diskgroup into inconsistent state
2029735	System panic while trying to create snapshot
1897007	vxesd coredumps on startup when the system is connected to a switch which has more than 64 ports
1831969	VxVM: ddl log files are created with world write permission
2010426	Tag setting and removal do not handle wrong enclosure name

Table 1-9 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
1920894	vxcheckhbaapi can loop forever
1920761	I/O hang observed after connecting the storage back to master node incase of local detach policy
2034104	Unable to initialize a disk using vxdiskadm
1946941	vxsnap print shows incorrect year
1829337	Array firmware reversal led to disk failure and offlined all VCS resources
2034564	I/Os hung in serialization after one of the disk which formed the raid5 volume was pulled out
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
2126731	VxVM 5.1: vxdisk -p list output is not consistent with previous versions

Veritas Volume Manager: Issues fixed in 5.1 RP1

Table 1-10 Veritas Volume Manager 5.1 RP1 fixed issues

Fixed issues	Description
1938484	EFI: Prevent multipathing don't work for EFI disk
1915356	I/O stuck in vxvm caused cluster node panic
1899688	[VVR] Every I/O on smartsync enabled volume under VVR leaks memory
1884070	When running iotest on volume, primary node runs out of memory
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1860892	Cache Object corruption when replaying the CRECs during recovery

Table 1-10 Veritas Volume Manager 5.1 RP1 fixed issues (*continued*)

Fixed issues	Description
1857729	CVM master in the VVR Primary cluster panic when rebooting the slave during VVR testing
1857558	[CVM] Need to ignore jeopardy notification from GAB for SFCFS/RAC, since oracle CRS takes care of fencing in this stack
1840673	After adding new luns one of the nodes in 3 node CFS cluster hangs
1835139	CERT : pnate test hang I/O greater than 200 seconds during the filer giveback
1826088	After pulling out FC cables of local site array, plex became DETACHED/ACTIVE
1792795	supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.

LLT, GAB, and I/O fencing fixed issues

[Table 1-11](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-11 LLT, GAB, and I/O fencing fixed issues

Incident	Description
1908938	[GAB] In a large cluster, cascaded lowest node failures result in GAB panic during sequence space recovery.
1840826	[GAB] Prevent 'gabconfig -c' while port 'a' is in the middle of iofence processing.
1861439 1849527	[LLT] Removing the LLT links from a single node in a four-node cluster causes other nodes to panic.
2066020	[LLT] The <code>dlpiping</code> utility exits with an error similar to "dlpiping: send ECHO_REQ failed."
2005045	[LLT] The <code>hastart</code> command fails to start HAD on one of the nodes with message "GabHandle::open failed errno = 16" in syslog after HAD is stopped on all the nodes in the cluster simultaneously.

Table 1-11 LLT, GAB, and I/O fencing fixed issues (*continued*)

Incident	Description
1859023	[LLT] The <code>lltconfig -T query</code> command displays a partially incorrect output
1846387 2084121	[Fencing] The <code>vxfenswap</code> and the <code>vxfentsthdw</code> utilities fail when rsh or ssh communication is not set to the same node.
1922413	[Fencing] The <code>vxfentsthdw</code> utility should detect storage arrays which interpret NULL keys as valid for registrations/reservations.
1847517	[Fencing] The <code>vxfenswap</code> utility has an incorrect usage message for <code>-n</code> option
1992560	[Fencing] The <code>vxfentsthdw</code> utility uses <code>scp</code> to communicate with the local host.
2098065	[Fencing] The <code>vxfenclearpre</code> utility cannot clear keys from coordinator disks and data disks when there is a preexisting split brain.
1512956	[Fencing] The <code>vxfenclearpre</code> utility displays error messages
2143933	[VxCPS] For a four-node cluster, the installer fails to configure server-based fencing which uses three CP servers as its coordination points. The process fails while registering the CP clients on the third CP server.
2097935	[VxCPS] Need strict host name matching in coordination point installer.

