Veritas Storage Foundation™ for Oracle® RAC Release Notes

Linux

6.0

May 2013

Symantec™

Veritas Storage Foundation™ for Oracle RAC Release Notes

Product version: 6.0

Document version: 6.0.5

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 website at the following URL:

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

About Symantec Connect

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

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

Storage Foundation for Oracle RAC Release Notes

This document includes the following topics:

- About this document
- Component product release notes
- About Veritas Storage Foundation for Oracle RAC
- About Symantec Operations Readiness Tools
- Important release information
- Changes introduced in 6.0
- No longer supported
- System requirements
- Fixed issues
- Known issues
- Software limitations
- Documentation

About this document

This document provides important information about Veritas Storage Foundation for Oracle RAC (SF Oracle RAC) version 6.0 for Linux. Review this entire document before you install or upgrade SF Oracle RAC.

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

This is Document version: 6.0.5 of the *Veritas Storage Foundation for Oracle RAC Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

https://sort.symantec.com/documents

Component product release notes

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

`/product_name/docs`

Symantec recommends copying the `docs` directory on the software media that contains the product guides to the `/opt/VRTS` 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 (6.0)*.
- Storage Foundation (SF)
  See *Veritas Storage Foundation Release Notes (6.0)*.
- Storage Foundation Cluster File System High Availability (6.0)
  See *Veritas Storage Foundation Cluster File System High Availability Release Notes (6.0)*.

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 Storage Foundation Cluster File System High Availability (SFCFS), and Veritas Storage Foundation,

which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

Benefits of SF Oracle RAC

SF Oracle RAC provides the following benefits:

- Support for file system-based management. SF Oracle RAC provides a generic clustered file system technology for storing and managing Oracle data files as well as other application data.
- Support for high-availability of cluster interconnects.
  The PrivNIC/MultiPrivNIC agents provide maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.
  See the following Technote regarding co-existence of PrivNIC/MultiPrivNIC agents with Oracle RAC 11.2.0.2:
  http://www.symantec.com/business/support/index?page=content&id=TECH145261
- Use of Cluster File System and Cluster Volume Manager for placement of Oracle Cluster Registry (OCR) and voting disks. These technologies provide robust shared block interfaces for placement of OCR and voting disks. In the absence of SF Oracle RAC, separate LUNs need to be configured for OCR and voting disks.
- Support for a standardized approach toward application and database management. Administrators can apply their expertise of Veritas technologies toward administering SF Oracle RAC.
- Increased availability and performance using Veritas Dynamic Multi-Pathing (DMP). DMP provides wide storage array support for protection from failures and performance bottlenecks in the Host Bus Adapters (HBA), Storage Area Network (SAN) switches, and storage arrays.
- Easy administration and monitoring of multiple SF Oracle RAC clusters using Veritas Operations Manager.
- VCS OEM plug-in provides a way to monitor SF Oracle RAC resources from the OEM console.
  For more information, see the *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases* guide.
- Improved file system access times using Oracle Disk Manager (ODM).
- Ability to configure Oracle Automatic Storage Management (ASM) disk groups over CVM volumes to take advantage of Veritas 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, Storage Checkpoints, and Database Storage Checkpoints.
  For more information, see the *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases* guide.
- Support for space optimization using periodic deduplication in a file system to eliminate duplicate data without any continuous cost.
  For more information, see the Veritas Storage Foundation Administrator's documentation.
- Ability to fail over applications with minimum downtime using Veritas Cluster Server (VCS) and Veritas Cluster File System (CFS).
- 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 application data, in addition to Oracle database files, across nodes.
- Support for policy-managed databases in Oracle RAC 11g Release 2.
- 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.
- 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 detach 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 Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps

you manage your datacenter more efficiently and get the most out of your Symantec products.

SORT can help you do the following:

Prepare for your next installation or upgrade	■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture. ■ Analyze systems to determine if they are ready to install or upgrade Symantec products. ■ Download the latest patches, documentation, and high availability agents from a central repository. ■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems.
Manage risks	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository. ■ Identify and mitigate system and environmental risks. ■ Display descriptions and solutions for hundreds of Symantec error codes.
Improve efficiency	■ Find and download patches based on product version and platform. ■ List installed Symantec products and license keys. ■ Tune and optimize your environment.

Note: Certain features of SORT are not available for all products. Access to SORT is available at no extra cost.

To access SORT, go to:

https://sort.symantec.com

Important release information

- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website: http://www.symantec.com/docs/TECH164885
- For the latest patches available for this release, go to: http://sort.symantec.com/

- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:
  http://www.symantec.com/docs/TECH170013
  Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

Changes introduced in 6.0

This section lists the changes in 6.0.

Changes in SF Oracle RAC

This section describes the changes to SF Oracle RAC in this release.

Support for policy-managed database environments

SF Oracle RAC now supports policy-managed database environments in Oracle RAC 11g Release 2. In policy-managed database environments, the administrator specifies the server pool on which the database instances run. Oracle Grid Infrastructure determines the server on which the database instances run. You can use the VCS agent for Oracle to start, stop, and monitor policy-managed databases.

Support for Web-based installation and configuration

This release extends the Web-based installer capabilities to support the following tasks:

- SF Oracle RAC installation and configuration checks
- I/O fencing configuration
- Rolling upgrade
- Preparation, installation, and post-installation tasks of Oracle RAC
- Adding nodes to a cluster

Web-based installer capabilities for the following tasks already exist and were introduced in version 5.1SP1:

- Installation and configuration of SF Oracle RAC
- Uninstallation of SF Oracle RAC

SF Oracle RAC installer enhancements

This release introduces the following enhancements in the SF Oracle RAC installer:

- Mirroring is now optional for the creation of OCR and voting disk storage.
- You can now choose to create either single or separate file systems for OCR and voting disk.
- The application resource for cssd can now be named as preferrable. The name cssd is no longer embedded into the monitor script.
- The `SF Oracle RAC Installation and Configuration Checks` option now includes a new check that verifies whether or not the user `nobody` exists on all nodes in the cluster.

CRSResource agent support for Oracle RAC 11g Release 2

The CRSResource agent is now supported for Oracle RAC 11g Release 2.

Creating a backup boot disk group when the boot disk is encapsulated and mirrored during upgrades

When you upgrade from a 5.1 Service Pack (SP) 1 or later release, the installer can split a mirrored boot disk group to create a backup disk group. You can use this backup in case of an upgrade failure.

Support for product installation using yum on Linux

You can now install any of the Veritas products with yum. Yum installation is supported for Red Hat Enterprise Linux 5 and 6.

See the *Installation Guide* for more information.

The installer can now detect duplicate VCS cluster IDs and can automatically generate cluster IDs

The installer can now detect duplicate VCS cluster IDs and prompt you to select an unused one. It can also generate an unused ID during installation.

The installer can check product versions and hotfixes

You can check the existing product versions using the installer command with the `-version` option before or after you install. After you have installed the current version of the product, you can use the `showversion` script in the /opt/VRTS/install directory to find version information.

You can discover the following information with these commands:

- The installed version of all released Strorage Foundation and High Availability Suite of products
- The missing required RPMs or patches as applicable for platform
- The available updates (including patches or hotfixes) from SORT for the installed products

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

Using the installer's postcheck option

You can use the installer's postcheck option to diagnose installation-related problems and to provide troubleshooting information.

Rolling upgrade improvements

The rolling upgrade procedure has been streamlined and simplified.

Allow Response files to change tuning parameters

You can set non-default product and system tunable parameters using a tunables template file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing during or after the installation procedure.

