Symantec™ Storage Foundation for Oracle® RAC 6.1 Release Notes - Solaris

February 2015

Symantec™ Storage Foundation for Oracle RAC Release Notes

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.1

Document version: 6.1 Rev 9

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 Symantec Storage Foundation for Oracle RAC

- About Symantec Operations Readiness Tools

- Important release information

- Changes introduced in SF Oracle RAC 6.1

- No longer supported

- System requirements

- Fixed issues

- Known issues

- Software limitations

- Documentation

About this document

This document provides important information about Symantec Storage Foundation for Oracle RAC (SF Oracle RAC) version 6.1 for Solaris. 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: 6.1 Rev 9" of the *Symantec 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:

`/docs/product_name`

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

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

- Symantec Cluster Server (VCS)
  See *Symantec Cluster Server Release Notes (6.1)*.

- Storage Foundation (SF)
  See *Symantec Storage Foundation Release Notes (6.1)*.

- Storage Foundation Cluster File System High Availability (6.1)
  See *Symantec Storage Foundation Cluster File System High Availability Release Notes (6.1)*.

About Symantec Storage Foundation for Oracle RAC

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

Benefits of SF Oracle RAC

SF Oracle RAC provides the following benefits:

- Support for file system-based management. SF Oracle RAC provides a generic clustered file system technology for storing and managing Oracle data files as well as other application data.

- Support for high-availability of cluster interconnects.
  For Oracle RAC 10g Release 2:
  The combination of LMX/LLT protocols and the PrivNIC/MultiPrivNIC agents provides maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.
  For Oracle RAC 11.2.0.1:
  The PrivNIC/MultiPrivNIC agents provide maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.

- 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 all supported Oracle RAC versions) and raw interfaces (for Oracle RAC 10g Release 2) for placement of OCR and voting disks.

- Support for a standardized approach toward application and database management. Administrators can apply their expertise of Symantec technologies toward administering SF Oracle RAC.

- Increased availability and performance using Symantec 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.

- 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 Symantec 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.

- Support for space optimization using periodic deduplication in a file system to eliminate duplicate data without any continuous cost.
  For more information, see the Symantec Storage Foundation Administrator's documentation.

- Ability to fail over applications with minimum downtime using Symantec 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 and later versions.

- 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 site fails, clients that are attached to the failed site can reconnect to a surviving site and resume access to the shared database.

- Support for configuring and administering multiple non-global zones with high availability to Oracle RAC instances inside the non-global zone.

- 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 website 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 and generate an Installation and Upgrade custom report. ■ List patches by product or platform, and in the order they need to be installed. Display and download the most recent patches or historical patches. ■ Display Array Support Library (ASL) details by vendor, platform, or Storage Foundation and High Availability (SFHA) version. ASLs make it easier to manage arrays that are connected to SFHA-based servers. ■ List VCS and ApplicationHA agents, documentation, and downloads based on the agent type, application, and platform.
Identify risks and get server-specific recommendations	■ Analyze your servers for potential environmental risks. Generate a Risk Assessment custom report with specific recommendations about system availability, storage use, performance, and best practices. ■ Display descriptions and solutions for thousands of Symantec error codes.

Improve efficiency

- Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), documentation, product releases, Hardware Compatibility Lists (HCLs), and VCS/ApplicationHA agents.
- Quickly gather installed Symantec product and license key information from across your production environment. Generate a License/Deployment custom report that includes product names, versions, and platforms, server tiers, Symantec Performance Value Units (SPVUs), and End of Service Life dates.
- List and download Symantec product documentation including product guides, manual pages, compatibility lists, and support articles.
- Access links to important resources on a single page, including Symantec product support, SymConnect forums, customer care, Symantec training and education, Symantec FileConnect, the licensing portal, and my.symantec.com. The page also includes links to key vendor support sites.
- Use a subset of SORT features from your iOS device. Download the application at:
https://sort.symantec.com/mobile

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/TECH211540

- For the latest patches available for this release, go to:
https://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/TECH211575

- The software compatibility list summarizes each Storage Foundation and High Availability (SFHA) Solutions product stack and the product features, operating system versions, and third-party products it supports. For the latest information on supported software, visit the following URL:
  http://www.symantec.com/docs/TECH213121

Note: Before you install or upgrade SFHA Solutions products, review the current compatibility lists to confirm the compatibility of your hardware and software.

Changes introduced in SF Oracle RAC 6.1

This section lists the changes in SF Oracle RAC 6.1.

Support for Oracle RAC 12c

SF Oracle RAC now supports installation of and upgrade to Oracle RAC 12c. The release also includes support for existing features of Oracle RAC in 12c. Support for new features of Oracle RAC 12c will be announced later.

CSSD agent enhancements

The CSSD agent is no longer a generic Application agent. It now has its own CSSD type definition that allows simpler configuration and flexible resource-handling.

The remaining changes are as follows:

- New attribute `RestartDaemons` introduced for Oracle RAC 11g Release 2 and later versions.
  The default value is set to 1 and indicates whether or not the Oracle Grid Infrastructure processes `ohasd`, `cssd`,`crsd`, and `evmd` are restarted if the status of these processes is unhealthy.

- Intelligent Monitoring Framework (IMF) is now supported for the `ohasd`, `cssd`,`crsd`, and `evmd` daemons.
  By default, IMF monitoring is enabled with a monitoring value of 3.

- The `Clean` function now uses the `force` option to forcibly stop Oracle Grid Infrastructure on nodes running Oracle RAC 11g Release 2.

As a result of these improvements, you will see the following changes during an upgrade:

- The agent type is set to `CSSD`.

■ The installer prompts for the Oracle Clusterware home directory. This is optional. The agent uses this information to locate Oracle Clusterware process binaries. If the value is not provided, the agent reads the information from the Oracle configuration file.

Changes related to installation and upgrades

The product installer includes the following changes in SF Oracle RAC 6.1.

Support for centralized installations using the Deployment Server

The Deployment Server is a script that makes it easier to install or upgrade SFHA releases. The Deployment Server lets you store multiple release images in one central location and deploy them to systems of any supported UNIX or Linux operating system (6.1 or later). Prior to 6.1, releases still require the same platform, architecture, distribution, and version of the operating system. You can use the Deployment Server if you want to install or upgrade multiple releases and or multiple platforms.

The Deployment Server lets you do the following as described in Table 1-1.

Table 1-1 Deployment Server functionality

Feature	Description
Manage release images	■ View available Storage Foundation releases. ■ Download maintenance and hot fix release images from the Symantec Operations Readiness Tools (SORT) website into a repository. ■ Load the downloaded release image files from FileConnect and SORT into the repository. ■ View and remove release image files stored in the repository.
Check versions	■ Discovers packages and patches installed on designated systems and informs you of the product and version installed, including installed hot fixes. ■ Identify base, maintenance, and hot fix level upgrades to your system and download maintenance and hot fix releases. ■ Query SORT for the most recent updates.

Table 1-1   Deployment Server functionality *(continued)*

Feature	Description
Install or upgrade systems	- Install or upgrade a release stored in the repository on selected systems. - In release 6.1 and later: - Install hot fix level releases. - Install SFHA from any supported UNIX or Linux operating system to any other supported UNIX or Linux operating system. - Automatically load the script-based installer hot fixes that apply to that release.

Note: The Deployment Server is available only for the script-based installer, not the web-based installer.

See the *Installation Guide* for more information.

Improved patching and updating process

You can now download product maintenance releases and public hot fix releases directly from the Symantec Operations Readiness Tools (SORT) website using the installer. When you use the `installer` command with the `-version` option, the installer now lists the available GA releases, maintenance releases, and hot fix releases. If you have Internet access, you can follow the installer prompts to download available patches and hot fixes to your local system.

Downloading patches and hot fixes requires the installer to make outbound networking calls. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network attempts by running the installer using the no Internet patch center (`-noipc`) option. When using the `-noipc` option, the installer does not try to connect to SORT website. For example:

```
# ./installer -version -noipc system1 system2
```

See the *Installation Guide* for more information.

Automatic download of installer hot fixes

If you are running the 6.1 product installer, and your system has Internet access, the installer automatically imports any needed installer hot fix, and begins using it.

If your system does not have Internet access, you can still download installer hot fixes manually using the Symantec Operations Readiness Tools patch finder tool.

Automatic downloading of installer hot fixes requires the installer to make outbound networking calls. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network attempts by running the installer using the no Internet patch center (-noipc) option.

See the *Installation Guide* for more information.

Support for simultaneously installing or upgrading base releases, maintenance patches, and hot fixes

Beginning with version 6.1, Symantec offers you a method to easily install or upgrade your systems directly to a base, maintenance, or hot fix level in one step using Install Bundles. Install Bundles is the ability for installers to merge so customers can install or upgrade directly to maintenance or hot fix levels in one execution. Install Bundles consists of executing the installer from a GA release with a pointer to a higher maintenance or hot fix release. The installer installs them both as if they were combined in the same release image. The various scripts, packages, and patch components are merged and multiple releases are installed together as if they are one install entity.

Note: This feature is not supported by the Deployment Server.

There are five possible methods of integration. All upgrades must be executed from the highest level script.

- Base + maintenance
- Base + hot fix
- Maintenance + hot fix
- Base + maintenance + hot fix
- Base or maintenance + multiple hot fixes

See the *Installation Guide* for more information.

Web installation program supports phased upgrade

You can now perform a phased upgrade of your product with the web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems.