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 1-12 Veritas Storage Foundation for Databases tools fixed issues

Incident	Description
1873738	The <code>dbed_vmchecksnap</code> command may fail
1399393	Clone command fails on an Oracle RAC database
1736516	Clone command fails for instant checkpoint on Logical Standby database
1789290	<code>dbed_vmclondb -o recoverdb</code> for offhost fails for Oracle 10gr2 and prior versions

Table 1-12 Veritas Storage Foundation for Databases tools fixed issues
(continued)

Incident	Description
1810711	Flashsnap reverse resync command fails on offhost flashsnap cloning

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 87.

Symantec VirtualStore

This section describes the known issues in this release of Symantec VirtualStore (SVS).

Virtual machines created by the FileSnap wizard might not boot correctly if during the FileSnap process the VirtualStore cluster node, the ESX Server, or the vCenter Server reboots (2164664, 2374229)

In some cases when you clone using FileSnap, and one of the following servers crashes or reboots while the clone process is in progress, the clones might not get created correctly:

- SVS nodes
- ESX host on which the clones are being created
- vCenter Server

Even if the clones appear in the vCenter inventory as created, the clones GuestOS might not be able to boot.

Workaround

Delete all of the clones that were created when the servers crashed or were rebooted, and redo the FileSnap operation.

Issues related to installation

This section describes the known issues during installation and upgrade.

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

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

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

Workaround:

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

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

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

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

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, the user umounts all VxFS file systems including the one which hosts NetBackup binaries (*/usr/opensv*). While upgrading SF 5.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure packages VRTSspb, VRTSat, and VRTSsisco, which causes NetBackup to stop working.

Workaround: Before you umount the VxFS file system which hosts NetBackup, copy the two files */usr/opensv/netbackup/bin/version* and */usr/opensv/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/opensv/netbackup/bin` and `mkdir -p /usr/opensv/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 VRTSspb, VRTSat, VRTSsisco 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 RPMs and patches needs. During migration some RPMs are already installed and during migration some RPMs are removed. This releases disk space. The installer then claims more space than it actually needs.

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

The VRTSaclib RPM is deprecated (2032052)

The VRTSaclib RPM is deprecated. For installation, uninstallation, and upgrades, note the following:

- Fresh installs: Do not install VRTSaclib.

- Upgrade: Ignore VRTSacclib.
- Uninstall: Ignore VRTSacclib.

Error messages in syslog (1630188)

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

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

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

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

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

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

Or

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

Kernel image: /boot/vmlinuz-2.6.16.60-0.54.5-smp

Initrd image: /boot/initrd-2.6.16.60-0.54.5-smp

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

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

SFCFSHA upgrade shows partial upgrade warning

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

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

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 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
```

The cfsmntadm add command may fail with no errors (2169538)

The `cfsmntadm add` command fails, if one host name is a substring of another host name in the list.

Note: VOM is affected by this issue when adding a CFS mount to a cluster that has systems with host names that are substrings of each other.

Workaround

Run the `cfsmntadm` command with the `"all="` option on one of the nodes in the CFS cluster to add the cfsmounts to all nodes.

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.

NFS issues with VxFS Storage Checkpoint (1974020)

NFS clients mounting VxFS Storage Checkpoints that are NFS-exported by SFCFS cluster nodes using a Virtual IP may receive the following error message upon Virtual IP failover:

```
Stale NFS file handle
```

This is a result of major numbers of VxFS Storage Checkpoints not necessarily being the same on all SFCFS cluster nodes.

Application group attempts to come online on primary site before fire drill service group goes offline on the secondary site

The application service group comes online on the primary site while the fire drill service group attempts to go offline at the same time, causing the application group to fault. (2107386)

Workaround: Ensure that the fire drill service group is completely offline on the secondary site before the application service group comes online on the primary site.

Veritas Volume Manager known issues

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

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 `vxctl upgrade` command to upgrade the CVM cluster.

Work-around:

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

Post encapsulation of the root disk, system comes back up after first reboot unencapsulated (2119038)

In some cases, after encapsulating the root disk and rebooting the system, it may come up without completing the encapsulation. This happens because the `vxvm-reconfig` startup script is unable to complete the encapsulation process.

Workaround

Reboot the system or run the following command.

