

Veritas Storage Foundation™ for Oracle® RAC Release Notes

Linux

5.1 Service Pack 1



Veritas Storage Foundation™ for Oracle RAC 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: 5.1SP1.3

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 for Oracle RAC Release Notes

This document includes the following topics:

- [About this document](#)
- [Component product release notes](#)
- [About Veritas Storage Foundation for Oracle RAC](#)
- [About Symantec Operations Readiness Tools](#)
- [Important release information](#)
- [Changes introduced in version 5.1 Service Pack 1](#)
- [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 for Oracle RAC (SF Oracle RAC) version for Linux. Review this entire document before you install or upgrade SF Oracle RAC.

The information in the Release Notes supersedes the information provided in the product documents for SF Oracle RAC.

This is Document version: 5.1SP1.3 of the *Veritas Storage Foundation for Oracle RAC Release Notes*. Before you start, ensure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

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

Component product release notes

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

/product_name/docs

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

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 Service Pack 1)*.
- Storage Foundation (SF)
See *Veritas Storage Foundation Release Notes (5.1 Service Pack 1)*.
- Storage Foundation Cluster File System (5.1 Service Pack 1)
See *Veritas Storage Foundation Cluster File System Release Notes (5.1 Service Pack 1)*.

About Veritas Storage Foundation for Oracle RAC

Veritas Storage Foundation™ for Oracle® RAC (SF Oracle RAC) leverages proprietary storage management and high availability technologies to enable robust, manageable, and scalable deployment of Oracle RAC on UNIX platforms. The solution uses Veritas Cluster File System technology that provides the dual advantage of easy file system management as well as the use of familiar operating system tools and utilities in managing databases.

The solution stack comprises the Veritas Cluster Server (VCS), Veritas Cluster Volume Manager (CVM), Veritas Oracle Real Application Cluster Support (VRTSdbac), Veritas Oracle Disk Manager (VRTSodm), Veritas Cluster File System (CFS), and Veritas Storage Foundation, which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

Benefits of SF Oracle RAC

SF Oracle RAC provides the following benefits:

- Support for file system-based management. SF Oracle RAC provides a generic clustered file system technology for storing and managing Oracle data files as well as other application data.
- Support for high-availability of cluster interconnects.
The PrivNIC/MultiPrivNIC agents provide maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.
- Use of clustered file system and volume management technologies for placement of Oracle Cluster Registry (OCR) and voting disks. These technologies provide robust shared block interfaces for placement of OCR and voting disks. In the absence of SF Oracle RAC, separate LUNs need to be configured for OCR and voting disks.
- Support for a standardized approach toward application and database management. A single-vendor solution for the complete SF Oracle RAC software stack lets you devise a standardized approach toward application and database management. Further, administrators can apply existing expertise of Veritas technologies toward SF Oracle RAC.
- Increased availability and performance using dynamic multi-pathing (DMP). DMP provides wide storage array support for protection from failures and performance bottlenecks in the HBAs, SAN switches, and storage arrays.
- Easy administration and monitoring of SF Oracle RAC clusters from a single web console.
- Support for many types of applications and databases.
- Improved file system access times using Oracle Disk Manager (ODM).
- Ability to configure ASM disk groups over CVM volumes to take advantage of dynamic multi-pathing (DMP).
- Enhanced scalability and availability with access to multiple Oracle RAC instances per database in a cluster.
- Support for backup and recovery solutions using volume-level and file system-level snapshot technologies. SF Oracle RAC enables full volume-level

snapshots for off-host processing and file system-level snapshots for efficient backup and rollback.

- Ability to failover applications without downtime using clustered file system technology.
- Prevention of data corruption in split-brain scenarios with robust SCSI-3 Persistent Group Reservation (PGR) based I/O fencing or Coordination Point Server-based I/O fencing. The preferred fencing feature also enables you to specify how the fencing driver determines the surviving subcluster.
- Support for sharing all types of files, in addition to Oracle database files, across nodes.
- Fast disaster recovery with minimal downtime and interruption to users. Users can transition from a local high availability site to a wide-area disaster recovery environment with primary and secondary sites. If a node fails, clients that are attached to the failed node can reconnect to a surviving node and resume access to the shared database. Recovery after failure in the SF Oracle RAC environment is far quicker than recovery for a failover database.
- Verification of disaster recovery configuration using fire drill technology without affecting production systems.
- Support for a wide range of hardware replication technologies as well as block-level replication using VVR.
- Support for campus clusters with the following capabilities:
 - Consistent reattach with Site Awareness
 - Site aware reads with VxVM mirroring
 - Monitoring of Oracle resources
 - Protection against split-brain scenarios