See the *Installation Guide* for more information.

Changes related to SFDB tools

The following sections describe the changes related to Storage Foundation for Databases (SFDB) tools in SF Oracle RAC 6.0.

Support for space-optimized snapshots for database cloning

You can use Storage Foundation for Databases (SFDB) tools to take space-optimized snapshots of your Oracle database and then create database clones by using those snapshots. SFDB tools use the underlying features of Storage Foundation for this operation.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Enhancements to Cached ODM Advisor (dbed_codm_adm)

You can use the Cached ODM Advisor command `dbed_codm_adm` to generate a variety of reports that help you determine which data files are suitable for enabling Cached ODM. The reports generated by Cached ODM Advisor are enhanced to use the historical data from Oracle Automatic Workload Repository (AWR).

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Support for space-optimized snapshots on DR site for database cloning

You can use Storage Foundation for Databases (SFDB) tools in a replicated environment to take space-optimized snapshots on a disaster recovery (DR) site. This functionality lets you create clones of your Oracle database on the DR site in a setup where the database on the primary site is being replicated on the DR site.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Single CLI for different point-in-time copy operations

You can use the new SFDB command `vxsfadm` to perform various point-in-time copy operations on your Oracle database. `vxsfadm` provides the following benefits:

- Uniform command line for multiple operations
- Use case based functionality
- Enhanced error handling

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Support for file-level snapshots for database cloning

You can use Storage Foundation for Databases (SFDB) tools to take file-level snapshots of your Oracle database and then create one or more clones based on those snapshots. SFDB tools use the underlying features of Storage Foundation for this operation.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

Enhanced authentication support

The authentication support for Storage Foundation for Databases (SFDB) tools is enhanced in this release. You can use the `sfae_auth_op` to set up and configure authentication for SFDB tools.

See *Veritas Storage Foundation: Storage And Availability Management for Oracle Databases*.

SmartTier integration with OEM

You can now view the following SmartTier related information in the Oracle Enterprise Manager (OEM) reports:

- Storage allocation and free space in each tier
- Space occupied by a data file in each tier
  This is useful when a part of a data file is moved from tier to tier when database objects such as table or index are moved.

Packaging updates

The following lists the package changes in this release.

- New `VRTSsfcpi60` RPM for product installer scripts
  The `VRTSsfcpi60` RPM is introduced in this release. The VRTSsfcpi60 RPM contains the installer scripts and libraries that the installer uses to install, configure and upgrade Veritas products.

For more information, see the *Installation Guide*.

Changes to SF Oracle RAC clusters in secure mode

In this release, the installation and configuration experience of secure cluster is considerably simplified. You can easily convert the cluster into secure cluster with this simplified secure cluster configuration model.

The new architecture is based on embedded VxAT, where the security components are installed as a part of the SF Oracle RAC package. The root broker is no longer a single-point-of-failure in the new architecture. There is no dependency on a separate VRTSat package. Non-root users who are already logged on SF Oracle RAC hosts are now not prompted for password. Additionally, a cluster-level user feature is introduced to simplify user administration in secure clusters.

See the *Installation Guide* and *Administrator's Guide* for more information.

Changes to LLT

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

- LLT now supports VLAN tagging (IEEE 802.1Q).
- The `lltconfig` command includes the following new options:
  - -N
    You can use this option to list all the used cluster IDs.
  - -M
    You can use this option to display the currently loaded LLT module version information.

  See the `lltconfig` manual page for more information.
  See the `llttab` manual page for more information.
- Link utilization statistics are enhanced that help in the root cause analysis of performance related issues.
- Periodic flushing of ARP cache is disabled.
- When MAC address of a NIC changes, LLT immediately relearns the new MAC address and also updates the peer nodes about the change.

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

Changes to GAB

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

Better GAB and I/O fencing integration to ensure application availability

In the event of a split-brain situation before VxFEN module implements the decision, sometimes GAB proceeds with attempting to resolve the join after the split-brain. GAB removes all but one joining subcluster. This behavior can cause the entire cluster to shut down. To avoid this scenario, GAB now gives priority to the fencing module.

With the GAB and I/O fencing integration in this release, if the I/O fencing module's decision is still pending before GAB initiates a join of the subcluster, GAB delays the iofence message. GAB wait depends on the value of the VxFEN tunable parameter *panic_timeout_offst* based on which VxFEN computes the delay value and passes to GAB.

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

GAB can now recognize clients with names in addition to ports

When kernel clients initialize GAB API, they can now define a client name string. GAB now adds a client name which enables GAB to track the client even before GAB port is registered. GAB also passes the client name information to LLT when registering the LLT port. The `lltstat -p` command also displays the GAB client names when providing the status details of the ports in use.

This feature is applicable only to GAB kernel clients, and not applicable for user-land GAB clients such as HAD.

The gabconfig command has new -C option

The -C option of the `gabconfig` command lists the names of the GAB clients that have registered with GAB. The `-C` option when used with `-a` option lists the client names along with the port membership details.

Changes to I/O fencing

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

Support for racer node re-election during I/O fencing race

At the time of a network partition, the VxFEN module elects the lowest node in each sub-cluster as the racer node to race for the coordination points on behalf of the sub-cluster. The other spectator nodes wait on the racer node to do the fencing.

In the previous releases, the I/O fencing race was entirely dependent on the single racer node as follows:

- If the racer node is not able to reach a majority of coordination points, then the VxFEN module on the racer node sends a LOST_RACE message and all nodes in the subcluster also panic when they receive the LOST_RACE message.
- If the racer node panics during the arbitration, then the spectator nodes in the sub-cluster assume that the racer node lost the race and the spectator nodes also panic.

With the new racer node re-election feature, the VxFEN module re-elects the node with the next lowest node id in the sub-cluster as the racer node. This feature optimizes the chances for the sub-cluster to continue with the race for coordination points.

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

Support for multiple virtual IP addresses in CP servers

You can now configure multiple network paths (virtual IP addresses) to access a CP server. CP server listens on multiple virtual IP addresses. If a network path fails, CP server does not require a restart and continues to listen on one of the other available virtual IP addresses.

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.

Support for Quorum agent in CP servers

With the support for multiple virtual IP addresses, you can now use the Quorum agent to configure CP server service group failover policies. You can specify the minimum number of IP resources that must be online for the Quorum resource to remain online.

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.

With fencing enabled, GAB can now automatically seed the cluster when some cluster nodes are unavailable

In the earlier releases, if some of the nodes are not up and running in a cluster, then GAB port does not come up to avoid any risks of preexisting split-brain. In such cases, you can manually seed GAB using the command `gabconfig -x` to bring the GAB port up. However, if you have enabled I/O fencing in the cluster, then I/O fencing can handle any preexisting split-brain in the cluster.

In this release, I/O fencing has extended this functionality to be able to automatically seed GAB as follows:

- If a number of nodes in a cluster are not up, GAB port (port a) still comes up in all the member-nodes in the cluster.
- If the coordination points do not have keys from any non-member nodes, I/O fencing (GAB port b) also comes up.

This new functionality is disabled by default. You must manually enable this automatic seeding feature of GAB in clusters where I/O fencing is configured in enabled mode.

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

You can still use the `gabconfig -x` command to manually seed the cluster.

Graceful shutdown of a node no longer triggers I/O fencing race condition on peer nodes

In the earlier releases, a gracefully leaving node clears its I/O fencing keys from coordination points. But the remaining sub-cluster races against the gracefully leaving node to remove its registrations from the data disks. During this operation, if the sub-cluster loses access to the coordination points, the entire cluster may panic if the racer loses the race for coordination points.