See the *Installation Guide* for more information.

Support for Solaris 11 Live Upgrade

You can use Live Upgrade on Solaris 11 systems to perform an upgrade of the product and the Solaris operating system. For the Live Upgrade process, an alternate boot environment is created on the primary boot disk by ZFS storage application. All boot environments are saved in the current disk. Thus, an alternate boot disk is not needed anymore.

See the *Installation Guide* for more information.

Changes related to Veritas File System

Veritas File System (VxFS) includes the following changes in 6.1:

Support for 64-bit quotas

Starting in release 6.1, 64-bit quotas are supported on disk layout Version 10. Users were earlier limited to set quota usage limits only up to 1 terabyte, restricting functionality in high data usage environments. With the support for 64-bit quotas, the quota usage limit can be set up to 4 exabytes.

As for 32-bit quotas, this continues to be supported on disk layout Version 9 or earlier. The same quota commands can be used for both 32-bit and 64-bit quotas.

As for 64-bit quotas, there are two new quotas files. For group quotas the file name is `quotas.grp.64` and for user quotas the file name is `quotas.64`. These files will be created on each file system after the disk layout version upgrade is completed.

See the *Administrator's Guide* for more information on quota files on Veritas File System.

See the *Installation Guide* for more information on upgrading disk layout versions.

maxlink support

Added support for more than 64K sub-directories. If maxlink is disabled on a file system, the sub-directory limit will be 32K by default. If maxlink is enabled on a file system, this allows you to create up to 4294967295(2^32 – 1) sub-directories.

By default maxlink is enabled.

See the *Administrator's Guide*.

Disk layout Version 10

In this release, disk layout Version 10 is now the default version.

Version 10 disk layout enables support for maxlink.

See the *Administrator's Guide*.

Support for direct mount inside non-global zones using Mount agent

You can mount VxFS directly inside a non-global zone. To mount VxFS inside a non-global zone, override the ContainerOpts attribute at the resource level and set the value of the RunInContainer attribute to 1.

Changes related to SFDB tools

The following sections describe the changes related to Storage Foundation for Databases (SFDB) tools in 6.1.

Reverse Resync for Oracle database recovery

In this release, the SFDB tools reintroduce the Reverse Resync feature for Oracle database recovery.

Reverse Resynchronization or Reverse Resync process helps in recovering a database from its volume snapshots using FlashSnap service.

Storage Foundation Database FlashSnap service is used to reverse resynchronize an online point-in-time copies image of a database in an Oracle environment.

Reverse Resync feature was supported in 5.X release. This feature was discontinued for 6.0 and 6.0.1 releases. In the current release, Reverse Resync feature is reintroduced with the following changes:

- You can perform ReverseResyncBegin operation after ReverseResyncAbort operation

- You can control the database recovery in ReverseResyncBegin operation using the new (optional) parameters:

  ```
  Reverse_Resync_Recovery
  ```

  ```
  Reverse_Resync_Archive_Log
  ```

Use the following commands for reverse resynchronization of the snapshot volume:

- `vxsfadm -o rrbegin` to start the Reverse Resync operation

- `vxsfadm -o rrcommit` to commit the Reverse Resync changes

- `vxsfadm -o rrabort` to abort or cancel the Reverse Resync operation and to go back to the original data volumes

---

Note: Reverse resync is not supported for RAC databases.

Supported Oracle configurations

In 6.1 release, SFDB tools support Oracle 12c release for Oracle databases.

Note: For Oracle 12c, the SFDB tools do not support the Multitenant database features, including the CDB and PDB databases.

Support for instant mode snapshots for Oracle RAC databases

In 6.1, the SFDB tools support instant mode snapshots for Oracle RAC databases.

Changes to LLT

Symantec Cluster Server includes the following changes to LLT in 6.1:

New SMF services avoid race conditions when you add or remove LLT driver on Solaris 11

On Solaris 11, Symantec has added two new SMF services, 'llt-postinstall', and 'llt-preremove' to manage addition and removal of LLT driver. With the addition of these new SMF services, the LLT driver is added only during package installation and removed on package removal. The new SMF services avoid failure to install the LLT driver during system restart.

LLT command changes

The following command changes are introduced in this release.

Updates in `lltconfig`:

- A new option `lltconfig -l` is introduced. When you add a new link, you can use the `-l` option to specify that the link is a low priority link.

Changes to GAB

Symantec Cluster Server (VCS) includes the following changes to GAB in 6.1:

Adaptive GAB tunables to prevent false failover

You can configure the VCS environment variables, `VCS_GAB_TIMEOUT_SECS` and `VCS_GAB_PEAKLOAD_TIMEOUT_SECS`, to make GAB adaptive to different load conditions on a node (per CPU load). GAB calculates the timeout range for the load

period based on the load average number provided by the operating system and the variable values that are set for HAD. GAB kills HAD after the timeout period.

For more information, see the *Symantec Cluster Server Administrator's Guide*.

New SMF services avoid race conditions when you add or remove GAB driver on Solaris 11

On Solaris 11, Symantec has added two new SMF services, 'gab-postinstall', and 'gab-preremove' to manage addition and removal of GAB driver. With the addition of these new SMF services, the GAB driver is added only during package installation and removed on package removal. The new SMF services avoid failure to install the GAB driver during system restart.

Changes to I/O fencing

Symantec Cluster Server (VCS) includes the following changes to I/O fencing in 6.1:

Set the order of coordination points while configuring I/O fencing

You can use the `-fencing` option in the installer to set the order of coordination points.

Decide the order of coordination points (coordination disks or coordination point servers) in which they participate in a race during a network partition. The order of coordination points you set in the installer is updated to the /etc/vxfenmode file. I/O fencing approaches the coordination points based on the order listed in the vxfenmode file.

So, the order must be based on the possibility of I/O Fencing reaching a coordination point for membership arbitration.

For more information, refer to the *Symantec Cluster Server Installation Guide*.

Refresh keys or registrations on the existing coordination points using the install program

You can use the `-fencing` option with the installer to refresh registrations on the existing coordination points.

Registration loss on the existing coordination points may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs. You must refresh registrations on coordination points when

the CoordPoint agent notifies VCS about the loss of registrations on any of the existing coordination points.

You can also perform a planned refresh of registrations on coordination points when the cluster is online without application downtime on the cluster.

For more information, refer to the *Symantec Cluster Server Installation Guide*.

New SMF services avoid race conditions when you add or remove I/O fencing driver on Solaris 11

On Solaris 11, Symantec has added two new SMF services, 'vxfen-postinstall', and 'vxfen-preremove' to manage addition and removal of I/O fencing driver. With the addition of these new SMF services, the I/O fencing driver is added only during package installation and removed on package removal. The new SMF services avoid failure to install the I/O fencing driver during system restart.

Support for HTTPS communication between CP server and application client cluster nodes

CP server and its application client cluster nodes can communicate securely over HTTPS, an industry standard protocol. Prior to release 6.1, communication between the CP server and its clients happened over the Inter Process Messaging (IPM) protocol, which is a Symantec proprietary protocol. Secure communication over IPM-based communication uses Symantec Product Authentication Services (AT) to establish secure communication between CP server and client nodes. With secure communication using HTTPS, CP server functionality is backward-compatible with previous releases. To support client nodes on releases before 6.1, CP server supports IPM-based communication in addition to HTTP-based communication. However, client nodes from 6.1 onwards only support HTTPS-based communication.

For more information, refer to the Symantec Cluster Server Installation Guide and Symantec Cluster Server Administrator's Guide.

The security attribute in `/etc/vxfenmode` file is obsolete

From VCS 6.1, the Coordination Point (CP) client will communicate with CP server using HTTPS protocol. The 'security' parameter in `/etc/vxfenmode` is therefore deprecated and setting it to 1 or 0 has no effect whatsoever.

Rolling upgrade of an application cluster to release version 6.1 requires CP server running release version 6.1

The application clusters and CP servers running on release version 6.1 communicate over the HTTPS protocol. Hence, an application cluster which is using CP server

as a fencing coordination point can no longer access the pre-6.1 CP server after the cluster is upgraded to 6.1. To ensure a smooth upgrade, either application cluster must use CP servers running release version 6.1 or the CP servers running an earlier release version must be upgraded to 6.1. Note that CP server running release version 6.1 can still work with pre-6.1 application clusters.

Checks introduced in `vxfentsthdw` utility for disk size and option to override errors

The `vxfentsthdw` utility is enhanced to check the disks for size compatibility and new error messages are introduced for better error evaluation. The utility also provides the override option (`-o`) to override size-related errors and continue testing.

New command for `hacli` in `vxfenswap` utility

A new option `-p` is introduced to specify a protocol value that `vxfenswap` utility can use to communicate with other nodes in the cluster. The supported values for the protocol can be `ssh`, `rsh`, or `hacli`.

DMP support for thin reclamation commands

In this release, Dynamic Multi-Pathing (DMP) adds support for the `UNMAP` command for thin reclamation. The Array Support Library (ASL) for each array uses the most suitable reclamation method supported for the array. In previous releases, DMP performed reclamation with the WRITE_SAME method for SCSI and the TRIM method for SSD devices. You can use the `vxdisk -p list` command to show the reclaim interface that is supported for a particular device.

For more information, see the *Administrator's Guide*.

DMP support for ZFS root on Solaris

Starting with this release, Dynamic Multi-Pathing (DMP) supports the ZFS root file system. When you install DMP with native support enabled, or enable native support with the tunable `dmp_native_support`, DMP also migrates the ZFS root pool to DMP. Reboot the system for the change to take effect.