About Symantec Operations Readiness Tools

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

- Collect, analyze, and report on server configurations across UNIX or Windows environments. You can use this data to do the following:
 - Assess whether your systems are ready to install or upgrade Symantec enterprise products

- Tune environmental parameters so you can increase performance, availability, and use
- Analyze your current deployment and identify the Symantec products and licenses you are using
- Upload configuration data to the SORT Web site, so you can share information with coworkers, managers, and Symantec Technical Support
- Compare your configurations to one another or to a standard build, so you can determine if a configuration has "drifted"
- Search for and download the latest product patches
- Get notifications about the latest updates for:
 - Patches
 - Hardware compatibility lists (HCLs)
 - Array Support Libraries (ASLs)
 - Array Policy Modules (APMs)
 - High availability agents
- Determine whether your Symantec enterprise product configurations conform to best practices
- Search and browse the latest product documentation
- Look up error code descriptions and solutions

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

To access SORT, go to:

<http://sort.symantec.com>

Important release information

- The latest product documentation is available on the Symantec Web site at: <http://www.symantec.com/business/support/overview.jsp?pid=15107>
- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website: <http://entsupport.symantec.com/docs/335001>
- For the latest patches available for this release, go to: <http://sort.symantec.com/>

Changes introduced in version 5.1 Service Pack 1

This section describes the new features and changes in version 5.1 SP1.

Changes related to the installation and upgrades

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 installer.

See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

Unencapsulation not required for some upgrade paths

Unencapsulation is no longer required for certain upgrade paths.

See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration 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 *Veritas Storage Foundation for Oracle RAC Installation and Configuration 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 *Veritas Storage Foundation for Oracle RAC Installation and Configuration 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 SF Oracle RAC configuration:

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

See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

Web-based installer supports configuring SF Oracle RAC cluster in secure mode

You can now configure the SF Oracle RAC cluster in secure mode using the Web-based installer.

See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration 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 *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

The Web-based installer enables you to install, configure, and uninstall

The Web-based installer has increased parity with the script-based installer. It now supports the ability to install, configure, and uninstall SF Oracle RAC. Note that the Web-based installer does not support Oracle 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.

Support for Oracle RAC 11g Release 2

SF Oracle RAC now supports Oracle RAC 11g Release 2.

Dynamic Storage Tiering is rebranded as SmartTier

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

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 SF Oracle RAC 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) SF Oracle RAC 5.1 SP1 removes this restriction. VxVM SF Oracle RAC 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.

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 commands from any node, even when the command changes the configuration of the shared disk group. You do not need to know which node is the master to issue the command. If you issue the command on the slave node, CVM ships the commands from the slave node to the master node. CVM then executes the command on the master node.

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

- The CVM protocol version must be at least 100.

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

```
# vxassist -g shareddg make sharedvol 20M ctlr:fscsi0
```
 - Commands that specify both a shared disk group and a private disk group. For example:

```
# vxdg destroy privatedg shareddg
```
 - Commands that include the defaults file as an argument. For example:

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

Changing the CVM master online

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

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

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

Changes related to 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
- 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 command for 5.1 SP1:

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

List of RPMs in SF Oracle RAC 5.1 SP1

[Table 1-1](#) lists the RPMs in SF Oracle RAC 5.1 SP1.

Table 1-1 List of RPMs

Name	Description
VRTSvlic	Symantec License Utilities
VRTSvxvm	Veritas Volume Manager binaries
VRTSfssdk	Veritas File System SDK - Manual Pages
VRTSvcs	Veritas Cluster Server
VRTSatClient	Symantec Product Authentication Service
VRTSatServer	Symantec Product Authentication Service
VRTSvcsdr	Veritas Cluster Server Disk Reservation Modules and Utilities by Symantec
VRTScavf	Veritas Cluster Server Agents for Storage Foundation Cluster File System
VRTSamf	Veritas Asynchronous Monitoring Framework by Symantec
VRTSvvmconv	Veritas Linux LVM to VxVM Converter
VRTSperl	Veritas Perl redistribution
VRTSaslapm	Volume Manager ASL/APM
VRTSllt	Veritas Low Latency Transport
VRTSspt	Veritas Software Support Tools
VRTSsfmh	Veritas Storage Foundation Managed Host
VRTSgab	Veritas Group Membership and Atomic Broadcast
VRTSob	Veritas File System Management Services Provider

