Symantec™ Storage Foundation for Oracle® RAC 6.2 Release Notes - Linux

January 2016

VERITAS™

Symantec™ Storage Foundation for Oracle RAC Release Notes

Product version: 6.2

Document version: 6.2 Rev 4

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_apj@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

This document provides important information about Symantec Storage Foundation for Oracle RAC (SF Oracle RAC) version 6.2 for Linux. Review this entire document before you install or upgrade SF Oracle RAC.

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

This is "Document version: 6.2 Rev 4" 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 website 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 (6.2)*.
- Storage Foundation (SF)
  See *Symantec Storage Foundation Release Notes (6.2)*.
- Storage Foundation Cluster File System High Availability (6.2)
  See *Symantec Storage Foundation Cluster File System High Availability Release Notes (6.2)*.

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 different storage configurations:
  Shared storage
  Flexible Storage Sharing (FSS): Sharing of Direct Attached Storage (DAS) and internal disks over network
- Faster performance and reduced costs per I/O per second (IOPS) using SmartIO. SmartIO supports read caching for the VxFS file systems that are mounted on VxVM volumes, in several caching modes and configurations. SmartIO also supports block-level read caching for applications running on VxVM volumes.
- Use of Cluster File System and Cluster Volume Manager for placement of Oracle Cluster Registry (OCR) and voting disks. These technologies provide robust shared block interfaces for placement of OCR and voting disks. In the absence of SF Oracle RAC, separate LUNs need to be configured for OCR and voting disks.
- Support for a standardized approach toward application and database management. Administrators can apply their expertise of 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.
- Support for container and pluggable databases in Oracle RAC 12c 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.
- 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/TECH225259
- 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/TECH225258

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 6.2

This section lists the changes in SF Oracle RAC 6.2.

Support for container and pluggable databases in Oracle 12c

SF Oracle RAC now supports the creation and configuration of container databases and pluggable databases in Oracle 12c environments. You can add the container and pluggable database resources to be managed by VCS. For steps on configuring these resources under VCS, see the *Symantec Storage Foundation for Oracle RAC Installation and Configuration Guide*.

Note: Oracle Dataguard is supported only with administrator-managed Oracle 12c databases.

SmartIO now supported in SF Oracle RAC environments

The SmartIO feature of Storage Foundation and High Availability Solutions (SFHA Solutions) enables data efficiency on your SSDs through I/O caching. Using SmartIO to improve efficiency, you can optimize the cost per IOPS. SmartIO does not require in-depth knowledge of the hardware technologies underneath. SmartIO uses advanced, customizable heuristics to determine what data to cache and how that data gets removed from the cache. The heuristics take advantage of SFHA Solutions' knowledge of the characteristics of the workload. SmartIO supports read and write caching for VxFS file systems mounted on VxVM volumes, in several caching modes and configurations.

- Read caching for applications running on VxVM volumes
- Read caching for applications running on VxFS file systems
- Database caching on VxFS file systems
- Database caching on VxVM volumes

Note: SmartIO writeback caching is not supported in SF Oracle RAC environments.

To use SmartIO, you set up a cache area on the target device. You can do this task simply with one command, while the application is online. When the application issues an I/O request, SmartIO checks to see if the I/O can be serviced from the cache. As applications access data from the underlying volumes or file systems, certain data is moved to the cache based on the internal heuristics. Subsequent I/Os are processed from the cache. You can also customize which data is cached, by adding advisory information to assist the SmartIO feature in making those determinations.

For more information, see the *Symantec. Storage Foundation and High Availability Solutions SmartIO for Solid State Drives Solutions Guide* for details.

Support for setting up ssh and rsh connection using the pwdutil.pl utility

The password utility, `pwdutil.pl`, is bundled in the 6.2 release under the scripts directory. The users can run the `pwdutil.pl` utility to set up the `ssh` and `rsh` connection automatically.

Release level terminology changes

With the 6.2 release, terms that are used to describe patch-based releases have changed as follows:

Table 1-1 Release level terminology changes

Pre 6.0.1	6.0.x, 6.1, 6.1.x	6.2 and forward	Status	Available from
P-Patch	Public hot fix	Patch	Official	SORT
Hot fix	Private hot fix	Hot fix	Unofficial	Customer support

Official patch releases are available from SORT. This release was previously referred to as a P-Patch or a Public hot fix and is now referred to as a Patch. Unofficial patch releases are available from customer support. Hot fix is the only unofficial patch release.

Package updates

The following lists the RPM changes in this release.

- The `VRTSlvmconv` RPM has been merged with the `VRTSvxvm` RPM.There is no separate RPM for `lvmconvert` now.
- The `VRTSvxvm` RPM adds dependency for bc -1.06.95-13.el7.x86_64, pcre-8.32-12.el7.i686 (pcre(x86-32)), and xz-libs-5.1.2-8alpha.el7.i686 (xz-libs(x86-32)) packages on RHEL 7 distribution. Newly required dependent RPM for `VRTSvxvm` is:

Package name	Version	Architecture
bc	Default version available with RHEL7	x86_64
pcre	Default version available with RHEL7	el7.i686
xz-libs	Default version available with RHEL7	el7.i686

For more information, see the *Installation Guide* for the complete list of RPMs for this release.

Atleast resource dependency

A new type of resource dependency has been introduced in this release wherein a parent resource can depend on a set of child resources. The parent resource is brought online or remains online only if a minimum number of child resources in this resource set are online. The system creates a set of child IP resources and the application resource depends on this set.