DMP support for ZFS root requires Solaris 11 update 1 or later.

See "Supported Solaris operating systems " on page 28.

For more information about configuring ZFS root, see the *Symantec Dynamic Multi-Pathing Administrator's Guide - Solaris*.

Changes related to product name branding

Beginning with the 6.1 release, Storage Foundation and High Availability Solutions product names are rebranded.

Table 1-2 lists the rebranded Storage Foundation and High Availability Solutions products.

Table 1-2　　　　Rebranded Storage Foundation and High Availability Solutions products

Old product name	New product names with Symantec branding
Veritas Storage Foundation	Symantec Storage Foundation (SF)
Veritas Dynamic Multi-Pathing	Symantec Dynamic Multi-Pathing (DMP)
Veritas Replicator Option	Symantec Replicator Option
Veritas Volume Replicator	Symantec Volume Replicator (VVR)
Veritas Storage Foundation Cluster File System HA	Symantec Storage Foundation Cluster File System HA (SFCFSHA)
Veritas Storage Foundation for Oracle RAC	Symantec Storage Foundation for Oracle RAC (SFRAC)
Veritas Storage Foundation HA	Symantec Storage Foundation HA (SFHA)
Veritas Cluster Server	Symantec Cluster Server (VCS)
Veritas Disaster Recovery Advisor	Symantec Disaster Recovery Advisor (DRA)
Veritas Storage Foundation and High Availability Solutions	Symantec Storage Foundation and High Availability Solutions (SFHAS)
Veritas High Availability Agent Pack	Symantec High Availability Agent Pack
Veritas File System Software Development Kit	Symantec File System Software Development Kit

Symantec rebranding does not apply to the following:

- Product acronyms
- Command names
- Error messages
- Alert messages

■ Modules and components

■ Feature names

■ License key description

■ Veritas Operations Manager product branding

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:

■ Oracle RAC 11g Release 1 Clusterware

■ PrivNIC and MultiPrivNIC agents for Oracle RAC 11.2.0.2 and later versions

■ Dissimilar version upgrades of SF Oracle RAC components
For example, if you have SF Oracle RAC 6.0 installed, you can not upgrade only VCS to version 6.1.

■ 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.

Symantec 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:

■ Storage Checkpoint policy and Storage Checkpoint quotas

■ Interactive modes in clone and rollback

System requirements

This section describes the system requirements for this release.

Supported Oracle VM Server for SPARC

Supported Oracle VM Server for SPARC (OVM) versions are OVM 2.0, OVM 2.1, OVM 2.2, OVM 3.0 and OVM 3.1.

For supported OS version for Oracle VM Server for SPARC, refer to *Oracle VM server for SPARC Release Notes*.

The version of the Oracle Solaris operating system (OS) that runs on a guest domain is independent of the Oracle Solaris OS version that runs on the primary domain. Therefore, if you run the Oracle Solaris 10 OS in the primary domain, you can still run the Oracle Solaris 11 OS in a guest domain. Likewise if you run the Oracle Solaris 11 OS in the primary domain, you can still run the Oracle Solaris 10 OS in a guest domain.

The only difference between running the Oracle Solaris 10 OS or the Oracle Solaris 11 OS on the primary domain is the feature difference in each OS.

Important preinstallation information for SF Oracle RAC

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

- Preinstallation checklist for your configuration. Go to the SORT installation checklist tool. From the drop-down lists, select the information for the Symantec product you want to install, and click **Generate Checklist**.

- Hardware compatibility list for information about supported hardware: http://www.symantec.com/docs/TECH211575

- For important updates regarding this release, review the Late-Breaking News Technote on the Symantec Technical Support website: http://www.symantec.com/docs/TECH211540

- Latest information on support for Oracle database versions: http://www.symantec.com/docs/DOC5081

- 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.

Table 1-3   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 1-3   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.
	You can also configure aggregated interfaces.
Fiber Channel or SCSI host bus adapters	At least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

Table 1-4 lists the hardware requirements for campus clusters in addition to the basic cluster requirements.

Table 1-4   Hardware requirements for campus clusters

Item	Description
Storage	▪ The storage switch (to which each host on a site connects) must have access to storage arrays at all the sites.
	▪ Volumes must be mirrored with storage allocated from at least two sites.
	▪ DWDM links are recommended between sites for storage links. DWDM works at the physical layer and requires multiplexer and de-multiplexer devices.
	▪ The storage and networks must have redundant-loop access between each node and each storage array to prevent the links from becoming a single point of failure.
Network	▪ Oracle requires that all nodes use the IP addresses from the same subnet.
	▪ Symantec recommends 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 Solaris operating systems

This section lists the supported operating systems for this release of Symantec products. For current updates, visit the Symantec Operations Readiness Tools Installation and Upgrade page: https://sort.symantec.com/land/install_and_upgrade.

Table 1-5 shows the supported operating systems for this release.

Table 1-5   Supported operating systems

Operating systems	Levels	Chipsets
Solaris 10	Update 9, 10, and 11	SPARC
Solaris 11	Solaris 11.1 and up to Support Repository Updates (SRU) 11.1.21.4.1	SPARC

This release (version 6.1) is not supported on the x86-64 architecture.

This release (version 6.1) supports Solaris and Solaris10 branded zones on the Solaris 11 operating system and native brand zones on the Solaris 10 operating system.

For SF Oracle RAC, all nodes in the cluster need to have the same operating system version and update level.

Supported database software

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

http://www.symantec.com/docs/DOC5081

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

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

https://support.oracle.com

Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x6000 (for Solaris 10) and 0x8000 (for Solaris 11). When you install the Veritas File System package, `VRTSvxfs`, the VRTSvxfs packaging scripts check the values of these variables in the kernel. If the values are less than the required values, VRTSvxfs increases the

values and modifies the `/etc/system` file with the required values. If the VRTSvxfs scripts increase the values, the installation proceeds as usual except that you must reboot and restart the installation program. A message displays if a reboot is required.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System. Use the following commands to check the values of the variables:

```
For Solaris 10:    # echo "lwp_default_stksize/X" | mdb -k
                   lwp_default_stksize:
                   lwp_default_stksize:           6000

                   # echo "svc_default_stksize/X" | mdb -k
                   svc_default_stksize:
                   svc_default_stksize:           6000


For Solaris 11:    # echo "lwp_default_stksize/X" | mdb -k
                   lwp_default_stksize:
                   lwp_default_stksize:           8000

                   # echo "svc_default_stksize/X" | mdb -k
                   svc_default_stksize:
                   svc_default_stksize:           8000
```

If the values shown are less than 6000 (for Solaris 10) and less than 8000 (for Solaris 11), you can expect a reboot after installation.

Note: The default value of the `svc_default_stksize` variable is 0 (zero), which indicates that the value is set to the value of the `lwp_default_stksize` variable. In this case, no reboot is required, unless the value of the `lwp_default_stksize` variable is too small.

To avoid a reboot after installation, you can modify the `/etc/system` file with the appropriate values. Reboot the system prior to installing the packages. Add the following lines to the `/etc/system` file:

```
For Solaris 10:    set lwp_default_stksize=0x6000
                   set rpcmod:svc_default_stksize=0x6000


For Solaris 11:    set lwp_default_stksize=0x8000
                   set rpcmod:svc_default_stksize=0x8000
```

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
- Volume Replicator
- Oracle Data Guard

Fixed issues

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

Issues fixed in SF Oracle RAC 6.1

Table 1-6 lists the issues fixed in SF Oracle RAC 6.1.

Table 1-6   Issues fixed in SF Oracle RAC 6.1

Incident	Description
3090447	The CRSResource agent does not support the C shell (csh) environment.

2873102	When you install, configure, or uninstall SF Oracle RAC, the installer prompts you to optionally upload installation logs to the Symantec Web site. If the installer encounters connectivity problems, you may see an error similar to the following:
	
	 `Status read failed: Connection reset by peer at`
	 `<media_path>/../perl/lib/5.14.2/Net/HTTP/Methods.pm line 269.`

3186840	Running the command `zoneadm -u` to upgrade zones removes some of the package binaries.
2851403	Veritas File System modules may fail to unload if SmartMove is enabled and a break-off snapshot volume has been reattached.
2689195	File system check daemon fails to restart after abnormal termination.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues

Table 1-7 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-7   SFDB tools fixed issues

Incident	Description
2591463	Database Storage Checkpoint unmount may fail with device busy.
2534422	FlashSnap validate reports snapshot unsplittable.
2580318	dbed_vmclonedb ignores new clone SID value after cloning once.
2579929	User authentication fails.
2479901	FlashSnap resync fails if there is an existing space-optimized snapshot.
2869268	Checkpoint clone fails in a CFS environment if cloned using same checkpoint and same clone name on both nodes.
2849540	Very long off-host cloning times for large number of data files.
2715323	SFDB commands do not work with the ZHS16GBK character set.

LLT, GAB, and I/O fencing fixed issues

Table 1-8 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-8   LLT, GAB, and I/O fencing fixed issues

Incident	Description
2869763	When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/llttab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files.
2991093	The preferred fencing node weight does not get reset to the default value when HAD is terminated. In spite of lack of high availability on that node, fencing may give preference to that node in a network partition scenario.
2995937	The default value of preferred fencing node weight that vxfen uses is 1 (one). However, when HAD starts without any service group or if HAD is stopped or terminated, the node weight is reset to 0 (zero). Since vxfen resets the preferred fencing weight to its default value when HAD gets terminated, stopping HAD and killing HAD shows different preferred fencing weight.