```
# service vxvm-reconfig start
```

This will reboot the system and complete the remaining stages of encapsulation.

Required attributes for Veritas Volume Manager (VxVM) devices to avoid boot failures (1411526)

To support iSCSI devices, Veritas Volume Manager (VxVM) does not start non-root devices until runlevel2. The boot process expects all local (non-NFS) mount points in the `/etc/fstab` file to be present at boot time. To avoid boot failures, all VxVM entries in the `/etc/fstab` file must have the `_netdev` attribute, and must not have the `fsck` required flag set. These attributes enable VxVM to defer mounting of VxVM devices until after VxVM has started.

vxrestored daemon fails to restore disabled paths (1663167)

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

Workaround:

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

To enable the `mpt_disable_hotplug_remove` tunable

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

```
options mptsas mpt_disable_hotplug_remove=0
```

- 2 Rebuild the `initrd` image:

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

- 3 Reboot the system.

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

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

Workaround:

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

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

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

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

Workaround:

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

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

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

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

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

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

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

Workaround:

To recover from this situation

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

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

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

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

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

Root disk encapsulation issue (1603309)

Encapsulation of root disk will fail if it has been assigned a customized name with vxdmpadm(1M) command. If you wish to encapsulate the root disk, make sure that you have not assigned a customized name to its corresponding DMP node. See vxdmpadm(1M) and the section "Setting customized names for DMP nodes" on page 173 for details.

VxVM starts before OS device scan is done (1635274)

While working with some arrays, VxVM may start before all devices are scanned by the OS. This slow OS device discovery may result in malfunctioning of VM, fencing and VCS due to partial disks seen by VxVM.

Workaround:

After the fabric discovery is finished, issue the `vxdisk scandisks` command to bring newly discovered devices into the VxVM configuration.

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 `vxdbg 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.

Converting LVM volumes to VxVM volumes fails when multipathed storage devices are present (1471781, 1931727)

The `vxvmconvert` utility cannot convert LVM volumes to VxVM volumes when multipathed storage devices are present. This issue occurs with SLES 11 and RHEL5, due to changes in the LVM utilities. If multipathed devices are detected, the `vxvmconvert` utility exits with the following error:

```
vxvmconvert cannot convert multipathed devices on SLES11 systems.  
... Exiting.
```

Workaround: There is no workaround for this issue.

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 1-13](#) shows the Hitachi arrays that have new array names.

Table 1-13 Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP

Table 1-13 Hitachi arrays with new array names (*continued*)

Previous name	New name
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

DMP disables subpaths and initiates failover when an iSCSI link is failed and recovered within 5 seconds. (2100039)

When using iSCSI S/W initiator with an EMC CLARiON array, iSCSI connection errors may cause DMP to disable subpaths and initiate failover. This situation occurs when an iSCSI link is failed and recovered within 5 seconds.

Work around:

When using iSCSI S/W initiator with an EMC CLARiON array, set the `node.session.timeo.replacement_timeout` iSCSI tunable value to 40 secs or higher.

Device discovery failure during VxVM startup in SUSE (2122771)

When the system boots up, some devices are not displayed in the `vxdisk list` output. This issue occurs if the `vold` daemon does the device discovery before the

operating system (OS) completes its device discovery. Therefore, the vold daemon may miss some devices.

Work-around:

Configure the vxvm-startup script to wait until the operating system discovery is completed before starting vold. In the `/etc/vx/vxvm-startup` file, set `DEV_DISCOVER_DELAY` to the expected device discovery time taken by the OS. By default, `DEV_DISCOVER_DELAY` is set to 0.

You must reboot the system before this configuration applies. To discover the missed devices, run the `vxdisk scandisks` command or the `vxctl enable` command.

DMP marks the subpaths as DISABLED while these subpaths are accessible from OS level (2037222)

For iSCSI devices on SLES 10 SP3, the DMP tunable parameter `dmp_fast_recovery` needs to be turned off.