Table 1-1 List of RPMs (*continued*)

Name	Description
VRTSvxfs	Veritas File System binaries
VRTSvxfen	Veritas I/O Fencing
VRTSscps	Veritas Cluster Server Coordination Point Server
VRTSvcssea	Veritas Cluster Server Enterprise Agents
VRTSgms	Veritas Group Messaging Services
VRTSvcsag	Veritas Cluster Server Bundled Agents
VRTSdbed	Veritas Storage Foundation Common Utilities for Databases
VRTSodm	Veritas Oracle Disk Manager
VRTSglm	Veritas Global Lock Manager
VRTSdbac	Veritas Oracle Real Application Cluster Support Package

Support for intelligent monitoring of VCS resources using IMF

VCS now supports intelligent resource monitoring in addition to poll-based monitoring. Intelligent Monitoring Framework (IMF) is an extension to the VCS agent framework. You can enable or disable the intelligent monitoring functionality of VCS agents as needed.

The benefits of intelligent monitoring over poll-based monitoring are as follows:

- Faster notification of resource state changes.
- Reduction in VCS system utilization which enables VCS to effectively monitor a large number of resources.

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

The following agents are IMF-aware in VCS 5.1 SP1:

- Mount
- Process
- Application
- Oracle
- Netlsnr

- CFSMount
- CVMVxconfig
- CFSfsckd

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 *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* and the *Veritas Storage Foundation for Oracle RAC Administrator's 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 *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* and the *Veritas Storage Foundation for Oracle RAC Administrator's 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 for CP server on AIX and HP-UX**
CP server now supports AIX and HP-UX in addition to Linux and Solaris operating systems.

Support to migrate between fencing modes when the cluster is running

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

See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for more details.

No longer supported

This section lists software versions and features that are no longer supported. Symantec advises customers to minimize the use of these features.

SF Oracle RAC does not support the following:

- **ASMInst agent**
The ASMInst agent is no longer supported in SF Oracle RAC environments. The ASM instances are managed by Oracle Clusterware.
- **Use of crossover cables**
Oracle does not support the use of crossover cables for cluster interconnects due to the possibility of data corruption and other software limitations.

Note: Crossover cables are however known to function without any issues in SF Oracle RAC. While the SF Oracle RAC Technical support team may continue to provide support on related issues for existing deployments, this support may be constrained in some respects as it is no longer a supported configuration by Oracle.

The use of crossover cables is discouraged for new deployments.

- **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`, `vxdmts_fstatsummary`, `dbdst_fiostat_collector`, `vxdmts_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.

Important preinstallation information

Before you install SF Oracle RAC, make sure you have reviewed the following information:

- Hardware compatibility list for information about supported hardware:
<http://entsupport.symantec.com/docs/330441>
- Disk storage array support information:
<http://entsupport.symantec.com/docs/283282>
- Latest information on support for Oracle database versions:
<http://www.symantec.com/docs/TECH44807>
- Oracle documentation for additional requirements pertaining to your version of Oracle.

Hardware requirements

Depending on the type of setup planned, make sure you meet the necessary hardware requirements.

For basic clusters See [Table 1-2](#) on page 24.

For campus clusters See [Table 1-3](#) on page 25.

Table 1-2 Hardware requirements for basic clusters

Item	Description
SF Oracle RAC systems	Two to sixteen systems with two or more CPUs at 2GHz or higher.
DVD drive	A DVD drive on one of the nodes in the cluster.
Disks	<p>SF Oracle RAC requires that all storage disks support SCSI-3 Persistent Reservations (PR).</p> <p>Note: The coordinator disk does not store data, so configure the disk as the smallest possible LUN on a disk array to avoid wasting space.</p>
Disk space	<p>You can evaluate your systems for available disk space by running the product installation program. Navigate to the product directory on the product disc and run the following command:</p> <pre data-bbox="538 722 995 748"># ./installsfrac -precheck node_name</pre> <p>For details on the additional space that is required for Oracle, see the Oracle documentation.</p>
RAM	<p>Each SF Oracle RAC system requires at least 2 GB.</p> <p>Symantec recommends additional amount of at least twice the Oracle SGA size.</p>
Swap space	See the Oracle Metalink document: 169706.1
Network links	<p>Two or more private links and one public link.</p> <p>Links must be 100BaseT or gigabit Ethernet directly linking each node to the other node to form a private network that handles direct inter-system communication. These links must be of the same type; you cannot mix 100BaseT and gigabit.</p> <p>Symantec recommends gigabit Ethernet using enterprise-class switches for the private links.</p>
Fiber Channel or SCSI host bus adapters	At least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

Table 1-3 lists the hardware requirements for campus clusters in addition to the basic cluster requirements.