Table 1-8   LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
3025931	There is a corner case where while shutting down the system, if the GAB service stop script is not run successfully, then, on the next reboot the GAB service fails to load. The GAB driver remains added into the system but module does not get loaded. In such cases, devlink entries are not created for the GAB driver and configuration of gab fails.
2802682	Server-based fencing may fail to start if you use the existing configuration files after reinstalling the stack.
2858190	If VRTSvxfen package is not installed on the system, then certain script files that are needed for the vxfentsthdw utility to function are not available. So, without the VRTSvxfen package installed on the system you cannot run the utility from the install media.
2724565	In SFRAC environments, sometimes GAB might fail to start because of the race between GAB and LMX in calling add_drv.
3140359	Port a does not come up due to race between `gabconfig -cx` and `gabconfig -x`.
3101262	GAB queue is overloaded causing memory pressure during I/O shipping.
3218714	GAB does not log messages about changing tunable values.
2858076	Changing the module parameter `gab_conn_wait` had no effect.

Known issues

This section covers the known issues in this release.

For Oracle RAC issues:

See "Oracle RAC issues" on page 32.

For SF Oracle RAC issues:

See "SF Oracle RAC issues" on page 35.

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 *Symantec 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

Enabling ODM in Oracle RAC 11 Release 2 installations causes errors (1913013)

Enabling ODM in Oracle RAC 11 Release 2 installations causes the following error:

```
'ODM ERROR V-41-4-1-253-12 Not enough space'
Oracle instance may also crash with same error.
```

The error is observed if the DISM (Dynamic Intimate Shared memory) feature is enabled. In Solaris, the Oracle database uses DISM if it is available on the system, and if the value of the `sga_max_size` initialization parameter is larger than the size required for all SGA components combined.

Workaround: Make sure that the file `ORACLE_HOME/bin/oradism` is owned by the root user with "execute" and "setuid" permissions. If the problem persists after correcting the permissions, uncomment the `sga_max_size` and `memory_target` `init.ora` parameters.

Oracle VIP Configuration Assistant fails with an error message (1182220)

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

```
The given interface(s), "net0" is not public.
Public interfaces should be used to configure virtual IPs.
```

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

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

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

Oracle Cluster Verification utility displays a warning message

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

For example:

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

```
3. Select the failed assistants and click the 'Retry' button
to retry them.
=============================================================
```

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

Oracle Database Configuration Assistant displays an error

The Database Configuration Assistant utility displays the following error:

```
SGA size cannot be greater than maximum shared memory
segment size (0).
```

Workaround: Ignore this message and manually configure the database memory parameters for Oracle. In the "Memory" tab of the Oracle Database Creation Assistant (DBCA), select a Custom and Manual shared memory management configuration and enter the appropriate values.

CP server service group fails to come online with the default database path after the CP server is upgraded from 6.0 to 6.1 on a multi-node cluster [3326639]

If the CP server is configured on a multi-node cluster before the upgrade with security enabled, you must reconfigure the CP server after the CP server upgrade. If you reuse the old credentials with the old database path, the CP server service group does not come online. Since the default database paths of CP server in 6.0 and 6.1 are different, reusing the old credentials and default database path prevents the CP server service group from coming online.

Workaround:

If the CP server multi-node cluster is configured with security enabled and if the old credentials such as database path are expected to be reused in reconfiguration of the CP server after the upgrade of the CP server, use the same database path before and after the upgrade.

SF Oracle RAC issues

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

Installation known issues

This section describes the known issues during installation and upgrade.

Verification of Oracle binaries incorrectly reports as failed during Oracle Grid Infrastructure installation

The verification of Oracle binaries may incorrectly report as failed during the Oracle Grid Infrastructure installation using the SF Oracle RAC installer. The message is erroneously reported due to a break in passwordless SSH communication. The SSH communication fails because execution of the `root.sh` script changes the owner of the operating system root directory to the grid user directory.

On Solaris 11 non-default ODM mount options will not be preserved across package upgrade (2745100)

On Solaris 11, before the package upgrade if Oracle Disk Manager (ODM) is mounted with non-default mount options such as nocluster, nosmartsync, etc., these mount options will not get preserved after package upgrade.

There is no workaround at this time.

On Solaris 10 xprtld will not be started if user use jumpstart to install product (3325954)

If you install the operating system plus the Symantec product using the JumpStart method and after installation, reboot the machine then configure and start the product, all the processes will be started except for `xprtld` process.

Workaround:

After reboot, manually execute the following command to start `xprtld`:

```
# /opt/VRTSsfmh/adm/xprtldctrl start
```

SF Oracle RAC installer does not support use of `makeresponsefile` option (2577669)

The SF Oracle RAC installer does not support the use of `makeresponsefile` option for configuring Oracle RAC settings. The following message is displayed when you attempt to configure Oracle RAC using the option:

```
Currently SFRAC installer does not support -makeresponsefile option.
```

Workaround: Configure Oracle RAC by editing the response file manually.

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups [2574731]

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
```

After Live Upgrade to Solaris 10 Update 10, boot from alternate boot environment may fail (2370250)

If your setup involves volumes in a shared disk group that are mounted as CFS in a cluster, then during Live Upgrade using the `vxlustart` command from any supported Solaris version to Solaris 10 Update 10, boot from an alternate boot environment may fail.

Workaround: Run the `vxlufinish` command. Before rebooting the system, manually delete the entries of all the volumes of shared disks that are mounted as CFS in the `/altroot.5.10/etc/vfstab` directory.

On Solaris 10, a flash archive installed through JumpStart may cause a new system to go into maintenance mode on reboot (2379123)

If a Flash archive is created on a golden host with encapsulated root disks, when this Flash archive is installed onto another host through JumpStart, the new system may go to maintenance mode when you initially reboot it.

This problem is caused by the predefined root disk mirror in the Flash archive. When the archive is applied to a clone system, which may have different hard drives, the newly cloned system may get stuck at root disk mirroring during reboot.

Workaround: Create the Flash archive on a golden host with no encapsulated root disks. Run `vxunroot` to clean up the mirrored root disks before you create the Flash archive.

Flash Archive installation not supported if the target system's root disk is encapsulated

Symantec does not support SF Oracle RAC installation using Flash Archive if the target system's root disk is encapsulated.

Make sure that the target system's root disk is unencapsulated before starting the installation.

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 performing a manual 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 manual 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 not 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:

```
# vxdctl -c mode
```

2 On the CVM master node, upgrade the CVM protocol:

```
# vxdctl upgrade
```

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.

Erroneous resstatechange trigger warning [2277819]

You may encounter the following warning when you restart resources:

```
CPI WARNING V-9-40-4317 The installer has detected that resstatechange
trigger is configured by setting TriggerResStateChange attributes.
```

Workaround: In future releases, the resstatechange trigger will not be invoked when a resource is restarted. Instead, the resrestart trigger will be invoked if you set the TriggerResRestart attribute. The resrestart trigger is available in the current release. Refer to the VCS documentation for details.

Upgrade with zones installed on CFS is not supported if CFS is under VCS control (3322276)

If CFS is under VCS control, then upgrade with zones installed on CFS is not supported if you perform phased upgrade.

Workaround: Unmount the CFS before performing the phased upgrade. After the upgrade is complete, re-mount the CFS and reinstall the zone(s).

Rolling upgrade may encounter a problem if open volumes from different disk groups have the same name (3326196)

When you perform a rolling upgrade, the installer may block the rolling upgrade even if all the open volumes are under VCS control. This may occur if there are volumes with the same name under different disk groups although they are not mounted.

Workaround: Avoid creating volumes from different disk groups with the same name. If they already exist, umount all the VxFS mount points. After the upgrade is finished, remount the volumes.

If you select rolling upgrade task from the Install Bundles menu, the CPI exits with an error (3442070)

If you try to perform rolling upgrade using Install Bundles and select the rolling upgrade task from the Install Bundle menu, the CPI exits with an error.

Workaround: Run the installer with `-rolling_upgrade` option instead of choosing the task from the menu.