```
# vxmpadm settune dmp_fast_recovery=off
```

DMP panics if a DDL device discovery is initiated immediately after loss of connectivity to the storage (2040929)

When using EMC Powerpath with VxVM 5.1SP1 on SLES11, set the `fast_io_fail_tmo` on the HBA port to any non-zero value that is less than the `dev_loss_tmo` value so as to avoid a panic in case a DDL device discovery is initiated by the `vxdisk scandisks` command or the `vxctl enable` command immediately after loss of connectivity to the storage.

A system can hang or panic in a SAN boot environment due to udev device removal after loss of connectivity to some paths on SLES 11 (2219626)

The issue may occur with NetApp LUNs in ALUA mode, with a SAN boot configuration. When a device fails with a `dev_loss_tmo` error, the operating system (OS) device files are removed by udev. After this removal, a system can hang or panic due to I/O disruption to the boot device. To avoid this issue, use the following workaround.

Workaround

To update the kernel and create the new rules file

- 1 Save the existing rules file.

```
# mkdir /savefiles
# cd /etc/udev/rules.d/
# mv 40-VxVM.rules /savefiles
```

- 2 Download and upgrade to kernel 2.6.27.45-0.1.1 or later from Novell.
- 3 Create the file `/etc/udev/rules.d/40-rport.rules` with the following content line:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports",
ACTION=="add",RUN+="/bin/sh -c 'echo 20 >
/sys/class/fc_remote_ports/%k/fast_io_fail_tmo;echo 864000
>/sys/class/fc_remote_ports/%k/dev_loss_tmo'"
```

- 4 Reboot the system.
- 5 If new LUNs are dynamically assigned to the host, run the following command:

```
# udevadm trigger --action=add --subsystem-match=fc_remote_ports
```

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.

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.

Panic occurs when VxFS module parameter `vxfs_hproc_ext` is set to 1 and you attempt to mount a clone promoted file system (2163931)

A system panic occurs if the following two conditions are met:

- The VxFS module parameter `vxfs_hproc_ext` is set to 1.
- A clone is promoted as a primary using the `fsckpt_restore` command, and then you attempt to mount the promoted file system.

Workaround: There is no workaround for this issue.

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

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

Workaround: There is no workaround for this issue.

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

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

```
# vxfilesnap source_file target_dir
```

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

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

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

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

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

Workaround: There is no workaround for this issue.

VxFS module loading fails when freevxfss module is loaded (1736305)

The following module loading error can occur during RPM installation if the `freevxfss` module is loaded:

```
Error in loading module "vxfss". See documentation.
```

```
ERROR: No appropriate VxFS drivers found that can be loaded.  
See VxFS documentation for the list of supported platforms.
```

Workaround: Ensure that the `freevxfss` module is not loaded before installing the `VRTSvxfss` RPM. The following command shows if the `freevxfss` module is loaded:

```
# lsmod | grep freevxfss
```

If the `freevxfss` module is loaded, unload the module:

```
# rmmod freevxfss
```

A mount can become busy after being used for NFS advisory locking

If you export a VxFS file system using NFS and you perform file locking from the NFS client, the file system can become unable to be unmounted. In this case, the `umount` command fails with the `EBUSY` error.

Workaround: Force unmount the file system:

```
# vxumount -o force /mount1
```

umount can hang when inotify watches are used (1590324)

If inotify watches are used, then an unmount can hang in the `vx_softcnt_flush()` call. The hang occurs because inotify watches increment the `i_count` variable and cause the `v_os_hold` value to remain elevated until the inotify watcher releases the hold.

