Veritas Storage Foundation™ for Oracle® RAC Release Notes

HP-UX 11i v3

5.1 Service Pack 1

Symantec

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: 5.1SP1.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

docs@symantec.com

About Symantec Connect

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

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

Storage Foundation for Oracle RAC Release Notes

This document includes the following topics:

About this document

This document provides important information about Veritas Storage Foundation for Oracle RAC (SF Oracle RAC) version 5.1 Service Pack 1 for HP-UX 11i v3. 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: 5.1SP1.0 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.
  For Oracle RAC 10g Release 2:
  The combination of LMX/LLT protocols and the PrivNIC/MultiPrivNIC agents provides maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.
  For Oracle RAC 11g Release 1/Oracle RAC 11g Release 2:
  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 and raw 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://www.symantec.com/docs/TECH144835
- 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.

The VRTScutil and VRTSacclib depots are no longer in use

For all high availability products, the VRTScutil and VRTSacclib depots 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 3.5 onward.

Changes in the SF Oracle RAC installer

This section lists the changes in the SF Oracle RAC installer.

Option to start and stop SF Oracle RAC

You can use the `-start` and `-stop` options with the SF Oracle RAC installer to start and stop SF Oracle RAC.

Installer resilience

If the installation or configuration of SF Oracle RAC fails, the installer discovers the presence of an installer instance when you run the installer again. The installer provides an option to resume the installation.

Adding nodes to a cluster

The SF Oracle RAC installer now provides the following capabilities for adding nodes to a cluster:

- Configures SF Oracle RAC and adds nodes to an existing SF Oracle RAC cluster.
- Performs the following Oracle RAC pre-installation tasks:
  - Creates Oracle user and groups on the new node.
  - Configures the private IP addresses and the PrivNIC resource for Oracle Clusterware (only if the IP addresses on the existing cluster are configured as PrivNIC resources).
  - Configures private IP addresses and the MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC (only if the IP addresses on the existing cluster are configured as MultiPrivNIC resources).
  - If the CFSMount and CVMVolDg resources for OCR and voting disk are configured under the cvm service group, the installer brings them online after adding the node to the cluster.
  - Starts the cvm group on the new node

Options to install minimum, recommended, or all SF Oracle RAC depots

The installer now provides the following options for installing SF Oracle RAC:

- Minimal
  Install only the basic functionality of SF Oracle RAC.

To view the list of depots for this option, use `-minpkgs` option with the installer script.

- Recommended
  Installs SF Oracle RAC without optional depots.
  To view the list of depots for this option, use `-recpkgs` option with the installer script.
- All
  Installs all SF Oracle RAC depots.
  To view the list of depots for this option, use `-allpkgs` option with the installer script.

Cross-product upgrades

The installer does not support direct upgrade from a previous SFCFS or SFHA version to SF Oracle RAC 5.1 SP1. You must upgrade SFCFS or SFHA to version 5.1 Service Pack 1, then install SF Oracle RAC 5.1 SP1.

Simplified SF Oracle RAC configuration

SF Oracle RAC configuration is faster and simpler with reduced manual interventions. The Cluster Volume Manager and Veritas Volume Replicator configuration is automatically performed by the installer.

SF Oracle RAC cluster verification checks

The option **SF Oracle RAC Installation and Configuration Checks** is introduced in the SF Oracle RAC installer. This option enables you to verify the cluster during various stages of SF Oracle RAC deployment.

For new installations, these checks can be performed before and after Oracle RAC installation.

For existing deployments, you can perform these checks after any of the following activities to check the sanity of the cluster:

- Upgrading the operating system
- Applying operating system patch updates
- Applying Oracle patch updates
- Applying SF Oracle RAC patch updates
- Adding a node to an SF Oracle RAC cluster
- Removing a node from an SF Oracle RAC cluster

- Updating the network configuration

Installation simulator

You can use the option `-makeresponsefile` with the installer script to simulate an installation, configuration, or uninstallation activity.

Oracle RAC installation

The SF Oracle RAC installer includes several enhancements for performing Oracle RAC pre-installation tasks.

Table 1-1 lists the changes and new features in the SF Oracle RAC installer for Oracle tasks.