```
# ./installer -hotfix_path /path/to/hotfix -rolling_upgrade
```

Upgrade fails because there is zone installed on the VxFS file system which is offline. The packages in the zone are not updated. (3319753)

If the zone installed on VxFS file system is under VCS control, and the VxFS file system is in offline state, the upgrade fails because it's not able to update the packages in the zones.

Workaround:

Check the status of the mounted file system which has the zones on it. If the file system is offline, you need to first bring it online, then do the upgrade, so that the packages in the local zone can be updated.

A Storage Foundation ABE upgrade fails while upgrading from 6.1PR1 to 6.1 (3325054)

If you perform an upgrade from Symantec Storage Foundation for Oracle RAC 6.1 PR1 to 6.1 in a Solaris Alternate Boot Environment (ABE), you may see the following error when you uninstall 6.1 PR1:

```
pkg: An unexpected error happened during uninstall:
[('/mnt/etc/vx/vxesd/vxesd.socket',
'/mnt/var/pkg/lost+found/etc/vx/vxesd-20130927T101550Z/vxesd.socket',
"[Errno122] Operation not supported on transport endpoint:
'/mnt/etc/vx/vxesd/vxesd.socket'")]
```

Workaround: Remove */mnt*/etc/vx/vxesd/vxesd.socket, where */mnt* is the ABE mountpoint. After you remove the socket, you can successfully uninstall 6.1 PR1 and install 6.1.

If you choose to upgrade nodes without zones first, the rolling upgrade or phased upgrade is not blocked in the beginning, but fails later (3319961)

If you choose to upgrade nodes without zones first, the rolling upgrade or phased upgrade is not blocked in the beginning, but fails later when you start to upgrade the nodes that have zones installed.

This issue occurs in the following scenarios:

■ A zone is installed on a Cluster File System (CFS) on one of the nodes.

■ A node is installed on a Veritas File System (VxFS) on one of the nodes, and node is under Symantec Cluster Server (VCS) control.

Workaround:

1 Before you upgrade, uninstall the zones on the nodes which have zones installed. Enter:.

   zoneadm -z *zonename* uninstall

2 Run the installer to run the upgrade.

3 After the upgrade completes, reinstall the zones.

LLT known issues

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

Cannot configure LLT if full device path is not used in the llttab file (2858159)

(Oracle Solaris 11) On virtual machines ensure that you use the full path of the devices corresponding to the links in llttab. For example, use /dev/net/net1 instead of /dev/net/net:1 in the llttab file, otherwise you cannot configure LLT.

Fast link failure detection is not supported on Solaris 11 (2954267)

Fast link failure detection is not supported on Solaris 11 operating system because the operating system cannot provide notification calls to LLT when a link failure occurs. If the operating system kernel notifies LLT about the link failure, LLT can detect a link failure much earlier than the regular link failure detection cycle. As Solaris 11 does not notify LLT about link failures, failure detection cannot happen before the regular detection cycle.

Workaround: None

GAB known issues

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.

Workaround: There is no workaround for this issue.

GAB may fail to stop during a phased upgrade on Oracle Solaris 11 (2858157)

While performing a phased upgrade on Oracle Solaris 11 systems, GAB may fail to stop. However, CPI gives a warning and continues with stopping the stack.

Workaround: Reboot the node after the installer completes the upgrade.

Cannot run pfiles or truss files on gablogd (2292294)

When pfiles or truss is run on gablogd, a signal is issued to gablogd. gablogd is blocked since it has called an gab ioctl and is waiting for events. As a result, the pfiles command hangs.

Workaround: None.

(Oracle Solaris 11) On virtual machines, sometimes the common product installer (CPI) may report that GAB failed to start and may exit (2879262)

GAB startup script may take longer than expected to start up. The delay in start up can cause the CPI to report that GAB failed and exits.

Workaround: Manually start GAB and all dependent services.

I/O fencing known issues

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

Rolling upgrade of VCS from pre-6.0 versions fails with CP server in secure mode [3262900]

If the CP server is configured in secure mode, rolling upgrade of VCS from versions lower than 6.0 to 6.1 is not supported. Since the `vxcpserv` process is not compatible with shared authentication, CP server service group fails to come online after performing phase 1 of the rolling upgrade.

Workaround: Use full upgrade or phased upgrade instead of rolling upgrade.

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@sys1,
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 cluster, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

When I/O fencing is not up, the svcs command shows VxFEN as online (2492874)

Solaris 10 SMF marks the service status based on the exit code of the start method for that service. The VxFEN start method executes the vxfen-startup script in the background and exits with code 0. Hence, if the vxfen-startup script subsequently exits with failure then this change is not propagated to SMF. This behavior causes the `svcs` command to show incorrect status for VxFEN.

Workaround: Use the `vxfenadm` command to verify that I/O fencing is running.

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 information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster 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 *Symantec 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 at least 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: There is no workaround for this issue.

CP server configuration fails while setting up secure credentials for CP server hosted on an SFHA cluster (2621029)

When you configure CP server using the `configure_cps.pl` utility, the configuration fails while setting up secure credentials for CP server that is hosted on an SFHA cluster. You may see the following error:

```
Creating softlink to credential directory /etc/VRTScps/db/CPSERVER
on node nodename.
Unable to connect to node nodename using /usr/bin/ssh.
Please configure ssh communication and retry. Exiting.
```

Workaround: You can use any of the following options:

- Before running the `configure_cps.pl` utility, change the default shell for root user to either KSH or bash.

- Perform the following steps after running the `configure_cps.pl` utility on each node of the cluster:

  - Manually remove the old credential directory or softlink. For example:

    ```
    # rm -rf /var/VRTSvcs/vcsauth/data/CPSERVER
    ```

  - Create a new soft-link to the shared location of the credential directory:

    ```
    # ln -s path_of_CP_server_credential_directory \
     /var/VRTSvcs/vcsauth/data/CPSERVER
    ```

  - Start the CPSSG service group:

    ```
    # hagrp -online CPSSG -any
    ```

The cpsadm command fails after upgrading CP server to 6.0 or above in secure mode (2846727)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSat package 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: To resolve this issue, perform the following procedure on all of the nodes of the CP server:

1   Rename `cpsadm` to `cpsadmbin`:

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

2   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 "$@"
```

3   Change the permissions of the new file to 775:

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

Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later [3226290]

The issue exists because the VCS 5.1SP1 release version does not support separate directories for truststores. However, VCS version 6.0 and later support separate directories for truststores. Because of this mismatch in support for truststores, you cannot set up trust between client systems and servers.

Workaround: Set up trust manually between the coordination point server and client systems using the `cpsat` or `vcsat` command so that the servers and client systems can communicate in a secure mode.

Hostname and username are case sensitive in CP server (2846392)

The hostname and username on the CP server are case sensitive. The hostname and username used by fencing to communicate with CP server must be in same case as present in CP server database, else fencing fails to start.

Workaround: Make sure that the same case is used in the hostname and username on the CP server.

Fencing may show the RFSM state as replaying for some nodes in the cluster (2555191)

Fencing based on coordination point clients in Campus cluster environment may show the RFSM state as replaying for some nodes in the cluster.

Workaround:

Restart fencing on the node that shows RFSM state as replaying.

Coordination point server-based fencing may fail if it is configured on 5.1SP1RP1 using 6.0.1 coordination point servers (2824472)

The 5.1SP1 installer (CPI) cannot set up trust between a 5.1SP1 client and a 6.0 or later server, because there are no separate directories for truststores in the 5.1SP1. When trust cannot be setup, the 5.1SP1 installer cannot configure 5.1SP1 clients to work with 6.0 or later CPS in secure mode.

Workaround:

Set up trust manually between the CPS and clients using the cpsat or the vcsat command. After that, CPS and client will be able to communicate properly in the secure mode.

The upper bound value of FaultTolerance attribute of CoordPoint agent should be less than the majority of the coordination points. (2846389)

The upper bound value of `FaultTolerance` attribute of `CoordPoint` agent should be less than the majority of the coordination points. Currently this value is less than the number of coordination points.

The vxfenswap utility deletes comment lines from the `/etc/vxfemode` file, if you run the utility with hacli option (3318449)

The vxfenswap utility uses RSH, SSH, or hacli protocol to communicate with peer nodes in the cluster. When you use vxfenswap to replace coordination disk(s) in disk-based fencing, vxfenswap copies `/etc/vxfenmode` (local node) to `/etc/vxfenmode` (remote node).

With the hacli option, the utility removes the comment lines from the remote `/etc/vxfenmode` file, but, it retains comments in the local `/etc/vxfenmode` file.

Workaround: Copy the comments manually from local `/etc/vxfenmode` to remote nodes.

When you configure CP server only for HTTPS-based communication, the `engine_A.log` displays a misleading message (3321101)

The `engine_A.log` file displays the following message when you configure CP server only for HTTPS-based communication but not for IPM-based communication.

```
No VIP for IPM specified in /etc/vxcps.conf
```

Workaround: Ignore the message.

The CoordPoint agent faults after you detach or reattach one or more coordination disks from a storage array (3317123)

After you detach or reattach a coordination disk from a storage array, the CoordPoint agent may fault because it reads an older value stored in the I/O fencing kernel module.

Workaround: Run the `vxfenswap` utility to refresh the registration keys on the coordination points for both server-based I/O fencing and disk-based I/O fencing. But, even if the registrations keys are not lost, you must run the `vxfenswap` utility to refresh the coordination point information stored in the I/O fencing kernel module.

For more information on refreshing registration keys on the coordination points for server-based and disk-based I/O fencing, refer to the *Symantec Cluster Server Administrator's Guide*.

Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups [3335137]

If SysDownPolicy of one or more online service groups is configured to AutoDisableNoOffline, fencing configurations such as server-based, disk-based and disable mode fail. Since the service groups is configured with `SysDownPolicy = { AutoDisableNoOffline }`, stopping VCS fails which leads to the failure of fencing configuration.

Workaround: When configuring fencing and before stopping VCS, you must offline the service groups configured with `SysDownPolicy = { AutoDisableNoOffline }` manually.

CP server does not allow adding and removing HTTPS virtual IP or ports when it is running [3322154]

CP server does not support adding and removing HTTPS virtual IPs or ports while the CP server is running. However, You can add or remove the IPM virtual IPs or ports.

Workaround: No workaround. If you want to add a new virtual IP for HTTPS, you must follow the entire manual procedure for generating HTTPS certificate for the CP server (server.crt), as documented in the *Symantec Cluster Server Installation Guide*.

The `vxfentsthdw` utility may not run on systems installed with partial SFHA stack [3333914]

The `vxfentsthdw` utility runs if the SFHA stack and VCS are fully installed with properly configured SF and VxVM. It also runs if the entire SFHA stack and VCS are not installed. However, partial installs where SF is installed and configured but

VCS is not installed is not supported. The utility will display an error with the `-g` or `-c` options.