Workaround: There is no workaround for this issue.

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).

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

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:

```
vradmin 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 `vradmin` on all the cluster nodes using the following commands:

```
# /etc/init.d/vras-vradmin.sh restart
```

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

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

Workaround:

To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT  
IPM_HEARTBEAT_TIMEOUT=120
```

- 2 Restart `vradmin` to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following:

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 bunkerdgrname 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 -t vxfs /dev/vx/dsk/dg/snapshot_volume
```

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

For example:

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

Running SUSE Linux and using Novell's YaST tool to configure an IPv6 address may result in an error

When Novell's YaST tool is invoked to configure an IPv6 address on a different network interface and if:

- the host name, the DNS server name and domain name are specified to the YaST tool.
- IPv6 address is assigned by the Dynamic Host Configuration Protocol (DHCP).
- the "Write Hostname to `/etc/hosts`" option is selected (this is selected by default).

This results in the `vradmin` command returning the following error:

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

This happens because the YaST tool can replace the `/etc/hosts` entry containing `127.0.0.2` from the IPv4 host name to the specified new IPv6 host name. For example:

```
127.0.0.2 v6hostname.space.ipv6.com v6hostname
```

Workaround: The following procedure resolves this issue.

To resolve this issue

- 1 Edit the `/etc/hosts` file to specify the correct IPv6 address.
- 2 Restart the `vradmin` daemon on all VVR hosts:

```
# /etc/init.d/vras-vradmin.sh restart
```

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

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

Workaround: There is no workaround for this issue.

vxassist layout removes the DCM (2162522)

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

- 1 Pause or stop the applications.
- 2 Wait for the RLINKs to be up to date. Enter the following:

```
# vxrlink -g diskgroup status rlink
```
- 3 Stop the affected RVG. Enter the following:

```
# vxrvrg -g diskgroup stop 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:

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

- 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.

The `vcsat` and `cpsat` commands may appear to be hung

The following commands may appear to be hung when you invoke them from the command shell:

- `/opt/VRTScps/bin/cpsat`
- `/opt/VRTSvc/bin/vcsat`

This issue occurs when the command requires some user interaction. [1841185]

Workaround:

- To fix the issue for `vcsat`, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTSvc  
# /opt/VRTSvc/bin/vssatvcs command_line_argument  
# unset EAT_HOME_DIR
```

- To fix the issue for `cpsat`, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTScps  
# /opt/VRTScps/bin/vssatcps command_line_argument  
# unset EAT_HOME_DIR
```

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 GAB

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

Trace messages from the gablogd daemon on the console for RHEL5 Update 5 or later

On RHEL5 Update 5 or later, the `gablogd` daemon prints informational and trace messages similar to the following [2139883]:

```
INFO: task gablogd:22812 blocked for more than 120 seconds.
"echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
gablogd      D ffff81004100daa0      0 22812      1      23106 22809 (NOTLB)
 ffff810faf539e38 0000000000000082 000000000000084c 0000000000000001
 ffff810faf539de8 0000000000000007 ffff810fc2a130c0 ffff810138ee8100
 000019f130082599 00000000000018572 ffff810fc2a132a8 000000001f76c3d63
Call Trace:
 [<ffffffffff88ee3690>] :gab:gab_linux_sv_wait+0x53/0x68
 [<ffffffffff8008e68d>] default_wake_function+0x0/0xe
 [<ffffffffff88ecd4c8>] :gab:gab_daemonlog+0xae1/0xc52
```

```
[<fffffffff88ee326c>] :gab:gab_linux_ioctl1+0x10e/0x1a3  
[<fffffffff88ee331d>] :gab:gab_linux_compat_ioctl1+0x1c/0x20  
[<fffffffff800fbe53>] compat_sys_ioctl1+0xc5/0x2b2  
[<fffffffff8006249d>] sysenter_do_call+0x1e/0x76
```

Workaround: As the operating system message indicates, set the following:

```
echo 0 > /proc/sys/kernel/hung_task_timeout_secs
```

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 *product administration 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.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

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, use the following steps to reattach the volumes.