Table 1-3 Hardware requirements for campus clusters

Item	Description
Storage	<ul style="list-style-type: none"> ■ The storage switch (to which each host on a site connects) must have access to storage arrays at all the sites. ■ Volumes must be mirrored with storage allocated from at least two sites. ■ DWDM links are recommended between sites for storage links. DWDM works at the physical layer and requires multiplexer and de-multiplexer devices. ■ The storage and networks must have redundant-loop access between each node and each storage array to prevent the links from becoming a single point of failure.
Network	<ul style="list-style-type: none"> ■ Oracle requires that all nodes use the IP addresses from the same subnet. ■ Symantec recommends two Network Interface Cards (NIC) per host for LLT heartbeats. Oracle Clusterware requires one private and one virtual IP for each host. ■ Symantec recommends a common cross-site physical infrastructure for storage and LLT private networks.
I/O fencing	I/O fencing requires placement of a third coordinator disk at a third site. The DWDM can be extended to the third site or the iSCSI LUN at the third site can be used as the third coordination point. Alternatively Coordination Point Server can be deployed at the third remote site as an arbitration point.

Supported Linux operating systems

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

The Veritas 5.1 Service Pack 1 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)
- 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 operating system patches

Oracle RAC requires the following patch on systems running SLES10 SP2:
kernel 2.6.16.60-0.39.3

Supported database software

For the latest information on supported Oracle database versions, see the following Technical Support TechNote:

<http://www.symantec.com/docs/TECH44807>

Note: SF Oracle RAC supports only 64-bit Oracle.

The following database versions are supported:

- Oracle RAC 11g Release 2

Additionally, see the Oracle documentation for patches that may be required by Oracle for each release.

I/O fencing requirements

Depending on whether you plan to configure disk-based fencing or server-based fencing, make sure that you meet the requirements for coordination points:

- Coordinator disks
See “[Coordinator disk requirements for I/O fencing](#)” on page 27.
- CP servers
See “[CP server requirements](#)” on page 27.

Coordinator disk requirements for I/O fencing

Make sure that the I/O fencing coordinator disks meet the following requirements:

- For disk-based I/O fencing, you must have three coordinator disks.
- The coordinator disks can be raw devices, DMP devices, or iSCSI devices.
- Each of the coordinator disks must use a physically separate disk or LUN.
Symantec recommends using the smallest possible LUNs for coordinator disks.
- Each of the coordinator disks should exist on a different disk array, if possible.
- The coordinator disks must support SCSI-3 persistent reservations.
- Symantec recommends using hardware-based mirroring for coordinator disks.
- Coordinator disks must not be used to store data or must not be included in disk groups that store user data.
- Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.

CP server requirements

SF Oracle RAC 5.1SP1 clusters (application clusters) support CP servers which are hosted on the following VCS and SFHA versions:

- VCS 5.1 or 5.1SP1 single-node cluster
CP server requires LLT and GAB to be configured on the single-node VCS cluster that hosts CP server. This requirement also applies to any single-node application cluster that uses server-based fencing.
- SFHA 5.1 or 5.1SP1 cluster

Warning: Before you upgrade CP server nodes to use VCS or SFHA 5.1SP1, you must upgrade all the application clusters that use this CP server to version 5.1SP1. Application clusters at version 5.1 cannot communicate with CP server that runs VCS or SFHA 5.1 SP1.

Make sure that you meet the basic hardware requirements for the VCS/SFHA cluster to host the CP server.

See the *Veritas Cluster Server Installation Guide* or the *Veritas Storage Foundation High Availability Installation Guide*.

Note: While Symantec recommends at least three coordination points for fencing, a single CP server as coordination point is a supported server-based fencing configuration. Such single CP server fencing configuration requires that the coordination point be a highly available CP server that is hosted on an SFHA cluster.

Make sure you meet the following additional CP server requirements which are covered in this section before you install and configure CP server:

- Hardware requirements
- Operating system requirements
- Networking requirements (and recommendations)
- Security requirements

[Table 1-4](#) lists additional requirements for hosting the CP server.

Table 1-4 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: <ul style="list-style-type: none"> ■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB) ■ 300 MB in /usr ■ 20 MB in /var
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the CP servers.
RAM	Each CP server requires at least 512 MB.
CP server to client node physical link	A secure TCP/IP connection is required to connect the CP servers to the SF Oracle RAC clusters (application clusters).