Workaround: Install VRTSvxfen package, then run the utility from either the install media or from the `/opt/VRTSvcs/vxfen/bin/` location.

When a client node goes down, for reasons such as node panic, I/O fencing does not come up on that client node after node restart (3341322)

This issue happens when one of the following conditions is true:

- Any of the CP servers configured for HTTPS communication goes down.

- The CP server service group in any of the CP servers configured for HTTPS communication goes down.

- Any of the VIPs in any of the CP servers configured for HTTPS communication goes down.

When you restart the client node, fencing configuration starts on the node. The fencing daemon, vxfend, invokes some of the fencing scripts on the node. Each of these scripts has a timeout value of 120 seconds. If any of these scripts fails, fencing configuration fails on that node.

Some of these scripts use `cpsadm` commands to communicate with CP servers. When the node comes up, `cpsadm` commands try to connect to the CP server using VIPs for a timeout value of 60 seconds. So, if the multiple `cpsadm` commands that are run within a single script exceed the timeout value, then the total timeout value exceeds 120 seconds, which causes one of the scripts to time out. Hence, I/O fencing does not come up on the client node.

Note that this issue does not occur with IPM-based communication between CP server and client clusters.

Workaround: Fix the CP server.

spfile created on VxFS and ODM may contain uninitialized blocks at the end (3320906)

spfile created on VxFS and ODM may contain uninitialized blocks at the end due to space allocation with file system block size alignment. This is harmless and does not cause any problem to Oracle startup.

Installation of Oracle Clusterware using Oracle response file fails (3321004)

The installation of Oracle Clusterware using Oracle response file fails with the following error:

```
There are issues using the DISPLAY value you provided.
Either the DISPLAY variable has not been set properly or
there are display connectivity problems.
```

This is because Oracle Clusterware response file does not require the DISPLAY environment variable whereas the SF Oracle RAC installer requires it.

Workaround: Before starting the SF Oracle RAC installer, export the DISPLAY environment variable as follows:

```
$ Export DISPLAY=10.200.58.255:4
```

PrivNIC and MultiPrivNIC agents not supported with Oracle RAC 11.2.0.2 and later versions

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

For more information, see the following Technote:

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

CSSD agent forcibly stops Oracle Clusterware if Oracle Clusterware fails to respond (3352269)

On nodes with heavy load, the CSSD agent attempts to check the status of Oracle Clusterware till it reaches the `FaultOnMonitorTimeouts` value. However, Oracle Clusterware fails to respond and the CSSD agent forcibly stops Oracle Clusterware. To prevent the CSSD agent from forcibly stopping Oracle Clusterware, set the value of the `FaultOnMonitorTimeouts` attribute to 0 and use the `AlertOnMonitorTimeouts` attribute as described in the following procedure.

Perform the following steps to prevent the CSSD agent from forcibly stopping Oracle Clusterware:

1   Change the permission on the VCS configuration file to read-write mode:

```
# haconf -makerw
```

2   Set the `AlertOnMonitorTimeouts` attribute value to 4 for the CSSD resource:

```
# hatype -display CSSD | grep AlertOnMonitorTimeouts
CSSD  AlertOnMonitorTimeouts 0
# hares -override cssd_resname AlertOnMonitorTimeouts
# hatype -modify CSSD AlertOnMonitorTimeouts 4
```

3 Set the `FaultOnMonitorTimeouts` attribute value to 0 for the CSSD resource:

```
# hatype -display CSSD | grep FaultOnMonitorTimeouts
CSSD  FaultOnMonitorTimeouts 4
# hares -override cssd_resname FaultOnMonitorTimeouts
# hatype -modify CSSD FaultOnMonitorTimeouts 0
```

4 Verify the `AlertOnMonitorTimeouts` and `FaultOnMonitorTimeouts` settings:

```
# hatype -display CSSD | egrep "AlertOnMonitorTimeouts|FaultOnMonitorTime
CSSD  AlertOnMonitorTimeouts 4
CSSD  FaultOnMonitorTimeouts 0
```

5 Change the permission on the VCS configuration file to read-only mode:

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

Unable to configure Highly Available IP (HAIP) using the web installer (3348812)

During HAIP configuration, the SF Oracle RAC web installer fails to update the `/etc/hosts` file with the HAIP alias.

Workaround:

Use one of the following options:

- Use the SF Oracle RAC script installer to configure HAIP.

- Use the SF Oracle RAC web installer, but add the IP addresses and aliases in the `/etc/hosts` file manually.

Intelligent Monitoring Framework (IMF) entry point may fail when IMF detects resource state transition from online to offline for CSSD resource type (3287719)

When IMF detects a state transition from ONLINE to OFFLINE state for a registered online resource, it sends a notification to the CSSD agent. The CSSD agent schedules a monitor to confirm the state transition of the resource. The resources of type CSSD takes more time to go online or offline fully. Therefore, if this immediate monitor finds the resource still in online state, it assumes that the IMF notification is false and attempts to register the resource in online state again.

In such partial state transitions, the agent repeatedly attempts to register the resource until the `RegisterRetryLimit` is reached (default value is 3) or the resource

registration is successful. After the resource is completely offline, the next resource registration with IMF will be successful.

Workaround: Increase the value of the `RegisterRetryLimit` attribute if multiple registration attempts fail.

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.

Warning message displayed on taking cssd resource offline if LANG attribute is set to "eucJP" (2123122)

When you take the cssd resource offline using the `hares -offline cssd` command and the LANG attribute is set to "eucJP", the following message may be observed in the `hamsg engine_A` command output:

```
VCS INFO V-16-2-13716 Could not find message V-16-2-13716
```

You may ignore the message.

PrivNIC resource faults in IPMP environments on Solaris 11 systems (2838745)

The PrivNIC resource faults on Solaris 11 systems when private interfaces used by IPMP are configured under PrivNIC resource.

Workaround: Avoid using PrivNIC or MultiPrivNIC agents in IPMP environments.

Error displayed on removal of VRTSjadba language package (2569224)

Removal of the VRTSjadba language package displays the following error on the screen:

```
Executing postremove script.
Generating BMC map file...
bmcmap ERROR V-33-1000-10001 Unable to create BMC map
```

You may ignore the error.

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 [2653695]

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.

Oracle service group faults on secondary site during failover in a disaster recovery scenario [2653704]

Oracle service group fails to go online in the DR site when disaster strikes the primary site. This happens if the AutoFailover attribute on the Service Group is set to 1 and when the corresponding service group's FireDrill is online in the DR site. Firedrill Service group may remain ONLINE on the DR site.

Workaround: If the service group containing the Oracle (or any database) resource faults after attempting automatic DR failover while FireDrill is online in the DR site, manually offline the FireDrill Service Group. Subsequently, attempt the online of the Oracle Service Group in the DR site.

Veritas Volume Manager can not identify Oracle Automatic Storage Management (ASM) disks (2771637)

Veritas Volume Manager (VxVM) commands can not identify disks that are initialized by ASM. Administrators must use caution when using the VxVM commands to avoid accidental overwriting of the ASM disk data.

When all Primary/Optimized paths between the server and the storage array are disconnected, ASM disk group dismounts and the Oracle database may go down (3289311)

The Oracle database shows an I/O error on the control file, but there was no I/O error seen on any DMP device. When all Primary/Optimized paths are disconnected, DMP fails over to other available paths but the failover takes time. In the meantime, the application (ASM/Oracle database) times out the I/O.

The ASM alert log file displays messages such as the following:

```
Errors in file /u01/app/oracle/diag/rdbms/orcl/orcl2/trace/orcl2_ckpt_6955.trc:
ORA-00221: error on write to control file
ORA-00206: error in writing (block 4, # blocks 1) of control file
ORA-00202: control file: '+DATA_P6/ORCL/CONTROLFILE/current.261.826783133'
ORA-15081: failed to submit an I/O operation to a disk
ORA-15081: failed to submit an I/O operation to a disk
Wed Oct 09 14:16:07 2013
WARNING: group 2 dismounted: failed to read virtual extent 0 of file 261
Wed Oct 09 14:16:07 2013
USER (ospid: 6955): terminating the instance due to error 221
Wed Oct 09 14:16:07 2013
WARNING: requested mirror side 2 of virtual extent 0 logical extent 1 offset
16384
is not allocated; I/O request failed
```

```
WARNING: requested mirror side 3 of virtual extent 0 logical extent 2 offset
16384
is not allocated; I/O request failed
```

The above issue may occur when the server is configured as follows:

DB: Oracle 12c

Volume Manager: ASM

Multipathing Solutions: DMP

OS: Solaris

Disk Array : HP EVA in ALUA mode

Workaround:

The following workaround can reduce the probability of this issue, and when you see this issue, you could use Oracle commands to start the database manually.

Increase the application time out and make the following changes to reduce the time taken to mark the path as offline:

- In the `/kernel/drv/fp.conf` file, add `fp_offline_ticker=15`.
- In the `/kernel/drv/fcp.conf` file, add `fcp_offline_delay=10`.

Oracle Universal Installer fails to start on Solaris 11 systems (2784560)

The Oracle Universal Installer (OUI) fails to start when the SF Oracle RAC installer invokes the OUI for the installation of Oracle Clusterware/Grid Infrastructure software.