In this release, this behavior has changed. When a node leaves gracefully, the CVM or other clients on that node are stopped before the VxFEN module is unconfigured. Hence, data disks are already clear of its keys. The remaining sub-cluster tries to clear the gracefully leaving node's keys from the coordination points but does not panic if it is not able to clear the keys.

Installer support to migrate between fencing configurations in an online cluster

You can now use the installer to migrate between disk-based and server-based fencing configurations. You can also replace the coordination points for any I/O fencing configuration in an online cluster using the same installer option. The installer uses the `vxfenswap` script internally.

You can also use response files to perform these I/O fencing reconfiguration operations.

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

Availability of shared disk group configuration copies

If the Cluster Volume Manager (CVM) master node loses access to a configuration copy, CVM redirects the read or write requests over the network to another node that has connectivity to the configuration copy. This behavior ensures that the disk group stays available.

In previous releases, CVM handled disconnectivity according to the disk group failure policy (dgfail_policy). This behavior still applies if the disk group version is less than 170. The dgfail_policy is not applicable to disk groups with a version of 170 or later.

Enhancements to CVM detach policies

In this release, the following changes have been made to the detach policies:

- A failure is considered global only if it affects all nodes in the cluster. Otherwise, a failure is treated as a local failure. Previously, any failure that affected more than one node was considered to be global.
- When the global detach policy is set, local failure of all plexes at the same time does not trigger plex detach. In this case, the volume remains enabled and I/Os fail on the node.
- When a node loses local connectivity to a disk, the disk is put in the lfailed state.

Enhancements to master node selection for failover

If the Cluster Volume Manager (CVM) master node leaves the cluster, CVM fails over the master role to another node in the cluster. In this release, CVM selects the node for failover based on the node's connectivity to the disks in the disk group. This behavior is an enhancement over previous releases of CVM.

During regular operations, CVM dynamically assigns an offset preference value to each node. The preference assignment is automatic, and generally does not require any intervention from the administrator.

If you need greater control over the master selection, you can also set customized preference values.

When a master failover occurs, CVM uses the custom node preferences together with the offset preference values to select the new master node.

Node join with DGDISABLED disk groups

In this release, a node can join the cluster even if there is a shared disk group that is in the DGDISABLED state. In previous releases, the node join would fail.

Licensing changes in the SFHA Solutions 6.0 release

Storage Foundation and High Availability Solutions 6.0 introduces the following licensing changes:

- The Cluster File System license is deprecated. CFS customers are entitled to the Storage Foundation Cluster File System High Availability (SFCFS HA) functionality.
- The VVR Option is renamed as Veritas Replicator Option. This option includes VVR (volume-based replication) and the new file-based replication solution.

- The VVR Enterprise license is deprecated; you can use Storage Foundation Enterprise and add Veritas Replicator Option to get this functionality. VVR Enterprise customers are entitled to Storage Foundation Enterprise with Replicator Option.
- The VCS license enables full cluster functionality as well as the limited start/stop functionality.
- Storage Foundation Enterprise CFS for Oracle RAC (Linux/x64) customers are entitled to Storage Foundation Enterprise for Oracle RAC (Linux/x64.)

The following functionality is included in the Standard and Enterprise licenses:

- The Compression feature is available with the Standard license.
- The SmartTier feature is now available with the Standard license.
- The Deduplication feature is available with the Enterprise license.

The following products are included in this release:

- Dynamic Multi-Pathing
- VirtualStore
- Storage Foundation Basic
- Storage Foundation Standard
- Storage Foundation Enterprise
- Veritas Cluster Server
- Veritas Cluster Server HA/DR
- Storage Foundation Standard HA: Storage Foundation Standard plus Veritas Cluster Server
- Storage Foundation Enterprise HA: Storage Foundation Enterprise plus Veritas Cluster Server
- Storage Foundation Enterprise HA/DR
- Storage Foundation Enterprise Cluster File System HA
- Storage Foundation Enterprise Cluster File System HA/DR
- Storage Foundation Enterprise for Oracle RAC
- Storage Foundation Enterprise HA/DR for Oracle RAC
- Storage Foundation Enterprise for Sybase ASE CE
- Storage Foundation Enterprise HA/DR for Sybase CE

HA: High Availability

HA/DR: High Availability and Disaster Recovery

Veritas Replicator Option can be added to all Storage Foundation and High Availability products, except Dynamic Multi-Pathing and Veritas Cluster Server.

Note that products, features, and options may differ by operating system and platform. Please see the product documentation for information on supported platforms.

Enhancements to collecting a VxExplorer troubleshooting archive

The Symantec Operations Readiness Tools (SORT) data collector contains functionality to collect and submit a VxExplorer archive. You can send this archive to Symantec Technical Support for problem diagnosis and troubleshooting. VxExplorer does not collect customer data.

The legacy `VxExplorer` script now works differently. When you run the script, it launches the SORT data collector on the specified local host with the `-vxexplorer` option.

To learn more about using the data collector to collect a VxExplorer archive, see:

www.symantec.com/docs/HOWTO32575

Changes related to product documentation

The Storage Foundation and High Availability Solutions 6.0 release includes the following changes to the product documentation.

Table 1-1 lists the documents introduced in this release.

Table 1-1 New documents

New documents	Notes
Veritas Storage Foundation Installation Guide	Installation and upgrade information for Storage Veritas Foundation.
Veritas Storage Foundation Administrator's Guide	Administration information for Veritas Storage Foundation.
Veritas Storage Foundation and High Availability Release Notes	Release-specific information for Veritas Storage Foundation and High Availability users.
Veritas Storage Foundation and High Availability Solutions Solutions Guide	Solutions and use cases for Veritas Storage Foundation and High Availability Solutions.

Table 1-1 New documents *(continued)*

New documents	Notes
Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide	Troubleshooting information for Veritas Storage Foundation and High Availability Solutions.
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	Virtualization-related information for Veritas Storage Foundation and High Availability Solutions.
Symantec VirtualStore Release Notes	Release-specific information Symantec VirtualStore.
Veritas Storage Foundation for Sybase ASE CE Release Notes	Release-specific information for Veritas Storage Foundation for Sybase ASE CE.
Veritas Storage Foundation for Sybase ASE CE Installation Guide	Installation information for Veritas Storage Foundation for Sybase ASE CE.
Veritas Storage Foundation for Sybase ASE CE Administrator's Guide	Administration information for Veritas Storage Foundation for Sybase ASE CE.
Virtual Business Services–Availability User's Guide	Information about Virtual Business Services. This document is available online.

Table 1-2 lists the documents that are deprecated in this release.

Table 1-2 Deprecated documents

Deprecated documents	Notes
Veritas File System Administrator's Guide	Content now appears in the *Veritas Storage Foundation Administrator's Guide* and in the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.
Veritas Volume Manager Administrator's Guide	Content now appears in the *Veritas Storage Foundation Administrator's Guide* and in the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.
Veritas Storage Foundation Advanced Features Administrator's Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Solutions Guide*.

Table 1-2 Deprecated documents *(continued)*

Deprecated documents	Notes
Veritas Volume Manager Troubleshooting Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide.*
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	Content now appears in the *Veritas Cluster Server Bundled Agents Reference Guide.*
Veritas Volume Replicator Planning and Tuning Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide.*
Veritas Volume Replicator Advisor User's Guide	Content now appears in the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide.*

Table 1-3 lists documents that are no longer bundled with the binaries. These documents are now available online.