Table 1-1 Changes in the SF Oracle RAC installer for Oracle tasks

Oracle task	Description
Oracle user and group creation	■ The installer checks the Oracle user and group IDs in use and suggests unused values.
OCR and voting disk configuration	■ The installer creates CVM volume mirrors for OCR and voting disk ■ The installer creates the OCR and voting disk volumes and sets the ownership ■ The installer starts the volumes ■ If you choose to create the storage on CFS, the installer creates the mount point, mounts it on all the nodes, and sets the ownership for the CFS mount point ■ The installer adds the volume and mount point resources to the VCS configuration so that the resources are brought online automatically when the node starts. ■ The installer brings the resources online
PrivNIC and MultiPrivNIC configuration	■ The installer no longer asks for the NIC information for the IP addresses, but uses all the available LLT links in the cluster for high availability purposes. ■ The installer provides the option to automatically update the `/etc/hosts` file during the private network configuration to include IP addresses used by PrivNIC or MultiPrivNIC resources.

Table 1-1 Changes in the SF Oracle RAC installer for Oracle tasks *(continued)*

Oracle task	Description
CSSD agent configuration	A new option **Configure CSSD agent** is introduced in the SF Oracle RAC installer that enables you to configure the CSSD agent after installing Oracle RAC. During the configuration, the installer creates the CSSD resource in the VCS configuration file and sets the OCR, voting disk, and IP address dependencies, such that they are online before Oracle Clusterware starts.

Silent installation of Oracle RAC using response files

SF Oracle RAC supports the completion of the following Oracle RAC tasks using SF Oracle RAC response files:

- Creating Oracle user and groups
- Creating storage for OCR and voting disk
- Configuration of private IP addresses and corresponding PrivNIC/MultiPrivNIC resources
- Installation of Oracle Clusterware and Oracle database by specifying the path of the Oracle RAC response files in the appropriate response file variable definitions. For more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.
- Configuration of the CSSD resource
- Relinking the SF Oracle RAC libraries with Oracle RAC

Note: You can now use the SF Oracle RAC response file in tandem with the Oracle RAC response files for Oracle Clusterware and Oracle database to perform silent end-to-end installation of SF Oracle RAC and Oracle RAC.

Support to enable rolling upgrades to future releases

This release of SF Oracle RAC establishes the necessary framework to support rolling upgrades in future releases. The framework enables you to use the installer to perform a phased upgrade of your cluster with minimal application and server downtime during the upgrade process.

I/O fencing configuration

You can use the `-fencing` option with the `installsfrac` program to configure I/O fencing.

Based on the fencing mechanism you want to use in the cluster, the installer provides the following options to configure I/O fencing:

- Disk-based I/O fencing - when you want to use disks as coordination points
- Server-based I/O fencing - when you want to use at least one CP server as coordination point

For more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

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.

List of depots in SF Oracle RAC 5.1 SP1

Table 1-2 lists the depots in SF Oracle RAC 5.1 SP1.

Table 1-2 List of depots

Name	Description
VRTSvlic	Symantec License Utilities
VRTSvxvm	Veritas Volume Manager binaries
VRTSfssdk	Veritas File System SDK - Manual Pages
VRTSvcs	Veritas Cluster Server
VRTSat	Symantec Product Authentication Service

Table 1-2 List of depots *(continued)*

Name	Description
VRTScavf	Veritas Cluster Server Agents for Storage Foundation Cluster File System
VRTSamf	Veritas Asynchronous Monitoring Framework by Symantec
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
VRTSvxfs	Veritas File System binaries
VRTSvxfen	Veritas I/O Fencing
VRTScps	Veritas Cluster Server Coordination Point Server
VRTSvcsea	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

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.

Dynamic Storage Tiering is rebranded as SmartTier

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

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.

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
- CVMVxconfigd
- CFSfsckd

Support for 16 nodes in a cluster

SF Oracle RAC now supports 16 nodes in a cluster.

Support for Oracle RAC 11g Release 2

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

Changes related to Storage Foundation for Databases (SFDB) tools

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

- Storage Foundation for Cluster File (HA) System support
- Multiple disk group support for FlashSnap
- Sub-file storage tiering is supported for SmartTier for Oracle (previously known as Database Dynamic Storage Tiering)
- SQLite repository
- Oracle Enterprise Manager (OEM) Plugin
- Oracle 11gR2 support

- HP-UX Service Guard support for replicated environments

Commands which have changed:

- `sfua_db_config` functionality is changed: this command is no longer needed to create a SFDB repository. The functionality of `sfua_db_config` is now used to set user and group access to various SFDB directories.
- Use the `dbed_update` command to create a new SQLite SFDB repository.
- `sfua_rept_adm` was used in 5.0 to perform repository backup and restore and this command will be obsolete in 5.1 SP1.
- The `sfua_rept_util` command is used to perform SQLite repository backup and restore.
- The Database Dynamic Storage Tiering (DBDST) feature is rebranded as SmartTier for Oracle and includes expanded functionality to support management of sub-file objects.
- Cached ODM: functionality is expanded with Cache ODM Manager

Commands which continue to be supported:

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

New commands:

- SmartTier for Oracle: commands added to support storage tiering of sub-file objects: `dbdst_obj_view`, `dbdst_obj_move`
- Cached ODM: command added to support Cached ODM Manager: dbed_codm_adm

Support for server-based fencing

The Coordination Point Server provides an alternative fencing mechanism that integrates with the existing VCS I/O fencing module.

The Coordination Point server (CPS) is a software solution running on a remote system or cluster that provides arbitration functionality by allowing client cluster nodes to perform the following tasks:

- Self-register to become a member of an active client cluster with access to the data drives
- Check which other nodes are registered as members of this active client cluster
- Self-unregister from this active client cluster
- Forcefully unregister other nodes (preempt) as members of this active client cluster

Multiple client clusters running different operating systems are able to simultaneously access the CP server. TCP/IP based communication is used between the CP server and client clusters.

Enhanced sample configuration files

This release of SF Oracle RAC includes a number of sample configuration files that illustrate various deployment scenarios. Each sample file contains a high-level description of the deployment scenario. The files are located at `/etc/VRTSvcs/conf/sample_rac/`.

Support for load balancing in MultiPrivNIC

The MultiPrivNIC agent supports load-balanced fail over if multiple links are identified as failover targets. A new attribute `UseLoadBalance` is introduced to support load balanced fail over of links. In the event that a preferred link goes down, the IP address is failed over to a private link on which maximum number of peer nodes are visible. If multiple links see maximum nodes and if load-balancing is enabled, the agent considers the current traffic on all devices and calculates a "winner" device with lower traffic. If load balancing is not enabled, the IP address is failed over to the link with the lower network ID.

New CRSResource agent

The CRSResource agent provides an alternative mechanism for monitoring the Oracle database in the absence of the VCS Oracle agent. It is useful in scenarios where the database is not managed by VCS and the applications need to be started using VCS after Oracle Clusterware starts the database. It checks the status of the

Oracle Clusterware resources, which include the Oracle database instance, the listener, and the virtual IP address (VIP). The agent supports multiple database configurations and ensures that the Oracle database is online and available to an application when it starts.

Support for health checks on SF Oracle RAC cluster

SF Oracle RAC provides a health check utility that evaluates the components and configuration in an SF Oracle RAC cluster. It gathers real-time operational information on cluster components and reports any deviations or warnings that may degrade performance or cause cluster issues.

The utility evaluates the health of the following components in the cluster:

- Low Latency Transport (LLT)
- LLT Multiplexer (LMX)
- I/O fencing
- Oracle Clusterware
- PrivNIC

Utility to view LMX bandwidth used for database traffic

The `lmxdbstat` utility can be used to determine the LMX bandwidth used for database traffic for each database. The utility is located at `/sbin/lmxdbstat`.

The utility reports the following information:

- Status of the LMX protocol
- The LMX port and buffer traffic received and transmitted at periodic intervals in packets and kilobytes for each database instance.
- The LMX port and buffer traffic received and transmitted at periodic intervals in packets and kilobytes for a database process.

For more information:

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

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

Enhancements to VRTSexplorer utility

The VRTSexplorer utility includes enhancements to collect more logs, including Oracle Clusterware and Oracle database log information.

Changes in the SF Oracle RAC documentation

This section provides a high-level list of the key changes in the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*:

- The sections "Oracle RAC 10g installation in an SF Oracle RAC environment" and "Oracle RAC 11g installation in an SF Oracle RAC environment" are now consolidated into a single section "Installing and configuring Oracle RAC".
  The Oracle RAC chapters containing preparatory and installation instructions are common for Oracle RAC 10g and Oracle RAC 11g. Differences in instructions, wherever applicable, are indicated in the relevant steps.
  The chapter formerly titled "Installing Oracle RAC 11g manually" is now placed in the chapter "Installing Oracle RAC" and is re-titled as follows:
  - Installing Oracle Clusterware using the Oracle Universal Installer
  - Installing the Oracle RAC database using the Oracle Universal Installer
- The minimum downtime upgrade procedure is revised for this release and is now termed as "phased upgrade".
- The list of SF Oracle RAC sample files that were reproduced in the appendix of the guide is now replaced by service group configuration illustrations of various scenarios and corresponding descriptions. The actual files can be referenced at `/etc/VRTSvcs/conf/sample_rac/`.