[Table 1-5](#) displays the CP server supported operating systems and versions. An application cluster can use a CP server that runs any of the following supported operating systems.

Table 1-5 CP server supported operating systems and versions

CP server	Operating system and version
CP server hosted on a VCS single-node cluster or on an SFHA cluster	<p>CP server supports any of the following operating systems:</p> <ul style="list-style-type: none"> ■ AIX 5.3 and 6.1 ■ HP-UX 11i v3 ■ Linux: <ul style="list-style-type: none"> ■ RHEL 5 ■ SLES 10 ■ SLES 11 ■ Solaris 9 and 10 <p>Review other details such as supported operating system levels and architecture for the supported operating systems.</p> <p>See the <i>Veritas Cluster Server Installation Guide</i> or the <i>Veritas Storage Foundation High Availability Installation Guide</i>.</p>

Following are the CP server networking requirements and recommendations:

- Symantec recommends that network access from the application clusters to the CP servers should be made highly-available and redundant. The network connections require either a secure LAN or VPN.
- The CP server uses the TCP/IP protocol to connect to and communicate with the application clusters by these network paths. The CP server listens for messages from the application clusters using TCP port 14250. This is the default port that can be changed during a CP server configuration.
- The CP server supports either Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with the application clusters. If the CP server is configured to use an IPv6 virtual IP address, then the application clusters should also be on the IPv6 network where the CP server is being hosted.
- When placing the CP servers within a specific network configuration, you must take into consideration the number of hops from the different application cluster nodes to the CP servers. As a best practice, Symantec recommends that the number of hops from the different application cluster nodes to the CP servers should be equal. This ensures that if an event occurs that results in an I/O fencing scenario, there is no bias in the race due to the number of hops between the nodes.

For secure communications between the SF Oracle RAC cluster and CP server, consider the following requirements and suggestions:

- In a secure communication environment, all CP servers that are used by the application cluster must be configured with security enabled. A configuration where the application cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.
- The CP server and application clusters should also use the same root broker. If the same root broker is not being used, then trust can be established between the cluster nodes and CP server for the secure communication. Trust can be established by the installer when configuring fencing.
- For non-secure communication between CP server and application clusters, there is no need to configure Symantec Product Authentication Service. In non-secure mode, authorization is still provided by CP server for the application cluster users. The authorization that is performed only ensures that authorized users can perform appropriate actions as per their user privileges on the CP server.

For information about establishing secure communications between the application cluster and CP server, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

Supported replication technologies for global clusters

SF Oracle RAC supports the following hardware-based replication and software-based replication technologies for global cluster configurations:

- | | |
|----------------------------|---|
| Hardware-based replication | <ul style="list-style-type: none">■ EMC SRDF■ Hitachi TrueCopy■ IBM Metro Mirror■ IBM SAN Volume Controller (SVC)■ EMC MirrorView |
| Software-based replication | <ul style="list-style-type: none">■ Veritas Volume Replicator■ Oracle Data Guard |

Fixed issues

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

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

See [“Documentation”](#) on page 53.

Issues fixed in SF Oracle RAC 5.1 SP1

[Table 1-6](#) lists the issues fixed in SF Oracle RAC 5.1 SP1.

Table 1-6 Fixed issues in SF Oracle RAC 5.1 SP1

Incident number	Description
1795151	Global group fails to come online on the DR site with a message that it is in the middle of a group operation.
1844422	The CSSD agent configuration fails if the OCR files are placed in a directory on CFS.
1855800	The SF Oracle RAC installer may fail to configure the CSSD resource if the <code>/etc/hosts</code> file contains commented IP address and host name entries.
1927920	The PrivNIC and MultiPrivNIC agents do not support MTU size settings in the <code>main.cf</code> configuration file.
2038617	The installation and configuration check "LLT links' speed and auto negotiation settings" fails when LLT is configured over UDP/TCP.
2042817	Oracle Clusterware fails to restart due to incorrect registration with VCSMM.
2045700	The SF Oracle RAC installer fails to validate the length of the disk group and volumes names used for OCR and voting disk.
2058424	The SF Oracle RAC installer relinks the Oracle database binaries only for the first node in the cluster.
2089351	The CSSD agent incorrectly reports OFFLINE even when one of the <code>cssd</code> , <code>crsd</code> , or <code>evmd</code> daemons is still running, causing the nodes to panic with the following message: <code>Oracle CRS failure. Rebooting for cluster integrity.</code>
2138574	In a node with 16 clusters, the PrivNIC agent fails to fail over the IP address for nodes with the NodeID value greater than 10.