Table 1-3 Online documents

Document
Veritas Cluster Server Agent Developer's Guide
Veritas File System Programmer's Reference Guide

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:

- Several documents are deprecated in this release.
  See "Changes related to product documentation" on page 23.
- Oracle RAC 11g Release 1 Clusterware
- ASMInst agent
  The ASMInst agent is no longer supported in SF Oracle RAC environments. The ASM instances are managed by Oracle Clusterware.
- Use of crossover cables
  Oracle does not support the use of crossover cables for cluster interconnects due to the possibility of data corruption and other software limitations.

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

The use of crossover cables is discouraged for new deployments.

- Bunker replication is not supported in a Cluster Volume Manager (CVM) environment.

Veritas Storage Foundation for Databases (SFDB) tools features which are no longer supported

The following Storage Foundation for Databases (SFDB) tools features are not supported in this release:

- FlashSnap reverse resync
- Checkpoint policy and Checkpoint quotas
- Interactive modes in clone and rollback

System requirements

The following topics describe 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://www.symantec.com/docs/TECH170013
- Latest information on support for Oracle database versions: www.symantec.com/docs/DOC4039
- 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 1-4 on page 27.

For campus clusters See Table 1-5 on page 28.

Table 1-4 Hardware requirements for basic clusters

Item	Description
SF Oracle RAC systems	Two to sixteen systems with two or more CPUs. For details on the additional requirements for Oracle, see the Oracle documentation.
DVD drive	A DVD drive on one of the nodes in the cluster.
Disks	SF Oracle RAC requires that all shared 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. The minimum size required for a coordinator disk is 128 MB.
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` You can also use the Veritas Web-based installation program to determine the available disk space. For details on the additional space that is required for Oracle, see the Oracle documentation.
RAM	Each SF Oracle RAC system requires at least 2 GB. For Oracle RAC requirements, see the Oracle Metalink document: 169706.1
Swap space	See the Oracle Metalink document: 169706.1

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

Item	Description
Network	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. Oracle requires that all nodes use the IP addresses from the same subnet.
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 1-5 lists the hardware requirements for campus clusters in addition to the basic cluster requirements.

Table 1-5 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 a common cross-site physical infrastructure for storage and LLT private networks.
I/O fencing	I/O fencing requires placement of a third coordinator point at a third site. The DWDM can be extended to the third site or the iSCSI LUN at the third site can be used as the third coordination point. Alternatively Coordination Point Server can be deployed at the third remote site as an arbitration point.

Supported Linux operating systems

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

Table 1-6 shows the supported Linux operating systems for this release.

Note: SF Oracle RAC is not yet supported on Oracle Linux 6 and RHEL 6.

Refer to the following TechNote for the latest information on supported operating systems and Oracle database versions:

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

The TechNote will be updated after Oracle supports Oracle Linux 6 and RHEL 6 and Symantec approves support for SF Oracle RAC on these platforms.

Table 1-6 Supported Linux operating systems

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 6	Update 1, 2	2.6.32-131.0.15.el6 2.6.32-220.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only
Red Hat Enterprise Linux 5	Update 5, 6, 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 11	SP1	2.6.32.12-0.7	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 10	SP4	2.6.16.60-0.85.1	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Linux 6	**6.1	2.6.32-131.0.15.el6	64-bit x86, EMT*/Opteron
Oracle Linux 5	**Update 5, 6, 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron

* Extended Memory Technology

** RHEL-compatible mode only.

Note: Only 64-bit operating systems are supported.

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

For DMP, SF, SFHA, SFCFSHA, SFRAC, VCS, and VirtualStore, Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

On Linux, Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel Application Binary Interface (ABI) compatibility.

For Storage Foundation for Oracle RAC, all nodes in the cluster must have the same operating system version and update level.

Supported database software

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

The following database versions are supported:

- Oracle RAC 11g Release 2

  Note: If you are running SLES 10 SP4, install the Oracle patch `12311357`.

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

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

Support for minor database versions is also documented in the afore-mentioned Technical Support TechNote.

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

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 57.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 6.0

Table 1-7 Veritas Storage Foundation for Oracle RAC fixed issues

Incident number	Description
2386570	The CSSD agent embeds `cssd` as the resource name in the cssd monitor script. This causes the agent to log the following warning message in the engine log file if the resource name is not set to cssd: `VCS WARNING V-16-1-10260 Resource does not exist: cssd`
2336374	The PrivNIC/MultiPrivNIC agents fail to remove the original IP address after a failover or failback operation.
2436063	The CRSResource agent fails to come online due to incorrect paths in the CRSResource monitor script.
2212272	Incorrect ownership assigned to the parent directory of ORACLE_BASE causes Oracle Clusterware/Grid Infrastructure installations to fail. When you use the SF Oracle RAC script-based installation program to install Oracle Clusterware/Grid Infrastructure, the ownership of the parent directory of ORACLE_BASE/GRID_BASE that is created by the installation program is incorrectly set to root. This causes the Oracle Universal Installer to display errors when it creates the oraInventory directory.

Table 1-7 Veritas Storage Foundation for Oracle RAC fixed issues *(continued)*

Incident number	Description
2491788	Failure to update the `DBHome` attribute while configuring the VIP ResType attribute logs an error message in the engine log file.
2555319	Failure to set the MTU (Maximum Transmission Unit) size in LLT over UDP environments causes issues with the PrivNIC/MultiPrivNIC agents. If the MTU size field is not set explicitly when you configure the PrivNIC/MultiPrivNIC agents in an LLT over UDP environment, the agents may fail to plumb the private IP addresses during their operations or may configure incorrect MTU size on the LLT interfaces.
2565842	During configuration of MultiPrivNIC, the SF Oracle RAC installer plumbs the IP address on a virtual interface despite the availability of the base interface.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP2

Table 1-8 describes the incidents that are fixed in Veritas Storage Foundation for Oracle RAC in 5.1 SP1 RP2.

Table 1-8 Veritas Storage Foundation for Oracle RAC fixed issues

Fixed issues	Description
2429449	The cssd agent explicitly uses hard-coded string "cssd" as resource name.
2390892	Starting the VCSMM driver on two or more nodes in the cluster causes a memory leak in the vcsmm_set_cluster_proto function during memory allocation
2374987	Failed to remove original IP address by PrivNIC and MultiPrivNIC agents during failover/failback operation
2374970	CRSResource agent support for 11gR2

Veritas Storage Foundation for Databases (SFDB) tools: Issues fixed in 6.0

Table 1-9 SFDB tools fixed issues

Fixed issues	Description
1840672	In a multiple disk group environment, if the snapshot operation fails then `dbed_vmsnap` fails to reattach all the volumes.
1469310	If the database fails over during FlashSnap operations, various error mesages appear.

Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP1

Table 1-10 describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in this release.

Table 1-10 Storage Foundation for Databases fixed issues

Incident	Description
2203917	Process table has been changed to use per-hash-bucket locks, and the number of buckets has been increased from 32 to 256.
2237709	The `dbdst_preset_policy` command no longer aborts when you specify the volume class as MEDIUM.

Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP2

There are no SFDB fixed issues in 5.1 SP1 RP2.

LLT, GAB, and I/O fencing: Issues fixed in 6.0

Table 1-11 lists the fixed issues for LLT, GAB, and I/O fencing.

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

Incident	Description
2515932	[GAB] gabconfig ioctl behaviour changed to return EALREADY if GAB is already configured.