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:

- Dissimilar version upgrades of SF Oracle RAC components
  For example, if you have SF Oracle RAC 4.1 installed with Oracle RAC 9i Release 2, you can not upgrade only VCS to version 5.1 Service Pack 1.
- Option `hawizard -rac` for configuring service groups
- 32-bit Oracle architectures
- Oracle RAC 9i, Oracle RAC 10g Release 1
- Storage Checkpoint commands for checking and saving the configuration of the Oracle RAC database: `dbed_checkconfig, dbed_saveconfig`
- 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`, `vxdbts_fstatsummary`, `dbdst_fiostat_collector`, `vxdbts_get_datafile_stats`
- `dbed_saveconfig`, `dbed_checkconfig`
- `dbed_ckptplan`, `dbed_ckptpolicy`
- `qio_convertdbfiles -f` option which is used to check for file fragmentation
- `dbed_scheduler`
- `sfua_rept_migrate` with `-r` and `-f` options

System requirements

This section describes the system requirements for this release.

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 1-3 on page 29.

For campus clusters See Table 1-4 on page 30.

Table 1-3 Hardware requirements for basic clusters

Item	Description
SF Oracle RAC systems	Two to sixteen HP-UX systems that are connected to the public network.
DVD drive	A DVD drive on one of the nodes in the cluster.
Disks	SF Oracle RAC requires that all storage disks support SCSI-3 Persistent Reservations (PR). **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. DMP devices are mandatory for use as coordinator disks.
Disk space	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: `# ./installsfrac -precheck node_name` For details on the additional space that is required for Oracle, see the Oracle documentation.

Table 1-3 Hardware requirements for basic clusters *(continued)*

Item	Description
RAM	Each SF Oracle RAC system requires at least 2 GB. Symantec recommends additional amount of at least twice the Oracle SGA size.
Swap space	See the Oracle Metalink document: 169706.1
Network links	Two or more private links and one public link. 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. Symantec recommends gigabit Ethernet using enterprise-class switches for the private links. You can also configure aggregated interfaces.
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 1-4 lists the hardware requirements for campus clusters in addition to the basic cluster requirements.

Table 1-4 Hardware requirements for campus clusters

Item	Description
Storage	■ 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	■ 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.

Table 1-4 Hardware requirements for campus clusters *(continued)*

Item	Description
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 HP-UX operating systems

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

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

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

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

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

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

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 10g Release 2
- Oracle RAC 11g Release 1
- 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 32.
- CP servers
  See "CP server requirements" on page 32.

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 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 1-5 lists additional requirements for hosting the CP server.

Table 1-5 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: ■ 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.

Table 1-5 CP server hardware requirements *(continued)*

Hardware required	Description
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 1-6 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 1-6 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	CP server supports any of the following operating systems:
	■ AIX 5.3 and 6.1
	■ HP-UX 11i v3
	■ Linux:
	  ■ RHEL 5
	  ■ SLES 10
	  ■ SLES 11
	■ Solaris 9 and 10
	Review other details such as supported operating system levels and architecture for the supported operating systems.
	See the *Veritas Cluster Server Installation Guide* or the *Veritas Storage Foundation High Availability Installation Guide*.

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	■ EMC SRDF ■ Hitachi TrueCopy ■ IBM Metro Mirror ■ IBM SAN Volume Controller (SVC) ■ EMC MirrorView
Software-based replication	■ 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 56.

Issues fixed in SF Oracle RAC 5.1 SP1

Table 1-7 lists the issues fixed in SF Oracle RAC 5.1 SP1.