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

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

```
# vxdg join snapshot_disk_group_name  
           primary_disk_group_name
```

- 2 Start all the volumes in primary disk group.

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

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

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

Repeat this step for all the volumes.

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

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

Workaround

There is no workaround for this issue.

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

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

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

```
dbed_vmclonedb started at 2009-08-26 11:32:16
Editing remote_login_passwordfile in initclone2.ora.
Altering instance_name parameter in initclone2.ora.
Altering instance_number parameter in initclone2.ora.
Altering thread parameter in initclone2.ora.
SFORA dbed_vmclonedb ERROR V-81-4918 Database clone2 has not been
correctly recovered.
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.20569/recover.log.
```

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

Workaround

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

To resolve this known issue

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

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

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

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

- 4 Start the database.

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 87.

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.

Obtaining information about mounted file system states (1764098)

For accurate information about the state of mounted file systems on Linux, refer to the contents of `/proc/mounts`. The `mount` command may or may not reference this source of information depending on whether the regular `/etc/mtab` file has been replaced with a symbolic link to `/proc/mounts`. This change is made at the discretion of the system administrator and the benefits are discussed in the mount online manual page. A benefit of using `/proc/mounts` is that changes to SFCFS mount options are accurately displayed for all nodes.

Veritas File System software limitations

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

Linux I/O Scheduler for Database Workloads (1446361)

Symantec recommends using the Linux deadline I/O scheduler for database workloads on both Red Hat and SUSE distributions.

To configure a system to use this scheduler, include the `elevator=deadline` parameter in the boot arguments of the GRUB or LILO configuration file.

The location of the appropriate configuration file depends on the system’s architecture and Linux distribution:

Configuration File

/boot/grub/menu.lst

Architecture and Distribution

RHEL5 x86_64, SLES10 x86_64, and SLES11 x86_64

For the GRUB configuration files, add the `elevator=deadline` parameter to the kernel command. For example, change:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2
    initrd /boot/initrd-2.6.18-128.el5.img
```

To:

```
title RHEL5UP3
    root (hd1,1)
    kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2 \
    elevator=deadline
    initrd /boot/initrd-2.6.18-128.el5.img
```

A setting for the elevator parameter is always included by SUSE in its LILO and GRUB configuration files. In this case, change the parameter from `elevator=cfq` to `elevator=deadline`.

Reboot the system once the appropriate file has been modified.

See the Linux operating system documentation for more information on I/O schedulers.

Recommended limit of number of files in a directory

To maximize VxFS performance, do not exceed 100,000 files in the same directory. Use multiple directories instead.

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 1-14

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 behavior on Linux SLES11 when connectivity to a path is lost (2049371)

On SLES 11, when the connectivity to a path is lost, the SLES 11 kernel removes the device path from its database. DMP reacts to the UDEV event that is raised in this process, and marks the device path as DISABLED[M]. DMP will not use the path for further I/Os. Unlike on other flavours of Linux, the path state is DISABLED[M] instead of DISABLED. Subsequently, if the path comes back online, DMP responds to the UDEV event to signal the addition of device path into SLES 11 kernel. DMP enables the path and changes its state to ENABLED.

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 errata

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

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

See [“Documentation”](#) on page 87.

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

Correction for setting up a disaster recovery fire drill

Topic: Setting up a disaster recovery fire drill

Issue: The content below is incorrect:

After the fire drill service group is taken offline, reset the value of the ReuseMntPt attribute to 1 for all Mount resources.

Use the following corrected information:

After the fire drill service group is taken offline, reset the value of the ReuseMntPt attribute to 0 for all Mount resources.

Updates for the procedure to set up non-SCSI3 fencing in virtual environments manually

Topic: Setting up non-SCSI3 fencing in virtual environments manually

Some information is missing in the procedure that is documented in the Installation Guide. Refer to the following procedure.

Setting up non-SCSI3 fencing in virtual environments manually

To manually set up I/O fencing in a non-SCSI-3 PR compliant setup

- 1 Configure I/O fencing in customized mode with only CP servers as coordination points.
- 2 Make sure that the Storage Foundation Cluster File System HA cluster is online and check that the fencing mode is customized.