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

Incident	Description
2495020	[Fencing] vxfend does not terminate if you run the `vxfenswap` command to change the fencing mode from 'scsi3' to 'customized', and chooses to rollback when vxfenswap prompts for confirmation.
2442402	[LLT] Reduce lltd CPU consumption by reducing the wakeup calls.
2437022	[Fencing] Fails to run the `vxfenswap` command to the same diskgroup when the disk policy changed.
2426664	[Fencing] vxfend does not terminate when you run the `vxfenswap` command to migrate from the customized mode to the scsi3 mode.
2411652	[GAB] Add a check in GAB for MAX message size of 64KB before enqueuing the message.
2386325	[Fencing] Fencing configuration fails and vxfenadm prints same serial number for all LUNs which have more than 96 bytes of SCSI Inquiry data in page 0x83.
2369742	[Fencing] Once vxfenconfig -c with a particular mode (say customized) has returned EFAULT ("1036 Unable to configure..."), all subsequent runs of vxfenconfig -c with a different mode (say scsi3) fail with error EBADMSG ("1050 Mismatched modes...").
2351011	[Fencing] The vxfenswap utility fails to accurately check for the exit status of the vxfenconfig commands run on the other nodes in the background. This may lead to the vxfenswap utility appearing indefinitely hung if the vxfenconfig process does not succeed for any reason.
2337916	[Fencing] Fencing shutdown script does not retry stopping the fencing module if fencing fails to unconfigure because of clients being registered.
2311361	[Fencing] Fencing details are printed in the engine logs every five minutes if fencing is running and the CoordPoint resource is configured.
2253321	[Fencing] Fencing fails to start if any of the coordination points is unavailable at the startup time.
2252470	[Fencing] Provide options to force the fencing library to obtain serial numbers using standard inquiry or extended inquiry using a variety of ID types.
2218448	[VxCPS] The cpsadm command fails if LLT is not installed or configured on a single-node cluster which hosts the CP server.

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

Incident	Description
2209664	[VxCPS] Configuring fencing is successful with three disks even when single_cp=1 and the formatting of warning messages aer required in vxfend_A.log.
2209144	[VxCPS] There is syntax error while unconfiguring CP server using the configure_cps.pl script.
2203070	[Fencing] Failed to configure fencing on a 64-node cluster, fencing comes up only on first 33 nodes.
2178126	[GAB] GAB fails to start if it is unable to allocate memory in atomic manner in low memory situations, typically in under-provisioned virtual machine setups.
2161816	[Fencing] 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.

2139883	[GAB] On RHEL5 Update 5 and later, messages similar to the following are repeatedly seen on the console:
	`INFO: task gablogd:22812 blocked for more than 120 seconds.`
	`"echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.`

2112742	[VxCPS] Server-based I/O fencing fails to start after configuration on nodes with different locale settings.
2100896	[Fencing] There is failure message even the migration from server-based to disk-based using vxfenswap succeeded.
2085941	[VxCPS] 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.

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

Incident	Description
2076240	[VxCPS] 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.
1973713	[Fencing] The agent XML files are missing for CP server agent.

Known issues

This section covers the known issues in this release.

For Oracle RAC issues:

See "Oracle RAC issues" on page 36.

For SF Oracle RAC issues:

See "SF Oracle RAC issues" on page 37.

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

See "Documentation" on page 57.

Oracle RAC issues

This section lists the known issues in Oracle RAC.

Oracle Grid Infrastructure installation may fail with internal driver error

The Oracle Grid Infrastructure installation may fail with the following error:

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

Workaround:

Perform one of the following steps depending on the type of installer you use for the installation:

- Script-based installer

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

- Web-based installer
  When you run the Web-based installer, in the **Enter the arguments to be passed to the Oracle installer** text box, enter the value `-ignoreInternalDriverError`.
  For more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

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

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.

SF Oracle RAC installer does not support use of fully qualified domain names (2585899)

The SF Oracle RAC installer does not support the use of fully qualified domain names (FQDN). Specifying the fully qualified domain name of a system results in the following error:

```
The node galaxy doesn't seem to be part of the cluster,
or CVM is not running on the node galaxy.
```

Workaround: Use only the host name of the system when you specify the system name.

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups (2591399)

The service groups freeze due to upgrading using the product installer if you stopped the installer after the installer already stopped some of the processes and then resumed the upgrade.

Workaround: You must unfreeze the service groups manually after the upgrade completes.

To unfreeze the service groups manually

1 List all the frozen service groups

```
# hagrp -list Frozen=1
```

2 Unfreeze all the frozen service groups:

```
# haconf -makerw
# hagrp -unfreeze service_group -persistent
# haconf -dump -makero
```

Unable to stop some SF Oracle RAC processes (2329580)

If you install and start SF Oracle RAC, but later configure SF Oracle RAC using `installvcs`, some drivers may not stop successfully when the installer attempts to stop and restart the SF Oracle RAC drivers and processes. The reason the drivers do not stop is because some dependent SF Oracle RAC processes may be in the running state.

Workaround: To re-configure the product, use the corresponding `install`*`product`* command to re-configure the product. Otherwise some processes may fail to stop or start.

For example, use `installsfrac` to re-configure SF Oracle RAC rather than using `installvcs`.

Web installer does not ask for authentication after the first session if the browser is still open (2509330)

If you install or configure SF Oracle RAC and then close the Web installer, if you have other browser windows open, the Web installer does not ask for authentication in the subsequent sessions. Since there is no option to log out of the Web installer, the session remains open as long as the browser is open on the system.

Workaround: Make sure that all browser windows are closed to end the browser session and subsequently log in again.

After finishing a kernel upgrade on a master node the cvm group on a slave node does not come online (2439439)

After successfully finishing a kernel upgrade on one node, the cvm group does not come online on the second node.

Workaround: Check that your cluster is not in a jeopardy state before you perform a rolling upgrade.

After performing a rolling upgrade, make sure the CVM is online on all nodes without errors (2595441)

Make sure that the CVM is online on all nodes without errors after you perform the first phase of a rolling upgrade. The CVM protocol version will not upgrade successfully on the nodes where CVM is offline or has errors.

If the CVM protocol version does note upgrade successfully, upgrade the CVM protocol on the CVM master node.

To upgrade the CVM protocol on the CVM master node

1 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

2 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

sfmh discovery issue when you upgrade your Veritas product to 6.0 (2622987)

If a host is not reporting to any management server but sfmh discovery is running before you upgrade to 6.0, sfmh-discovery may fail to start after the upgrade.

Workaround:

If the host is not reporting to VOM, stop sfmh-discovery manually before upgrading to 6.0 by executing the following command on the managed host:

```
/opt/VRTSsfmh/adm/vxvmdiscovery-ctrl.sh stop
```

Incorrect server names sometimes display if there is a clock synchronization issue (2627076)

When you install a cluster with the Web-based installer, you choose to synchronize your systems with an NTP server due to a clock synchronization issue, you may see the NTP server name in messages instead of your server names.

Workaround:

Ignore the messages. The product is still installed on the correct servers.

Stopping the Web installer causes Device Busy error messages (2633924)

If you start the Web installer, and then perform an operation (such as prechecking, configuring, or uninstalling), you may get an error message saying the device is busy.

Workaround: Do one of the following:

- Kill the start.pl process.
- Start the webinstaller again. On the first Web page you see that the session is still active. Either take over this session and finish it or terminate it directly.

Issues related to LLT

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

LLT connections are not formed when a vlan is configured on a NIC (2484856)

LLT connections are not formed when a vlan is configured on a NIC that is already used to configure an LLT link.