Table 1-7 Fixed issues in SF Oracle RAC 5.1 SP1

Incident number	Description
1439223	The `lmxpollport` function times out with an incorrect timeout value on systems that have been running for more than 410 days. The issue was caused by the lbolt variable, which resets after 410 days.
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 `/etc/hosts` file contains commented IP address and host name entries.
1879412	The Low Latency Transport Multiplexer (LMX) module may cause the system to panic with the following message: `kernel heap corruption detected` Incorrect manipulation of the request queue corrupts the memory and causes the system to panic. When the last request is removed from the queue, the queue pointers are not updated correctly.
1927920	The PrivNIC and MultiPrivNIC agents do not support MTU size settings in the `main.cf` configuration file.
1938799	When LMX registers with LLT, LLT calls the `lmxlltxcanput` function before delivering any packet to LMX, causing performance overheads.
2038617	The installation and configuration check "LLT links' speed and auto negotiation settings" fails when LLT is configured over UDP/TCP.

Table 1-7 Fixed issues in SF Oracle RAC 5.1 SP1 *(continued)*

Incident number	Description
2042817	Oracle Clusterware fails to restart due to incorrect regisration 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.
2089351	The CSSD agent incorrectly reports OFFLINE even when one of the cssd, crsd, or evmd daemons is still running, causing the nodes to panic with the following message: `Oracle CRS failure.  Rebooting for cluster integrity.`
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.
2237829	When a node panics on a heavily loaded cluster, the Oracle instance on another node crashes.

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 dbed_vmchecksnap command may fail
1399393	Clone command fails on an Oracle RAC database
1736516	Clone command fails for instant checkpoint on Logical Standby database
1789290	dbed_vmclonedb -o recoverdb for offhost fails for Oracle 10gr2 and prior versions
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 38.

For SF Oracle RAC issues:

See "SF Oracle RAC issues" on page 40.

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

See "Documentation" on page 56.

Oracle RAC issues

This section lists the known issues in Oracle RAC.

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

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 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.

Oracle VIP Configuration Assistant fails with an error message

During Oracle RAC 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

```
The given interface(s), "lan0" is not public.
Public interfaces should be used to configure virtual IPs.
```

This message appears only when the VIP is not from the regular public IP range (for example, 200.). [1182220]

Workaround: Invoke the `vipca` utility manually as the superuser.

```
# export DISPLAY=nebula:0.0
# $CRS_HOME/bin/vipca
```

Oracle Cluster Verification utility displays a warning message

During the final stage of Oracle RAC 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification utility.

For example:

```
Utility
============================================================
OUI-25031: Some of the configuration assistants failed. It is
strongly recommended that you retry the configuration
assistants at this time. Not successfully running any "
Recommended" assistants means your system will not be correctly
configured.
1. Check the Details panel on the Configuration Assistant Screen
to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button
to retry them.
============================================================
```

Workaround: You may safely ignore this message if the cluster is operating satisfactorily.

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.

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

The Web-based installer fails to run.

Workarounds:

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

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

To start the Web-based installer workaround

1. Create a mount point.

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

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

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

3. Mount the disc to the mount point.

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

4. Create a temporary installation directory.

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

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

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

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

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

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

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

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

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

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

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

```
VRTSvlic depot not installed on system_name
```

Workaround:

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

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

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

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

EULA changes (2161557)

The locations for all EULAs have changed.

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

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

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

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

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

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

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

The VRTSacclib depot is deprecated (2032052)

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

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

Installer assigns duplicate node ID during `-addnode` procedure

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

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

Miscalculated file set usage (2123429)

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

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

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

Workaround

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

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

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

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

Workaround

Create a resource dependency between the various CFSmount resources.

VCS fails to go to the running state on HP-UX 11.31 with March 2011 release

Due to a regression caused by the patch PHKL_41700 (QXCR1001078659) that went into HP-UX 11.31 March 2011 release, the select() call takes long time to return from 'timeout sleep'. Due to this, _had misses the heartbeat with GAB resulting in SIGABRT by GAB. [2287383]

Workaround: You must tune 'hires_timeout_enable' kernel parameter to 1 before starting the cluster. Run the following command to set this variable to 1:

```
# kctune hires_timeout_enable=1
```

Note: HP is likely to deliver the resolution for this issue via PHKL_41967 patch post the March 2011 release.

Verification for VRTSat package or patch returns errors

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

Workaround: This message may be safely ignored.

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.

Stopping cluster nodes configured with I/O fencing

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect or “split brain.”

For more information, see *Veritas Cluster Server User’s Guide*.