LLT, GAB, and I/O fencing fixed issues

[Table 1-7](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-7 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.
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 -n option
1992560	[Fencing] The <code>vxfentsthdw</code> utility uses scp 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 1-8 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_vmclonedb -o recoverdb</code> for offhost fails for Oracle 10gr2 and prior versions
1810711	Flashsnap reverse resync command fails on offhost flashsnap cloning

Known issues

This section covers the known issues in this release.

For Oracle RAC issues:

See [“Oracle RAC issues”](#) on page 33.

For SF Oracle RAC issues:

See [“SF Oracle RAC issues”](#) on page 35.

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

See [“Documentation”](#) on page 53.

Oracle RAC issues

This section lists the known issues in Oracle RAC.

During installation or system startup, Oracle Grid Infrastructure may fail to start

After successful installation of Oracle RAC 11g Release 2 Grid Infrastructure, while executing the `root.sh` script, `ohasd` may fail to start. Similarly, during system startup, Oracle Grid Infrastructure may fail to start though the VCS engine

logs may indicate that the cssd resource started Oracle Grid Infrastructure successfully.

The following message may be displayed on running the `strace` command:

```
# /usr/bin/strace -ftt -p pid_of_ohasd.bin
14:05:33.527288 open("/var/tmp/.oracle/npohasd", O_WRONLY <unfinished ...>
```

For possible causes and workarounds, see the Oracle Metalink document: 1069182.1

Oracle Grid Infrastructure installation may fail with the SF Oracle RAC installer

When you run the `installsfrac -configure` command to install Oracle Grid Infrastructure for Oracle RAC 11g Release 2, the installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

Workaround: Export the `OUI_ARGS` environment variable, before you run the SF Oracle RAC installation program:

```
export OUI_ARGS=--ignoreInternalDriverError
```

For more information, see the Oracle Metalink document: 970166.1

The Oracle Interface Configuration (OIFCFG) command fails if private IP address is plumbed on a virtual interface

In Oracle RAC 11.2.0.1 installations, plumbing the private IP addresses on virtual interfaces cause the Oracle Interface Configuration (`oifcfg`) command to fail with the following error:

```
# oifcfg delif -global eth3/192.168.1.0:cluster_interconnect
PRIF-26: Error in update the profiles in the cluster
```

This issue may be observed during failover of IP addresses by the MultiPrivNIC agent. This is because the MultiPrivNIC agent plumbs the IP address on a virtual interface during failover. [2131345]

Oracle RAC 11g Release 2 installation on SLES 11 may fail if the RPM `libcap1-1.10-6.10.x86_64` is not present

Oracle RAC 11g Release 2 installations on SLES 11 fail if the RPM `libcap1-1.10-6.10.x86_64` is not present.

Workaround: Before installing Oracle RAC 11g Release 2, verify that the RPM `libcap1-1.10-6.10.x86_64.rpm` is present on the system:

```
# rpm -qa | grep libcap
libcap1-1.10-6.10
libcap2-2.11-2.15
```

If the RPM is not installed, install it using the instructions in the Oracle Metalink document: 952051.1

Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software

The Oracle Cluster Verification utility fails during the installation of the Oracle Grid Infrastructure software. If the failure indicates that the OCR and vote device locations are not shared, ignore the message.

Changing the Veritas agent for Oracle error handling

The Veritas agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file, `oraerror.dat`, which provides a list of Oracle errors and actions to address the errors.

For a description of the actions:

See the *Veritas High Availability Agent for Oracle Installation and Configuration Guide*.

Currently, the file specifies the NOFAILOVER action for the following Oracle errors: ORA-00061, ORA-02726, ORA-6108, ORA-06114

The NOFAILOVER action means that the agent sets the state of the resource to OFFLINE and freezes the service group. If you want to change this behavior, you can stop the agent, edit `oraerror.dat`, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

SF Oracle RAC issues

This section lists the known issues in SF Oracle RAC for this release.

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.

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.

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 VRTSaclib.
- Uninstall: Ignore VRTSaclib.

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.

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`

Incorrect ownership assigned to the parent directory of ORACLE_BASE causes Oracle Clusterware/Grid Infrastructure installations to fail

When you use the SF Oracle RAC installation program to install Oracle Clusterware/Grid Infrastructure, the ownership of the parent directory of `ORACLE_BASE`/`GRID_BASE` that is created by the installation program is incorrectly set to root. This causes the Oracle Universal Installer to display errors when it creates the `oraInventory` directory as the `oraInventory` directory must be created on behalf of the `oracle` user (Oracle RAC 10g Release 2/Oracle RAC 11g Release 1) or `grid` user (Oracle RAC 11g Release 2).