Workaround: Do not specify the MAC address of a NIC in the `lltab` file while configuring LLT if you want to configure a vlan later. If you have already specified the MAC address of a NIC, then delete the MAC address from the `lltab` file, and update the file before you restart LLT.

LLT may fail to detect when bonded NICs come up (2604437)

When LLT is configured over a bonded NIC and that bonded NIC is DOWN with the `ifconfig` command, LLT marks the corresponding link down. When the bonded NIC is UP again using the `ifconfig` command, LLT fails to detect this change and marks the link up.

Workaround: Close all the ports and restart LLT, then open the ports again.

Issues related to GAB

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

Cluster panics during reconfiguration (2590413)

While a cluster is reconfiguring, GAB broadcast protocol encounters a race condition in the sequence request path. This condition occurs in an extremely narrow window which eventually causes the GAB master to panic.

Issues related to I/O fencing

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

Installer is unable to split a cluster that is registered with one or more CP servers (2110148)

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

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.

Fencing does not come up on one of the nodes after a reboot (2573599)

If VxFEN unconfiguration has not finished its processing in the kernel and in the meantime if you attempt to start VxFEN, you may see the following error in the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1007 Vxfen already configured
```

However, the output of the `gabconfig -a` command does not list port b. The `vxfenadm -d` command displays the following error:

```
VXFEN vxfenadm ERROR V-11-2-1115 Local node is not a member of cluster!
```

Workaround: Start VxFEN again after some time.

The vxfenswap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)

The `vxfenswap` utility runs the `vxfenconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxfenswap` command using RSH (with the `-n` option), then RSH does not detect

the failure of validation of coordination points on a node. From this point, `vxfenswap` proceeds as if the validation was successful on all the nodes. But, it fails at a later stage when it tries to commit the new coordination points to the VxFEN driver. After the failure, it rolls back the entire operation, and exits cleanly with a non-zero error code. If you run `vxfenswap` using SSH (without the `-n` option), then SSH detects the failure of validation of coordination of points correctly and rolls back the entire operation immediately.

Workaround: Use the `vxfenswap` utility with SSH (without the `-n` option).

In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)

When you start server-based I/O fencing, the node may not join the cluster and prints error messages in logs similar to the following:

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1043
Detected a preexisting split brain. Unable to join cluster.
```

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
operation failed.
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@galaxy,
domaintype vx; not allowing action
```

The vxfend daemon on the application cluster queries the coordination point server (CP server) to check if the cluster members as seen in the GAB membership are registered with the CP server. If the application cluster fails to contact the CP server due to some reason, then fencing cannot determine the registrations on the CP server and conservatively assumes a pre-existing split-brain.

Workaround: Before you attempt to start VxFEN on the application, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

The cpsadm command fails if LLT is not configured on the application cluster (2583685)

The `cpsadm` command fails to communicate with the coordination point server (CP server) if LLT is not configured on the application cluster node where you run the `cpsadm` command. You may see errors similar to the following:

```
# cpsadm -s 10.209.125.200 -a ping_cps
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using
environment variable
```

```
CPS_NODEID
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

However, if you run the `cpsadm` command on the CP server, this issue does not arise even if LLT is not configured on the node that hosts CP server. The `cpsadm` command on the CP server node always assumes the LLT node ID as 0 if LLT is not configured.

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily, or if the node is a single-node VCS configuration where LLT is not configured, then the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

Workaround: Set the value of the `CPS_NODEID` environment variable to 255. The `cpsadm` command reads the `CPS_NODEID` variable and proceeds if the command is unable to get LLT node ID from LLT.

Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)

Even if the cluster nodes have no registration on the CP server and if you provide coordination point server (CP server) information in the `vxfenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

Workaround: Manually add the cluster nodes' and users' information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster nodes' and users' information to the CP server during configuration.

CP server repetitively logs unavailable IP addresses (2530864)

If coordination point server (CP server) fails to listen on any of the IP addresses that are mentioned in the `vxcps.conf` file or that are dynamically added using the command line, then CP server logs an error at regular intervals to indicate the failure. The logging continues until the IP address is bound to successfully.

```
CPS ERROR V-97-51-103 Could not create socket for host
10.209.79.60 on port 14250
CPS ERROR V-97-1400-791 Coordination point server could not
open listening port = [10.209.79.60]:14250
Check if port is already in use.
```

Workaround: Remove the offending IP address from the listening IP addresses list using the `rm_port` action of the `cpsadm` command.

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

Server-based fencing comes up incorrectly if default port is not mentioned (2403453)

When you configure fencing in customized mode and do no provide default port, fencing comes up. However, the `vxfenconfig -l` command output does not list the port numbers.

Workaround: Retain the "port=<port_value>" setting in the `/etc/vxfenmode` file, when using customized fencing with atleast one CP server. The default port value is 14250.

Secure CP server does not connect from localhost using 127.0.0.1 as the IP address (2554981)

The `cpsadm` command does not connect to the secure CP server on the localhost using 127.0.0.1 as the IP address

Workaround: Connect the secure CP server using any of the virtual IPs that is configured with the CP server and is plumbed on the local node.

Unable to customize the 30-second duration (2551621)

When the vxcpserv process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: No workaround.

NIC resource gets created with incorrect name while configuring CPSSG with the configure_cps.pl script (2585229)

The name of the NIC resource created by the `configure_cps.pl` script does not come out correct when, for example, m^{th} VIP is mapped to n^{th} NIC and every m is not equal to n. In this case, although CPSSG continues to function without any problem, when you unconfigure CPSSG using `configure_cps.pl`, it fails.

Workaround: To unconfigure CPSSG, you must remove the CPSSG configuration from the VCS configuration.

The cpsadm command fails after upgrading CP server to 6.0 in secure mode (2478502)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSat RPM is not removed from the system, the `cpsadm` command loads the old security libraries present on the system.

As the installer runs the `cpsadm` command on the CP server to add or upgrade the SF Oracle RAC cluster (application cluster), the installer also fails.

Workaround : Perform the following steps on all the nodes of the CP server:

- Rename cpsadm to cpsadmbin.

  ```
  # mv /opt/VRTScps/bin/cpsadm /opt/VRTScps/bin/cpsadmbin
  ```

- Create a file `/opt/VRTScps/bin/cpsadm` with the following content:

  ```
  #!/bin/sh
  EAT_USE_LIBPATH="/opt/VRTScps/lib"
  export EAT_USE_LIBPATH
  /opt/VRTScps/bin/cpsadmbin "$@"
  ```

- Provide the following permissions to the new file:

  ```
  # chmod 755 /opt/VRTScps/bin/cpsadm
  ```

PrivNIC and MultiPrivNIC agents not supported with Oracle RAC 11.2.0.2

The PrivNIC and MultiPrivNIC agents are not supported with Oracle RAC 11.2.0.2.

For more information, see the following Technote:

http://www.symantec.com/business/support/index?page=content&id=TECH145261

Node fails to join the SF Oracle RAC cluster if the file system containing Oracle Clusterware is not mounted (2611055)

The sequence number of the startup script for Oracle High Availability Services daemon (ohasd) is lower than some of the SF Oracle RAC components such as VXFEN and VCS. During system startup, if the file system containing Oracle Clusterware does not get mounted before the ohasd startup script is executed, the script continuously waits for the file system to become available. As a result, the other scripts (including those of SF Oracle RAC components) are not executed and the node being started does not join the SF Oracle RAC cluster.