I/O fencing uses SCSI-3 Persistent Reservation keys to implement data protection. The software places keys on I/O fencing coordinator and data disks. The 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 disks 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 and data disks. Depending on the order of reboot and subsequent startup events, the cluster might warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown command instead of the reboot command to perform a graceful reboot for systems.

```
# /usr/sbin/shutdown -r now
```

Stopping VCS does not unregister port f from GAB membership

In an SF Oracle RAC cluster with all the CFS resources under VCS control, when you stop VCS, all the CFS resources must go down cleanly and CFS must unregister port f from GAB membership. Oracle RAC 10g Clusterware does not clean up all

its processes when it is stopped. Now, when you stop VCS, all the CFS resources go down. However, due to the left over Oracle processes, CFS does not unregister port f from GAB membership.

Workaround: Perform the following steps to bring down port f.

To bring down port f

1 Kill all the Oracle processes.

```
# kill -9 'ps -u oracle|awk '{print $1}'`
```

2 Verify that all CFS file systems are unmounted.

```
# mount | grep cluster
```

3 Unregister port f from GAB membership.

```
# fsclustadm cfsdeinit
```

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.

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

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

if the node is attempting to join the cluster after you issue the `vxdctl upgrade` command to upgrade the CVM cluster.

Work-around:

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

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 drl=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
   ```

Issues related to LLT

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

LLT port stats sometimes shows recvcnt larger than recvbytes

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

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

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

Issues related to I/O fencing

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

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

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

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.

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.

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.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0MP2 to 5.1SP1 (2003131)

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

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

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

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

When upgrading from SF Oracle RAC version 5.0 or 5.0.1 to SF Oracle RAC 5.1SP1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File:  is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround

Before running `sfua_rept_migrate`, rename the startup script NO_S*vxdbms3 to S*vxdbms3.

Relinking ODM after upgrading from 5.0.x

The `VRTSodm` library path has changed from /opt/VRTSodm/lib/libodm.sl to /opt/VRTSodm/lib/libodm.so.

After upgrading to 5.1 Service Pack 1 from 5.0.x you must update the ODM link for your database to the new `VRTSodm` library path /opt/VRTSodm/lib/libodm.so.

Upgrading in an HP Serviceguard environment (2116455)

When upgrading SFDB to 5.1SP1 from the previous release in an HP Serviceguard environment, first verify that the `cmviewcl` command can be executed by a non-root user. This permission change must be done before executing SFDB upgrade commands.

Using SFDB tools after upgrading Oracle to 11.2.0.2 (2203228)

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

Database fails over during Flashsnap operations (1469310)

In an 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, ues the following steps to reattach the volumes.

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

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

   ```
   # vxdg join snapshop_disk_group_name
         primary_disk_group_name
   ```

2. Start all the volumes in primary disk group.

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

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

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

   Repeat this step for all the volumes.

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

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

Workaround

There is no workaround for this issue.

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

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

Workaround: Set MonitorOption attribute for Oracle resource to 0.

Software limitations

This section describes the software limitations in this release.

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.

Limited global clustering capabilities in Veritas Cluster Server Management Console

The Veritas Cluster Server Management Console (VCS MC) management server currently includes limited capabilities for global clustering. Use the VCS graphical user interface (hagui) to take advantage of global clustering capabilities such as Fire Drill - Readiness and Remote Group Agent for remote groups.

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.

Performance recommendation for space-optimized volume snapshots

For minimal performance impact, Symantec recommends that the Space Optimized Snapshots (SOS) be created only of data volumes. A mirror breakoff snapshot should be created of the Oracle log volume. The log volumes are typically small in size and do not have significant space overhead.

Unsupported volume location scenarios

The ocrvol volume and votevol volume cannot exist in the same shared disk group as that of the Oracle datafiles. However, you can allow for this scenario when you manually configure Oracle service groups.

Oracle Disk Manager (ODM) limitation