Workaround:

1. Log into each node in the cluster as the root user.
2. Perform the following operations:
 - If you have not yet installed Oracle Clusterware/Grid Infrastructure, create the directory and set the correct ownership as follows before you invoke the installation program:

```
# mkdir -p oracle_base
# chown user_name:oraInventory_group_name
  oracle_base/..
```

where:

oracle_base is the name of the Oracle base directory.

user_name is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).

oraInventory_group_name is the name of the oraInventory group.

Complete the other preparatory tasks before you invoke the installation program. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

- If you faced this issue during the installation of Oracle Clusterware/Grid Infrastructure, open another terminal session, and modify the ownership of the directory on all nodes in the cluster as follows:

```
# chown user_name:oraInventory_group_name
  oracle_base/..
```

where:

oracle_base is the name of the Oracle base directory.

user_name is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).

oraInventory_group_name is the name of the oraInventory group.

Return to the former session and proceed with the installation.

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.

Deporting issues with shared disk groups

If you manually deport a shared disk group, the CVMVolDg agent does not automatically reimport it as a shared disk group. You must manually reimport it as a shared disk group.

SCSI reservation errors during bootup

If you reboot a node of an SF Oracle RAC cluster, SCSI reservation errors may be observed during bootup. [255515]

For example:

```
Nov 23 13:18:28 galaxy kernel: scsi3 (0,0,6) : RESERVATION CONFLICT
```

This message is printed for each disk that is a member of any shared disk group which is protected by SCSI-3 I/O fencing. The message may be safely ignored.

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 SF Oracle RAC 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 SF Oracle RAC, 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.

Network interfaces change their names after restart

On SUSE systems, network interfaces change their names after restarting even with `HOTPLUG_PCI_QUEUE_NIC_EVENTS=yes` and `MANDATORY_DEVICES="..."` set.

Workaround: Set the interface names in the following files:

```
SLES 10 /etc/udev/rules.d/30-net_persistent_names.rules
```

SLES 11 `/etc/udev/rules.d/70-persistent-net.rules`

Kernel warning messages when Veritas modules load

For SLES 10, a warning message resembling the following may be displayed in the console or the system log when Veritas modules are loaded into the kernel.

```
Warning: module not supported by Novell, setting U taint flag.  
module license 'Proprietary. Send bug reports to  
support@symantec.com' taints kernel.
```

These warning messages are displayed because the Veritas modules are proprietary. They can safely be ignored.

DBED features are not integrated with GCO

DBED features are not integrated with Global Cluster Option (GCO). After GCO migration, be aware that DBED features will not be functional. [1241070]

Issue with format of the last 8-bit number in private IP addresses

The PrivNIC/MultiPrivNIC resources fault if the private IP addresses have a leading 0 in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1. or X.0X.X.1 or 0X.X.X.1, where X is an octet of the IP address. [1164506]

When you configure private IP addresses for Oracle Clusterware, ensure that the IP addresses have a format as displayed in the following two-node example:

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

Confirm the correct format by viewing the PrivNIC or MultiPrivNIC resource in the `/etc/VRTSvcs/conf/config/main.cf` file.

When master node loses access to complete storage, detached sites remain in RECOVER state even after reattaching and recovering the sites

In a campus cluster environment, if the master node loses access to complete storage, all but one of the sites is detached and the DCO volumes may get detached if the `dgfailpolicy` is set to `dgdisable`. If the detached sites are reattached and recovered, the site still remains in RECOVER state. [1828142]

Workaround: Change the status of the site as described in the following procedure to resolve the issue.

To change the status of the site

1 Log onto the CVM master node.

2 Reattach the detached sites:

```
# vxdg -g dg_name reattachsite site_name
```

The site remains in RECOVER state.

3 Restore DCO volumes by unpreparing and preparing the volumes.

Unprepare the volumes:

```
# vxsnap -g dg_name -f unprepare vol_name
```

Prepare the volumes:

```
# vxsnap -g dg_name prepare vol_name dnl=on
```

4 Reattach the detached sites:

```
# vxdg -g dg_name reattachsite site_name
```

5 Verify that the state of the detached sites is now ACTIVE:

```
# vxprint
```

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.

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.

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.