Workaround: If the rebooted node does not join the SF Oracle RAC cluster, the cluster can be started manually using the following command:

```
# installsfrac -start node1 node2
```

Issue with format of the last 8-bit number in private IP addresses (1164506)

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.

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.

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.

Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool (2392741)

If the cardinality of a policy-managed Oracle RAC database is set to a number lesser than the number of nodes in the server pool, and if the Oracle agent tries to bring the database online on all the nodes in the server pool, the operation fails on some of the nodes in the server pool. The resource on respective nodes move to the faulted state.

"Configuration must be ReadWrite : Use haconf -makerw" error message appears in VCS engine log when hastop -local is invoked (2609137)

A message similar to the following example appears in the `/var/VRTSvcs/log/engine_A.log` log file when you run the `hastop -local` command on any system in a SF Oracle RAC cluster that has `CFSMount`resources:

```
2011/11/15 19:09:57 VCS ERROR V-16-1-11335 Configuration must be
ReadWrite : Use haconf -makerw
```

The `hastop -local` command successfully runs and you can ignore the error message.

Workaround: There is no workaround for this issue.

Multiple system panics upon unmounting a CFS file system (2107152)

There is a system panic when you unmount a `mntlock`-protected VxFS file system, if that device is duplicate mounted on different directories.

Workaround: There is no workaround for this issue.

The vxassist maxsize option fails to report the maximum size of the volume that can be created with given constraints when the disk group has the siteconsistent flag set (2563195)

The `vxassist maxsize` option fails to report the maximum size of volume that can be created with given constraints when the disk group has the siteconsistent flag set. The following error is reported:

```
# vxassist -g dgname maxsize
VxVM vxassist ERROR V-5-1-752 No volume can be created within the given
constraints
```

Workaround:

Specify the size explicitly to the `vxassist make` command.

A controller can remain disabled due to udev device removal after loss of connectivity to some paths on RHEL6 and SLES11 (2697321)

The issue may occur with NetApp LUNs in ALUA mode. When a device fails with a `dev_loss_tmo` error, the operating system (OS) device files are removed by `udev`. After this removal, a controller will remain in the disabled state until a reboot is run on the host. To avoid this issue, use the following workaround.

Workaround

To create the new rules file

1 Create the file `/etc/udev/rules.d/40-rport.rules` with the following content line:

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

2 Reboot the system.

3 If new LUNs are dynamically assigned to the host, run the following command:

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

Application group attempts to come online on primary site before fire drill service group goes offline on the secondary site (2107386)

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

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

Oracle group fails to come online if Fire Drill group is online on secondary cluster [2556835]

If a parallel global service group faults on the local cluster and does not find a failover target in the local cluster, it tries to failover the service group to the remote cluster. However, if the firedrill for the service group is online on a remote cluster, offline local dependency is violated and the global service group is not able to failover to the remote cluster.

Workaround: Offline the Firedrill service group and online the service group on a remote cluster.

Removal of SAN cable from any node in a global cluster setup takes application service groups offline on all nodes (2580393)

In a replicated global cluster setup, the removal of SAN cable from any node in the cluster causes the CFS mount points to fault. As a result, dependent application groups are taken offline and replication to the secondary site is adversely affected.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

Database Storage Checkpoint unmount may fail with device busy (2591463)

In some cases, when a database that is cloned using a Database Storage Checkpoint is shut down, an error similar to the following may occur:

```
SFAE Error:0457: Failed to unmount device
/dev/vx/dsk/datadg/datavol:Ckpt_1317707593_rw_1317708154.
Reason: VxFS returned error : umount: /tmp/clonedb/data: device is busy
```

Workaround

As an Oracle user, force shut down the clone database if it is up and then retry the unmount operation.

Incorrect error message if wrong host name is provided (2585643)

If you provide an incorrect host name with the `-r` option of `vxsfadm`, the command fails with an error message similar to one of the following:

```
FSM Error: Can't use string ("") as a HASH ref while "strict refs"
in use at /opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 776.

SFDB vxsfadm ERROR V-81-0609 Repository location  is invalid.
```

The error messages are unclear.

Workaround

Provide the name of a host that has the repository database, with the `-r` option of `vxsfadm`.

FlashSnap validate reports snapshot unsplittable (2534422)

The FlashSnap validation operation fails with the following error if the mirrors for data volumes and archive log volumes share the same set of disks:

```
SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.
```

Workaround

Ensure that snapshot plexes for data volumes and snapshot plexes for archive log volumes reside on separate set of disks.

Attempt to use SmartTier commands fails (2332973)

The attempts to run SmartTier commands such as `dbdst_preset_policy` or`dbdst_file_move` fail with the following error:

```
fsppadm: ERROR: V-3-26551: VxFS failure on low level mechanism
with message - Device or resource busy
```

This error occurs if a sub-file SmartTier command such as `dbdst_obj_move` has been previously run on the file system.

There is no workaround for this issue. You cannot use file-based SmartTier and sub-file SmartTier simultaneously.

`dbed_vmclonedb` ignores new clone SID value after cloning once (2580318)

After you have done FlashSnap cloning using a snapplan, any further attempts to create a clone from the same snapplan using the `dbed_vmclonedb` continue to use the original clone SID, rather than the new SID specified using the *new_sid* parameter.

This issue is also observed when you resynchronize the snapplan, take a snapshot again without specifying the new clone SID, and then try to clone with the new SID.

Workaround

You can use one of the following workarounds:

- After the snapshot is resynchronized, delete the snapplan using the `dbed_vmchecksnap -o remove` command. You can then use a new clone SID by creating a new snapplan, which may have the same name, and using the snapplan for taking more snapshots.
- Use the `vxsfadm` command to take the snapshot again and specify the clone SID with the snapshot operation so that the clone operation can be done with the new clone SID.

Attempt to use certain names for tiers results in error (2581390)

If you attempt to use certain names for tiers, the following error message is displayed:

```
SFORA dbdst_classify ERROR V-81-6107 Invalid Classname BALANCE
```

This error occurs because the following names are reserved and are not permitted as tier names for SmartTier:

- BALANCE
- CHECKPOINT
- METADATA

Workaround

Use a name for SmartTier classes that is not a reserved name.

User authentication fails (2579929)

The `sfae_auth_op -o auth_user` command, used for authorizing users, fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username>
```

Reattempting the operation fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0372 AT broker failed to start:
```

The authentication setup might have been run with a strict umask value, which results in the required files and directories being inaccessible to the non-root users.

Workaround

If you have not done authentication setup, set umask to a less strict value before running the `sfae_auth_op -o setup` or `sfae_auth_op -o import_broker_config` commands.

To set umask to a less strict value

- Use the command:

```
# umask 022
```

If you have already done authentication setup, perform the following steps.

To resolve the problem if you have already done authentication setup

1 Shut down the authentication broker, if it is running.

```
# /opt/VRTSdbed/at-broker/bin/sfaeatd.sh stop
```

2 Change the permissions for files and directories that are required to be readable by non-root users.

```
# chmod o+r /etc/vx/vxdbed/admin.properties
# chmod o+rx /var/vx/vxdba/auth/users
# find /opt/VRTSdbed/at-broker -type d -exec chmod o+rx {} \;
```

Clone operation failure might leave clone database in unexpected state (2512664)

If the clone operation fails, it may leave the clone database in an unexpected state. Retrying the clone operation might not work.

Workaround

If retrying does not work, perform one the following actions depending on the point-in-time copy method you are using:

- For FlashSnap, resync the snapshot and try the clone operation again.
- For FileSnap and Database Storage Checkpoints, destroy the clone and create the clone again.
- For space-optimized snapshots, destroy the snapshot and create a new snapshot.