Oracle Disk Manager (ODM) uses the Quick I/O driver for asynchronous I/O. Do not turn off the Quick I/O mount option, which is the default.

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.

Replication in a shared environment

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

Storage Checkpoint and Database FlashSnap limitation

The following are the limitations of Storage Checkpoint and Database FlashSnap:

- You cannot create a clone database using a mounted Storage Checkpoint.

- If you create an Oracle instance using the spfile option, you must run the `dbed_update` command before you can successfully perform any Storage Checkpoint or Database FlashSnap functions.
- Storage Checkpoints require file system layout version 6 or version 7. Use the vxupgrade(1M) command to check the current layout version and to change the layout version, if necessary. When upgrading a CFS file system, issue the command from the primary node. Note that after you upgrade a system to layout version 6 or version 7, the file system is no longer compatible with the older VxFS file systems.
- When cloning a database using Database FlashSnap, the Oracle database must have at least one mandatory archive destination. For more information about Oracle parameters for archiving redo logs, see your Oracle documentation.
- Only online snapshots are supported for an Oracle RAC database, when using the `dbed_vmsnap`, `dbed_vmclonedb`, and `dbed_vmchecksnap` commands.
- After running `dbed_vmsnap -o reverse_resync_commit`, your primary database is started using a pfile. If your original primary database used an spfile, you need to shut down the database and restart it using spfile. Then, run `dbed_update` to update the repository.
- The Storage Checkpoint and Database FlashSnap features of SF Oracle RAC do not support the graphical user interface of the Veritas Storage Foundation for Oracle product.
- The Database FlashSnap feature does not support RAID-5 volumes.
- SF Oracle RAC does not support the Veritas FlashSnap agent for Symmetrix (EMC TimeFinder) mapping functionality (package:`VRTSfas`).

Documentation errata

The following sections, if present, cover additions or corrections for Document version: 5.1SP1.0 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 56.

See "About Symantec Operations Readiness Tools" on page 10.

Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide

Table 1-9 Documentation errata

Topic	Correction
Section 8: Uninstallation of SF Oracle RAC **Chapter 31:** Preparing to uninstall SF Oracle RAC from a cluster **Topic:** Preparing to uninstall SF Oracle RAC from a cluster > Removing root disk encapsulation	The root disk encapsulation cannot be removed using the steps described in the topic. Do not perform the step. You may restore or create an LVM root disk image from the VxVM root disk on an other disk. Boot through the LVM root disk and remove VxVM.

Documentation

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

Documentation set

Table 1-10 lists the documentation for Veritas Storage Foundation for Oracle RAC.

Table 1-10 Veritas Storage Foundation for Oracle RAC documentation

Document title	File name
Veritas Storage Foundation for Oracle RAC Release Notes	sfrac_notes_51sp1_hpux.pdf
Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide	sfrac_install_51sp1_hpux.pdf
Veritas Storage Foundation for Oracle RAC Administrator's Guide	sfrac_admin_51sp1_hpux.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sf_adv_ora_51sp1_hpux.pdf

Table 1-11 lists the documents for Veritas Cluster Server.

Table 1-11 Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_51sp1_hpux.pdf
Veritas Cluster Server Release Notes	vcs_notes_51sp1_hpux.pdf
Veritas Cluster Server Administrator's Guide	vcs_admin_51sp1_hpux.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_51sp1_hpux.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev_51sp1.pdf
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	vcs_vvr_agent_51sp1_hpux.pdf
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_51sp1_hpux.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_51sp1_hpux.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_51sp1_hpux.pdf

Table 1-12 lists the documentation for Veritas Storage Foundation.

Table 1-12 Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_51sp1_hpux.pdf
Veritas Storage Foundation and High Availability Installation Guide	sf_install_51sp1_hpux.pdf
Veritas Storage Foundation Cluster File System Release Notes	sfcfs_notes_51sp1_hpux.pdf
Veritas Storage Foundation Cluster File System Administrator's Guide	sfcfs_admin_51sp1_hpux.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sf_adv_ora_51sp1_hpux.pdf
Veritas Storage Foundation Advanced Features Administrator's Guide	sf_adv_admin_51sp1_hpux.pdf

Table 1-13 lists the documentation for Veritas Volume Manager and Veritas File System.

Table 1-13 Veritas Volume Manager and Veritas File System documentation

Document title	File name
Veritas Volume Manager Administrator's Guide	vxvm_admin_51sp1_hpux.pdf
Veritas Volume Manager Troubleshooting Guide	vxvm_tshoot_51sp1_hpux.pdf
Veritas File System Administrator's Guide	vxfs_admin_51sp1_hpux.pdf
Veritas File System Programmer's Reference Guide	vxfs_ref_51sp1_hpux.pdf

Table 1-14 lists the documentation for Veritas Volume Replicator.

Table 1-14 Veritas Volume Replicator documentation

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

Table 1-15 lists the documentation for Symantec Product Authentication Service (AT).

Table 1-15 Symantec Product Authentication Service documentation

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

Manual pages

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

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

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

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

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

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

See the `man`(1) manual page.