Workaround: Install the following packages before installing Oracle Clusterware/Grid Infrastructure.

```
SUNWxwplt
SUNWmfrun
```

For instructions, see the Oracle documentation.

Symantec Storage Foundation for Databases (SFDB) tools known issues

The following are known issues in this release of Symantec Storage Foundation for Databases (SFDB) tools.

Sometimes SFDB may report the following error message: SFDB remote or privileged command error (2869262)

While using SFDB tools, if you attempt to run commands, such as `dbed_update` then you may observe the following error:

```
$ /opt/VRTSdbed/bin/dbed_update
No repository found for database faildb, creating new one.
SFDB vxsfadm ERROR V-81-0450 A remote or privileged command could not
be executed on swpa04

Reason: This can be caused by the host being unreachable or the vxdbd
daemon not running on that host.

Action: Verify that the host swpa04 is reachable. If it is, verify
that the vxdbd daemon is running using the /opt/VRTS/bin/vxdbdctrl
status command, and start it using the /opt/VRTS/bin/vxdbdctrl start
command if it is not running.
```

Workaround: There is no workaround for this issue.

The information file that is generated after a DBED data collector operation reports an error (2795490)

When the VRTSexplorer DBED scripts use the old VRTSdbms3-specific scripts that are removed from the products, the information file reports the following error:

```
/opt/VRTSdbms3/vxdbms_env.sh: cannot open [No such file or directory]
```

Workaround:

1 Run the `cd /opt/VRTSspt/DataCollector/sort` command. If this directory does not exist, run `sh /opt/VRTSspt/DataCollector/*.sh`.

2 Run the `cd advanced/lib/VOS/v10/Collector/VxExpCollector/explorer_scripts` command.

3 In `dbed_rept_sql`, comment

$**VXDBMS_DIR/vxdbms_env.sh**

Or

Replace **$VXDBMS_DIR/vxdbms_env.sh** with

```
[[ -f $VXDBMS_DIR/vxdbms_env.sh ]] &&
    {
        . $VXDBMS_DIR/vxdbms_env.sh
    }
```

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SF Oracle RAC. There is no workaround at this point of time.

The dbdst_obj_move(1M) command moves all the extents of a database table (3277003)

The `dbdst_obj_move(1M)` command moves all the extents of a database table when:

■ The `dbdst_obj_move(1M)` command is run from the CFS secondary node.

■ The object is an Oracle database table (-t option)

■ A range of extents is specified for movement to a target tier (-s and -e options). The `dbdst_obj_move(1M)` command moves all extents of the specified table to a target tier when the extent size is greater than or equal to 32768. However, the expectation is to move only a specified range of extents.

Workaround: Run the `dbdst_obj_move(1M)` command from the CFS primary node.

Use the `fsclustadm showprimary` *<mountpoint>* and `fsclustadm idtoname` *<nodeid>*commands to determine the mode of a CFS node.

When you attempt to move all the extents of a table, the `dbdst_obj_move`(1M) command fails with an error (3260289)

When you attempt to move all the extents of a database table, which is spread across multiple mount-points in a single operation, the `dbdst_obj_move`(1M) command fails. The following error is reported:

```
bash-2.05b$ dbdst_obj_move -S sdb -H $ORACLE_HOME -t test3 -c MEDIUM
FSPPADM err : UX:vxfs fsppadm: WARNING: V-3-26543: File handling failure
on /snap_datadb/test03.dbf with message -
SFORA dst_obj_adm ERROR V-81-6414 Internal Error at fsppadm_err
```

Note: To determine if the table is spread across multiple mount-points, run the dbdst_obj_view(1M) command

Workaround: In the `dbdst_obj_move`(1M) command, specify the range of extents that belong to a common mount-point. Additionally, if your table is spread across "n" mount-points, then you need to run the `dbdst_obj_move`(1M) command "n" times with a different range of extents.

The ReverseResyncBegin (RRBegin) operation with recovery option as AUTO fails (3076583)

The RRBegin operation with the recovery option as AUTO fails when you perform the following sequence of operations:

1   Validate the FlashSnap setup using the validate operation.

2   In the database, take the tablespace offline.

3   Perform a snapshot operation.

4   Bring the tablespace online which was taken offline in 2.

5   Perform the Reverse Resync Begin operation.

Note: This issue is encountered only with Oracle version 10gR2.

Workaround: Perform one of the following:

- Make sure to bring the tablespace online only after performing the RRBegin and RRCommit operations. Otherwise, perform the Reverse Resync Begin operation while the tablespace is in the offline mode.

- To recover a database, specify the recovery option as **AUTO_UNTIL_SCN** in the RRBegin operation.

The ReverseResyncBegin (RRBegin) operation fails when performed on multiple snapshot configurations (3066532)

When you perform a Reverse Resync operation on multiple snapshot configurations, SFDB reports the following error message:

```
[oracle@dblxx64-3-vip3 ~]$ vxsfadm -a oracle -s flashsnap --name \
man -o rrbegin

SFDB vxsfadm ERROR V-81-0943 Repository already relocated to alternate
location.
```

As per the Reverse Resync design, the first RRBegin operation relocates the SFDB repository to a backup location, and the ReverseResyncAbort and ReverseResyncCommit operations restore it to the original location. When the second RRBegin operation attempts to relocate the same repository which is already relocated, SFDB reports the error message.

Workaround: Make sure to perform the RRAbort or RRCommit operation using the snapshot configuration that is in the RRBegin state.

Note: You must complete Reverse Resync operations for a particular configuration before you start with another configuration.

The ReverseResyncBegin (RRBegin) operation fails and reports an error message due to a missing binary control file (3157314)

When the RRBegin operation cannot find the binary control file that is used to recover a database instance, it reports the following error message:

```
[oracle@testbox ~]$ vxsfadm -a oracle -s flashsnap -name man -o rrbegin

SFDB vxsfadm ERROR V-81-0949 Binary Control file is not available for
recovery purposes
```

This issue is observed in the third-mirror break-off type (FlashSnap) snapshots that are created using the older SFDB version, which did not include the binary control file in the snapshot images.

Workaround:

There is no workaround for this issue.

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.

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.

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 of 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 Checkpoint, 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.

Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.1 (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 6.1.

When upgrading from SF Oracle RAC version 5.0 to SF Oracle RAC 6.1 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 (2844247)

If you have a parameter, such as `log_archive_dest_1`, in single line in the `init.ora` file, then `dbed_vmclonedb` works but `dbed_vmcloneb` fails if you put in multiple lines for parameter.

Workaround: Edit the PFILE to arrange the text so that the parameter values are on a single line. If the database uses a spfile and some parameter values are spread across multiple lines, then use the Oracle commands to edit the parameter values such as they fit in a single line.

Workaround

There is no workaround for this issue.

Data population fails after datafile corruption, rollback, and restore of offline checkpoint (2869259)

Sometimes when a datafile gets corrupted below its reservation size, the rollback may not pass and the file may not be rolled back correctly.

There is no workround at this point of time.

Checkpoint clone fails if the `archive log` destination is same as the datafiles destination (2869266)

Checkpoint cloning fails if the `archive log` destination is the same as the datafiles destination. The error is similar to:

```
Use of uninitialized value $path in hash element
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 121.
Use of uninitialized value $path in concatenation (.) or string
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 124.
Use of uninitialized value $path in pattern match (m//)
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 126.

SFDB vxsfadm ERROR V-81-0564 Oracle returned error.

Reason: ORA-02236: invalid file name (DBD ERROR: error possibly near
<*> indicator at char 172 in 'CREATE CONTROLFILE REUSE SET DATABASE
'TClone03' RESETLOGS NOARCHIVELOG
```

Workaround: For the 6.1 release, create distinct archive and datafile mounts for the checkpoint service.

FileSnap detail listing does not display the details of a particular snap (2846382)

FileSnap does not support displaying a detailed listing of a snapshot or clone. FileSnap only supports displaying a summary of all the snapshots or clones. For example, for the CLI `vxsfadm -s filesnap -a oracle --name=snap1 -o list`, a summary listing all the snapshots is displayed, instead of a detailed listing of a particular snapshot.

Workaround: There is no workaround for this issue.

Flashsnap clone fails under some unusual archivelog configuration on RAC (2846399)

In a RAC environment, when using FlashSnap, the archive log destination to snapshot must be a shared path, and must be the same across all the nodes. Additionally, all nodes must use the same archive log configuration parameter to specify the archive log destination. Configurations similar to the following are not supported:

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_2='location=/tpcc_arch'
tpcc3.log_archive_dest_3='location=/tpcc_arch'
```

Where tpcc1, tpcc2, and tpcc3 are the names of the RAC instances and /tpcc_arch
is the shared archive log destination.

Workaround: To use FlashSnap, modify the above configuration to
*.log_archive_dest_1='location=/tpcc_arch'. For example,

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_1='location=/tpcc_arch'
tpcc3.log_archive_dest_1='location=/tpcc_arch'
```

'vxdbd' process is online after Flash archive installation (2869269)

After a Flash archive installation of the SF stack, the `vxdbd` process is up, even if
the stack is not configured.

Workaround: You can ignore, or stop the `vxdbd` process using the
`/opt/VRTSdbed/common/bin/vxdbdctrl stop`command.