Contact Symantec support if retrying using the workaround does not succeed.

FlashSnap resync fails if there is an existing space-optimized snapshot (2479901)

If you try a FlashSnap resync operation when there is an existing space-optimized snapshot, the resync operation fails with the following error:

```
Error: VxVM vxdg ERROR V-5-1-4597 vxdg join FS_oradg oradg failed
datavol_snp : Record already exists in disk group
archvol_snp : Record already exists in disk group
```

Workaround

Destroy the space-optimized snapshot first and then perform the FlashSnap resync operation.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.0.

When upgrading from SF Oracle RAC version 5.0 to SF Oracle RAC 6.0 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.

Clone command fails if PFILE entries have their values spread across 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.

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 [1985055]

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

Workaround: Set MonitorOption attribute for Oracle resource to 0.

Software limitations

This section covers the software limitations of this release.

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

See “Documentation” on page 57.

Oracle Clusterware/Grid Infrastructure installation fails if the cluster name exceeds 14 characters

Setting the cluster name to a value that exceeds 14 characters during the installation of Oracle Clusterware/Grid Infrastructure causes unexpected cluster membership issues. As a result, the installation may fail.

Workaround: Restart the Oracle Clusterware/Grid Infrastructure installation and set the cluster name to a value of maximum 14 characters.

Parallel execution of `vxsfadm` is not supported (2515442)

Only one instance of the `vxsfadm` command can be run at a time. Running multiple instances of `vxsfadm` at a time is not supported.

Stale SCSI-3 PR keys remain on disk after stopping the cluster and deporting the disk group

When all nodes present in the SF Oracle RAC cluster are removed from the cluster, the SCSI-3 Persistent Reservation (PR) keys on the data disks may not get preempted. As a result, the keys may be seen on the disks after stopping the cluster or after the nodes have booted up. The residual keys do not impact data disk fencing as they will be reused or replaced when the nodes rejoin the cluster. Alternatively, the keys can be cleared manually by running the `vxfenclearpre` utility.

For more information on the `vxfenclearpre` utility, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

Creating point-in-time copies during database structural changes is not supported (2496178)

SFDB tools do not support creating point-in-time copies while structural changes to the database are in progress, such as adding or dropping tablespaces and adding or dropping data files.

However, once a point-in-time copy is taken, you can create a clone at any time, regardless of the status of the database.

SELinux supported in disabled and permissive modes only

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

Policy-managed databases not supported by CRSResource agent

The CRSResource agent supports only admin-managed database environments in this release. Policy-managed databases are not supported.

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.

DMP settings for NetApp storage attached environment

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

Table 1-12 describes the DMP tunable parameters and the new values.

Table 1-12 DMP settings for NetApp storage attached environment

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

The change is persistent across reboots.

To change the tunable parameters

1 Issue the following commands:

```
# vxdmpadm settune dmp_restore_internal=60

# vxdmpadm settune dmp_path_age=120
```

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

```
# vxdmpadm gettune dmp_restore_internal

# vxdmpadm gettune dmp_path_age
```

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

The preferred fencing feature gives preference to more weighted or larger subclusters by delaying the smaller subcluster. This smaller subcluster delay is effective only if the initial RACER node in the larger subcluster is able to complete the race. If due to some reason the initial RACER node is not able to complete the race and the VxFEN driver activates the racer re-election algorithm, then the smaller subcluster delay is offset by the time taken for the racer re-election and the less weighted or smaller subcluster could win the race. This limitation though not desirable can be tolerated.

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.

In a cluster using SCSI-3 based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on both the data disks and coordinator disks. In a cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. 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 coordination 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 coordination 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.

Uninstalling VRTSvxvm causes issues when VxFEN is configured in SCSI3 mode with dmp disk policy (2522069)

When VxFEN is configured in SCSI3 mode with dmp disk policy, the DMP nodes for the coordinator disks can be accessed during system shutdown or fencing arbitration. After uninstalling VRTSvxvm RPM, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm RPM is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

Veritas Storage Foundation for Databases (SFDB) tools software limitations

The following are the SFDB tools software limitations in this release.

Oracle Data Guard in an Oracle RAC environment

Database snapshots and Database Storage Checkpoints are not supported in a Data Guard with Oracle RAC environment.

Upgrading if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 6.0: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 6.0.

Documentation

Product guides are available in the PDF format on the software media in the `/product_name/docs` directory. Additional documentation is available online.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The latest product documentation is available on the Symantec website.

http://sort.symantec.com/documents

Documentation set

Table 1-13 lists the documentation for Veritas Storage Foundation for Oracle RAC.

Table 1-13 Veritas Storage Foundation for Oracle RAC documentation

Document title	File name
Veritas Storage Foundation for Oracle RAC Release Notes	sfrac_notes_60_lin.pdf
Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide	sfrac_install_60_lin.pdf
Veritas Storage Foundation for Oracle RAC Administrator's Guide	sfrac_admin_60_lin.pdf

Table 1-14 lists the documentation for Veritas Storage Foundation Cluster File System High Availability.

Table 1-14 Veritas Storage Foundation Cluster File System High Availability documentation

Document title	File name
Veritas Storage Foundation Cluster File System High Availability Release Notes	sfcfs_notes_60_lin.pdf
Veritas Storage Foundation Cluster File System High Availability Installation Guide	sfcfs_install_60_lin.pdf
Veritas Storage Foundation Cluster File System High Availability Administrator's Guide	sfcfs_admin_60_lin.pdf

Table 1-15 lists the documents for Veritas Cluster Server.

Table 1-15 Veritas Cluster Server documentation

Title	File name
Veritas Cluster Server Installation Guide	vcs_install_60_lin.pdf
Veritas Cluster Server Release Notes	vcs_notes_60_lin.pdf
Veritas Cluster Server Administrator's Guide	vcs_admin_60_lin.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_60_lin.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev_60_unix.pdf

Table 1-15 Veritas Cluster Server documentation *(continued)*

Title	File name
Veritas Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_60_lin.pdf
Veritas Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_60_lin.pdf
Veritas Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_60_lin.pdf

Table 1-16 lists the documentation for Veritas Storage Foundation.

Table 1-16 Veritas Storage Foundation documentation

Document title	File name
Veritas Storage Foundation Release Notes	sf_notes_60_lin.pdf
Veritas Storage Foundation Installation Guide	sf_install_60_lin.pdf
Veritas Storage Foundation Administrator's Guide	sf_admin_60_lin.pdf
Veritas Storage Foundation: Storage and Availability Management for Oracle Databases	sf_adv_ora_60_lin.pdf
Veritas File System Programmer's Reference Guide	vxfs_ref_60_lin.pdf

Table 1-17 lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-17 Veritas Storage Foundation and High Availability Solutions products documentation

Document title	File name
Veritas Storage Foundation and High Availability Solutions Solutions Guide	sfha_solutions_60_lin.pdf
Veritas Storage Foundation and High Availability Solutions Virtualization Guide	sfha_virtualization_60_lin.pdf
Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide	sf_replication_admin_60_lin.pdf

If you use Veritas Operations Manager (VOM) to manage Veritas Storage Foundation and High Availability products, refer to the VOM product documentation at:

http://sort.symantec.com/documents

Manual pages

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

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

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

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

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

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

See the `man`(1) manual page.

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

To edit the man(1) configuration file

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

   ```
   export LC_ALL=C
   ```

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

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

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

   where other man paths are specified in the configuration file.

3. Add new section numbers. Change the line:

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

   to

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