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
```

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.

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
```

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 00000000000000082 0000000000000084c 0000000000000001
ffff810faf539de8 00000000000000007 ffff810fc2a130c0 ffff810138ee8100
000019f130082599 00000000000018572 ffff810fc2a132a8 00000001f76c3d63
Call Trace:
[<fffffffff88ee3690>] :gab:gab_linux_sv_wait+0x53/0x68
[<fffffffff8008e68d>] default_wake_function+0x0/0xe
[<fffffffff88ecd4c8>] :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 *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for more information on preferred fencing.

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

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

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

Reconfiguring SF Oracle RAC 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 SF Oracle RAC 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.

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
```

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

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

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

When upgrading from SF Oracle RAC version 5.0 to SF Oracle RAC 5.1SP1 the S*vxdbs3 startup script is renamed to NO_S*vxdbs3. The S*vxdbs3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbs3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbs3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround

Before running `sfua_rept_migrate`, rename the startup script `NO_S*vxdbs3` to `S*vxdbs3`.

Database fails over during Flashsnap operations (1469310)

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

Workaround

There is no workaround for this issue.

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

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

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

Workaround

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

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

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

```
# vxdg join snapshot_disk_group_name  
           primary_disk_group_name
```

- 2 Start all the volumes in primary disk group.

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

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

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

Repeat this step for all the volumes.

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

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

Workaround

There is no workaround for this issue.

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

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

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

Software limitations

This section covers the software limitations of this release.

See the corresponding Release Notes for a complete list of software limitations related to that component or product.

See [“Documentation”](#) on page 53.

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

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

```
# vxrvg -g diskgroup stop rvg
```
- 4 Disassociate the volumes from the RVG. Enter the following:

```
# vxvol -g diskgroup dis vol
```
- 5 Resize the volumes. In this example, the volume is increased to 10 GB. Enter the following:

```
# vxassist -g diskgroup growto vol 10G
```
- 6 Associate the data volumes to the RVG. Enter the following:

```
# vxvol -g diskgroup assoc rvg vol
```
- 7 Start the RVG. Enter the following:

```
# vxrvg -g diskgroup start rvg
```
- 8 Resume or start the applications.

Health checks may fail on clusters that have more than 10 nodes

If there are more than 10 nodes in a cluster, the health check may fail with the following error:

```
vxgettext ERROR V-33-1000-10038  
Arguments exceed the maximum limit of 10
```

The health check script uses the `vxgettext` command, which does not support more than 10 arguments.[2142234]

Cached ODM not supported in SF Oracle RAC environments

Cached ODM is not supported for files on Veritas local file systems and on Cluster File System.

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.

SELinux supported in disabled and permissive modes only

SELinux (Security Enhanced Linux) is supported only in "Disabled" and "Permissive" modes. After you configure SELinux in "Permissive" mode, you may see a few messages in the system log. You may ignore these messages.

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: 5.1SP1.3 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 53.

See ["About Symantec Operations Readiness Tools"](#) on page 10.

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.

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 1-9](#) lists the documentation for Veritas Storage Foundation for Oracle RAC.

Table 1-9 Veritas Storage Foundation for Oracle RAC documentation

Document title	File name
<i>Veritas Storage Foundation for Oracle RAC Release Notes</i>	sfrac_notes_51SP1_lin.pdf
<i>Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide</i>	sfrac_install_51SP1_lin.pdf
<i>Veritas Storage Foundation for Oracle RAC Administrator's Guide</i>	sfrac_admin_51SP1_lin.pdf
<i>Veritas Storage Foundation: Storage and Availability Management for Oracle Databases</i>	sf_adv_ora_51SP1_lin.pdf

[Table 1-10](#) lists the documents for Veritas Cluster Server.

Table 1-10 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_51sp1.pdf

Table 1-10 Veritas Cluster Server documentation (*continued*)

Title	File name
<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 1-11](#) lists the documentation for Veritas Storage Foundation.

Table 1-11 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
<i>Veritas Storage Foundation Cluster File System Release Notes</i>	sfdfs_notes_51SP1_lin.pdf
<i>Veritas Storage Foundation Cluster File System Administrator's Guide</i>	sfdfs_admin_51SP1_lin.pdf
<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 1-12](#) lists the documentation for Veritas Volume Manager and Veritas File System.

Table 1-12 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

Table 1-12 Veritas Volume Manager and Veritas File System documentation (continued)

Document title	File name
<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 1-13](#) lists the documentation for Veritas Volume Replicator.

Table 1-13 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 1-14](#) lists the documentation for Symantec Product Authentication Service (AT).

Table 1-14 Symantec Product Authentication Service documentation

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