`sfua_rept_migrate` fails after phased SF Oracle RAC upgrade from 5.0MP3RP5 to 6.0.1 (2874322)

Command `sfua_rept_migrate` sometimes gives an error when upgrading to 6.0.1,
and fails to unmount the repository volume. The error message is similar to:

```
# ./sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
Unmounting SFUA Sybase ASA repository.
UX:vxfs umount: ERROR: V-3-26388: file system /rep has been mount
locked
SFORA sfua_rept_migrate ERROR V-81-5550 umount /dev/vx/dsk/repdg/repvol
failed.
SFORA sfua_rept_migrate ERROR V-81-9162 Failed to umount repository.
```

Workaround: The error does not hamper the upgrade. The repository migration
works fine, but the old repository volume does not get unmounted. Unmount the
mount using the manual option.

For example, use `/opt/VRTS/bin/umount -o mntunlock=VCS /rep`.

For more information, see TECH64812.

Instant mode clone fails in RAC environment for all FSMs with data loading (3517782)

When you use the instant clone mode for RAC databases, the clone operation may fail during Oracle recovery. The issue is more likely to be seen when there is load activity on some of the RAC nodes.

Workaround: Use either online or offline snapshot mode.

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 68.

Limitations of CSSD agent

The limitations of the CSSD agent are as follows:

- For Oracle RAC 11g Release 2 and later versions: The CSSD agent restarts Oracle Grid Infrastructure processes that you may manually or selectively take offline outside of VCS.
  Workaround: First stop the CSSD agent if operations require you to manually take the processes offline outside of VCS.
  For more information, see the topic "Disabling monitoring of Oracle Grid Infrastructure processes temporarily" in the *Symantec Storage Foundation for Oracle RAC Installation and Configuration Guide*.

- The CSSD agent detects intentional offline only when you stop Oracle Clusterware/Grid Infrastructure outside of VCS using the following command: `crsctl stop crs [-f]`. The agent fails to detect intentional offline if you stop Oracle Clusterware/Grid Infrastructure using any other command.
  Workaround: Use the `crsctl stop crs [-f]` command to stop Oracle Clusterware/Grid Infrastructure outside of VCS.

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.

Rolling upgrade from version 5.1 SP1 RP2 requires VCSMM point patch

Before performing rolling upgrades from 5.1 SP1 RP2, apply the VCSMM point patch on the SF Oracle RAC cluster.

For instructions, see: https://sort.symantec.com/patch/detail/6093

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 *Symantec 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.

Web installation program does not support native operating system upgrade mechanisms

The Web installation program does not support native operating system upgrade mechanisms such as Live Upgrade.

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.

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 *Symantec 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 package, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm package is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

Symantec 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

SFDB tools cannot be used with RAC standby databases. SFDB tools can still be used with the primary database, even in a Data Guard Oracle RAC environment.

Upgrading to Oracle 10.2.0.5 is required if using SFDB tools

If you are running Oracle version 10.2.0.4 and upgrading a Storage Foundation product with SFDB tools to 6.1, you must upgrade the Oracle binaries and database to version 10.2.0.5, before upgrading to 6.1.

Documentation

Product guides are available in the PDF format on the software media in the `/docs/product_name` 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

Each product in the Storage Foundation and High Availability Solutions product line includes release notes, an installation guide, and additional documents such as administration and agent guides. In most cases, you may also need to refer to the documentation for the product's components.

The SFHA Solutions documents describe functionality and solutions that apply across the product line. These documents are relevant whichever SFHA Solutions product you use.

Note: The GNOME PDF Viewer is unable to view Symantec documentation. You must use Adobe Acrobat to view the documentation.

Symantec Storage Foundation for Oracle RAC documentation

Table 1-9 lists the documentation for Symantec Storage Foundation for Oracle RAC.

Table 1-9 Symantec Storage Foundation for Oracle RAC documentation

Document title	File name	Description
Symantec Storage Foundation for Oracle RAC Release Notes	sfrac_notes_61_sol.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.

Table 1-9   Symantec Storage Foundation for Oracle RAC documentation
(continued)

Document title	File name	Description
Symantec Storage Foundation for Oracle RAC Installation and Configuration Guide	sfrac_install_61_sol.pdf	Provides information required to install and configure the product.
Symantec Storage Foundation for Oracle RAC Administrator's Guide	sfrac_admin_61_sol.pdf	Provides information required for administering and troubleshooting the product.

The SFHA Solutions documents describe functionality and solutions relevant to the SF Oracle RAC product.

See Table 1-13 on page 71.

Symantec Storage Foundation Cluster File System High Availability documentation

Table 1-10 lists the documentation for Symantec Storage Foundation Cluster File System High Availability.

The SFHA Solutions documents describe functionality and solutions relevant to the SFCFSHA product.

See Table 1-13 on page 71.

Table 1-10   Symantec Storage Foundation Cluster File System High Availability
documentation

Document title	File name	Description
Symantec Storage Foundation Cluster File System High Availability Release Notes	sfcfs_notes_61_sol.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.
Symantec Storage Foundation Cluster File System High Availability Installation Guide	sfcfs_install_61_sol.pdf	Provides information required to install the product.
Symantec Storage Foundation Cluster File System High Availability Administrator's Guide	sfcfs_admin_61_sol.pdf	Provides information required for administering the product.

Symantec Cluster Server documentation

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

Table 1-11   Symantec Cluster Server documentation

Title	File name	Description
Symantec Cluster Server Release Notes	vcs_notes_61_sol.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.
Symantec Cluster Server Installation Guide	vcs_install_61_sol.pdf	Provides information required to install the product.
Symantec Cluster Server Administrator's Guide	vcs_admin_61_sol.pdf	Provides information required for administering the product.
Symantec Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_61_sol.pdf	Provides information about bundled agents, their resources and attributes, and more related information.
Symantec Cluster Server Agent Developer's Guide (This document is available online only.)	vcs_agent_dev_61_unix.pdf	Provides information about the various Symantec agents and procedures for developing custom agents.
Symantec Cluster Server Application Note: Dynamic Reconfiguration for Oracle Servers (This document is available online only.)	vcs_dynamic_reconfig_61_sol.pdf	Provides information on how to perform dynamic reconfiguration operations on VCS clustered system domains of Oracle servers.
Symantec Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_61_sol.pdf	Provides notes for installing and configuring the DB2 agent.
Symantec Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_61_sol.pdf	Provides notes for installing and configuring the Oracle agent.
Symantec Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_61_sol.pdf	Provides notes for installing and configuring the Sybase agent.

Symantec Storage Foundation documentation

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

Table 1-12   Symantec Storage Foundation documentation

Document title	File name	Description
Symantec Storage Foundation Release Notes	sf_notes_61_sol.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.
Symantec Storage Foundation Installation Guide	sf_install_61_sol.pdf	Provides information required to install the product.
Symantec Storage Foundation Administrator's Guide	sf_admin_61_sol.pdf	Provides information required for administering the product.
Symantec Storage Foundation: Storage and Availability Management for Oracle Databases	sfhas_oracle_admin_61_unix.pdf	Provides information about the deployment and key use cases of the SFDB tools with Storage Foundation High Availability (SFHA) Solutions products in Oracle database environments. It is a supplemental guide to be used in conjunction with SFHA Solutions product guides.
Veritas File System Programmer's Reference Guide (This document is available online only.)	vxfs_ref_61_sol.pdf	Provides developers with the information necessary to use the application programming interfaces (APIs) to modify and tune various features and components of the Veritas File System.

Symantec Storage Foundation and High Availability Solutions products documentation

Table 1-13 lists the documentation for Symantec Storage Foundation and High Availability Solutions products.

Table 1-13   Symantec Storage Foundation and High Availability Solutions products documentation

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions—What's new in this release (This document is available online.)	sfhas_whats_new_61_unix.pdf	Provides information about the new features and enhancements in the release.

Table 1-13   Symantec Storage Foundation and High Availability Solutions
products documentation *(continued)*

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions Getting Started Guide	getting_started.pdf	Provides a high-level overview of installing Symantec products using the Veritas script-based installer. The guide is useful for new users and returning users that want a quick refresher.
Symantec Storage Foundation and High Availability Solutions Solutions Guide	sfhas_solutions_61_sol.pdf	Provides information about how SFHA Solutions product components and features can be used individually and in concert to improve performance, resilience and ease of management for storage and applications.
Symantec Storage Foundation and High Availability Solutions Virtualization Guide (This document is available online.)	sfhas_virtualization_61_sol.pdf	Provides information about Symantec Storage Foundation and High Availability support for virtualization technologies. Review this entire document before you install virtualization software on systems running SFHA products.
Symantec Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide (This document is available online.)	sfhas_dr_impl_61_sol.pdf	Provides information on configuring campus clusters, global clusters, and replicated data clusters (RDC) for disaster recovery failover using Storage Foundation and High Availability Solutions products.
Symantec Storage Foundation and High Availability Solutions Replication Administrator's Guide	sfhas_replication_admin_61_sol.pdf	Provides information on using Volume Replicator (VVR) for setting up an effective disaster recovery plan by maintaining a consistent copy of application data at one or more remote locations.
Symantec Storage Foundation and High Availability Solutions Troubleshooting Guide	sfhas_tshoot_61_sol.pdf	Provides information on common issues that might be encountered when using Symantec Storage Foundation and High Availability Solutions and possible solutions for those issues.

Veritas Operations Manager (VOM) is a management tool that you can use to manage Symantec Storage Foundation and High Availability Solutions products. If you use VOM, refer to the VOM product documentation at:

https://sort.symantec.com/documents

Manual pages

The manual pages for Symantec 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 Symantec 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.

The latest manual pages are available online in HTML format on the Symantec website at:

https://sort.symantec.com/documents