```
# vxfenadm -d
```

- 3 Make sure that the cluster attribute UseFence is set to SCSI3.

```
# haclus -value UseFence
```

- 4 On each node, edit the /etc/vxenviron file as follows:

```
data_disk_fencing=off
```

- 5 On each node, edit the /etc/sysconfig/vxfen file as follows:

```
vxfen_vxfnd_tmt=25
```

- 6 On each node, edit the /etc/vxfenmode file as follows:

```
loser_exit_delay=55  
vxfen_script_timeout=25
```

Refer to the sample /etc/vxfenmode file.

- 7 On each node, set the value of the LLT sendhbcap timer parameter value as follows:
 - Run the following command:

```
lltconfig -T sendhbcap:3000
```

- Add the following line to the `/etc/llttab` file so that the changes remain persistent after any reboot:

```
set-timer senhbcap:3000
```

8 On any one node, edit the VCS configuration file as follows:

- Make the VCS configuration file writable:

```
# haconf -makerw
```

- For each resource of the type `DiskGroup`, set the value of the `MonitorReservation` attribute to `0` and the value of the `Reservation` attribute to `NONE`.

```
# hares -modify <dg_resource> MonitorReservation 0
```

```
# hares -modify <dg_resource> Reservation "NONE"
```

- Run the following command to verify the value:

```
# hares -list Type=DiskGroup MonitorReservation!=0
```

```
# hares -list Type=DiskGroup Reservation!="NONE"
```

The command should not list any resources.

- Modify the default value of the `Reservation` attribute at type-level.

```
# haattr -default DiskGroup Reservation "NONE"
```

- Make the VCS configuration file read-only

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

9 Make sure that the `UseFence` attribute in the VCS configuration file `main.cf` is set to `SCSI3`.

10 To make these `VxFEN` changes take effect, stop and restart `VxFEN` and the dependent modules

- On each node, run the following command to stop VCS:

```
# /etc/init.d/vcs stop
```

- After VCS takes all services offline, run the following command to stop VxFEN:

```
# /etc/init.d/vxfen stop
```

- On each node, run the following commands to restart VxFEN and VCS:

```
# /etc/init.d/vxfen start
# /etc/init.d/vcs start
```

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 1-15](#) lists the documentation for Veritas Storage Foundation Cluster File System.

Table 1-15 Veritas Storage Foundation Cluster File System documentation

Document title	File name
<i>Veritas Storage Foundation Cluster File System Release Notes</i>	sfdfs_notes_51sp1_lin.pdf
<i>Veritas Storage Foundation Cluster File System Installation Guide</i>	sfdfs_install_51sp1_lin.pdf
<i>Veritas Storage Foundation Cluster File System Administrator's Guide</i>	sfdfs_admin_51sp1_lin.pdf

[Table 1-16](#) lists the documentation for Symantec VirtualStore.

Table 1-16 Symantec VirtualStore documentation

Document title	File name
<i>Veritas Storage Foundation Cluster File System Release Notes</i>	sfdfs_notes_51sp1_lin.pdf
The Symantec VirtualStore content is documented in the <i>Veritas Storage Foundation Cluster File System Release Notes</i> .	

Table 1-16 Symantec VirtualStore documentation (*continued*)

Document title	File name
<i>Symantec VirtualStore Installation and Configuration Guide</i>	virtualstore_install_51sp1_lin.pdf
<i>Symantec VirtualStore Administrator's Guide</i>	virtualstore_admin_51sp1_lin.pdf

[Table 1-17](#) lists the documents for Veritas Cluster Server.

Table 1-17 Veritas Cluster Server documentation

Title	File name
<i>Veritas Cluster Server Installation Guide</i>	vcs_install_51sp1_lin.pdf
<i>Veritas Cluster Server Release Notes</i>	vcs_notes_51sp1_lin.pdf
<i>Veritas Cluster Server Administrator's Guide</i>	vcs_admin_51sp1_lin.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents_51sp1_lin.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev_51sp1pr4.pdf
<i>Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide</i>	vcs_vvr_agent_51sp1_lin.pdf
<i>Veritas Cluster Server Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_agent_51sp1_lin.pdf
<i>Veritas Cluster Server Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_agent_51sp1_lin.pdf
<i>Veritas Cluster Server Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_agent_51sp1_lin.pdf

[Table 1-18](#) lists the documentation for Veritas Storage Foundation.

Table 1-18 Veritas Storage Foundation documentation

Document title	File name
<i>Veritas Storage Foundation Release Notes</i>	sf_notes_51sp1_lin.pdf
<i>Veritas Storage Foundation and High Availability Installation Guide</i>	sf_install_51sp1_lin.pdf

Table 1-18 Veritas Storage Foundation documentation (*continued*)

Document title	File name
<i>Veritas Storage Foundation: Storage and Availability Management for Oracle Databases</i>	sf_adv_ora_51sp1_lin.pdf
<i>Veritas Storage Foundation Advanced Features Administrator's Guide</i>	sf_adv_admin_51sp1_lin.pdf

[Table 1-19](#) lists the documentation for Veritas Volume Manager and Veritas File System.

Table 1-19 Veritas Volume Manager and Veritas File System documentation

Document title	File name
<i>Veritas Volume Manager Administrator's Guide</i>	vxvm_admin_51sp1_lin.pdf
<i>Veritas Volume Manager Troubleshooting Guide</i>	vxvm_tshoot_51sp1_lin.pdf
<i>Veritas File System Administrator's Guide</i>	vxfs_admin_51sp1_lin.pdf
<i>Veritas File System Programmer's Reference Guide</i>	vxfs_ref_51sp1_lin.pdf

[Table 1-20](#) lists the documentation for Veritas Volume Replicator.

Table 1-20 Veritas Volume Replicator documentation

Document title	File name
<i>Veritas Volume Replicator Administrator's Guide</i>	vvr_admin_51sp1_lin.pdf
<i>Veritas Volume Replicator Planning and Tuning Guide</i>	vvr_planning_51sp1_lin.pdf
<i>Veritas Volume Replicator Advisor User's Guide</i>	vvr_advisor_users_51sp1_lin.pdf

[Table 1-21](#) lists the documentation for Symantec Product Authentication Service (AT).

Table 1-21 Symantec Product Authentication Service documentation

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

Table 1-21 Symantec Product Authentication Service documentation (*continued*)

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

Manual pages

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

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

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

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

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

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

See the `man(1)` manual page.

Manual pages are divided into sections 1, 1M, 3N, 4, and 4M. Edit the `man(1)` configuration file `/etc/man.config` to view these pages.

To edit the man(1) configuration file

- 1 If you use the man command to access manual pages, set `LC_ALL` to “C” in your shell to ensure that the pages are displayed correctly.

```
export LC_ALL=C
```

See incident 82099 on the Red Hat Linux support website for more information.

- 2 Add the following line to `/etc/man.config`:

```
MANPATH /opt/VRTS/man
```

where other man paths are specified in the configuration file.

- 3 Add new section numbers. Change the line:

```
MANSECT          1:8:2:3:4:5:6:7:9:tcl:n:l:p:o
```

to

```
MANSECT          1:8:2:3:4:5:6:7:9:tcl:n:l:p:o:3n:1m
```