For example, if an application depends on five IPs and if this application has to be brought online or has to remain online, at least two IPs must be online. If two or more IP resources come online, the application attempts to come online. If the number of online resources falls below the minimum requirement, resource fault is propagated up the resource dependency tree.

For more information, refer to the *Administrator's Guide*.

Changes related to installation and upgrades

The product installer includes the following changes in SF Oracle RAC 6.2.

Connecting to the SORT website through a proxy server

The product installer connects to the Symantec Operations Readiness Tools (SORT) website for several purposes, such as downloading the latest installer patches, and

uploading installer logs. Deployment Server can connect to SORT to automatically download Maintenance or Patch release images. In this release, before running the product installer or Deployment Server, you can use the following proxy settings to connect to SORT through proxy servers:

```
# https_proxy=http://proxy_server:port
# export https_proxy
# ftp_proxy=http://proxy_server:port
# export ftp_proxy
```

Support for installation using the Red Hat Satellite server

You can install SF Oracle RAC using the Red Hat Satellite server. Red Hat Satellite is supported for Red Hat Enterprise Linux 6 (RHEL6) and Red Hat Enterprise Linux 7 (RHEL7). You can install RPMs and rolling patches on the systems which the Red Hat Satellite server manages.

In a Red Hat Satellite server, you can manage the system by creating a channel. A Red Hat Satellite channel is a collection of software packages. Using channels, you can segregate the packages by defining some rules.

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 platform (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 1-2.

Table 1-2 Deployment Server functionality

Feature	Description
Install or Upgrade systems with Install Bundle and Install Template	■ Install or upgrade systems with an Install Bundle. ■ Install packages on systems based on the information stored in the Install Template.
Define or modify Install Bundles	Define or modify Install Bundles and save them using the Deployment Server.

Table 1-2 Deployment Server functionality *(continued)*

Feature	Description
Create Install Templates	Discover installed components on a running system that you want to replicate on new systems.
Connecting the Deployment Server to SORT using a proxy server	Use a proxy server, a server that acts as an intermediary for requests from clients, for connecting the Deployment Server to the Symantec Operations Readiness Tools (SORT) website.
Platform Filtering	On the Set Preference menu, choose Selected Platforms to filter the platforms that are currently being used in the deployment environment.

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

See the *Installation Guide* for more information.

New `ProcessOnOnly` resource added to VCS configuration file during upgrades

During upgrade, the following new `ProcessOnOnly` resource is added to the VCS configuration file:

```
ProcessOnOnly vxattachd (
   Critical = 0
   Arguments = "- /usr/lib/vxvm/bin/vxattachd root"
   PathName = "$SH"
   RestartLimit = 3
   )
```

Symantec Storage Foundation for Oracle RAC gets installed in secure mode by default

Symantec Storage Foundation for Oracle RAC gets installed in secure mode by default. You are advised to install SF Oracle RAC in secure mode to be able to control guest user access to secure clusters and encrypt communication between SF Oracle RAC components. You can choose the non-secure mode during installation; however, the product installer warns you during the installation with the following message:

```
Symantec recommends that you install the cluster
in secure mode. This ensures that communication between
cluster components is encrypted and cluster information
is visible to specified users only.
```

The upgrade from non-secure mode continues to happen in non-secure mode. The upgrade from secure mode advises you to control user access to secure clusters.

Changes related to Veritas File System

There are no changes related to VxFS in this release.

VxVM SmartIO support for SF Oracle RAC installations

VxVM SmartIO is supported for SF Oracle RAC installations. When SmartIO is enabled on multiple nodes, Group Lock Manager (GLM) library keeps cache on each node coherent.

See the *Symantec™ Storage Foundation and High Availability Solutions SmartIO for Solid State Drives Solutions Guide* for details.

Changes to I/O fencing

Symantec Cluster Server (VCS) includes the following changes to I/O fencing in 6.2:

I/O fencing supports majority-based fencing mechanism, a new fencing mechanism that does not need coordination points

I/O fencing supports a new fencing mode called majority-based I/O fencing. Majority-based I/O fencing mode does not need coordination points to provide protection against data corruption and data consistency in a clustered environment. Use majority-based I/O fencing when there are no additional servers and or shared SCSI-3 disks to be used as coordination points. It provides a reliable arbitration method and does not require any additional hardware setup, such as CP Servers or shared SCSI3 disks.

In the event of a network failure, the majority sub-cluster wins the fencing race and survives the race. Note that even if the majority sub-cluster is hung or unresponsive, the minority sub-cluster loses the fencing race and the cluster panics. The cluster remains unavailable till the issue is resolved.

For more information, refer to the *Installation Guide* and *Administrator's Guide*.

Clear coordination point server registrations using the vxfenclearpre utility

The vxfenclearpre utility is enhanced to clear registrations from coordination point servers for the current cluster in addition to the existing functionality to remove SCSI3 registrations and reservation keys from the set of coordinator disks and shared data disks. The local node from where you run the utility must have the UUID of the current cluster at `/etc/vx/.uuids` directory in the `clusuuid` file.

Note that you may experience delays while clearing registrations on the coordination point servers because the utility tries to establish a network connection with IP addresses used by the coordination point servers. The delay may occur because of a network issue or if the IP address is not reachable or is incorrect.

For more information, refer to the *Administrator's Guide*.

Raw disk I/O fencing policy is not supported

Symantec does not support raw disk policy for I/O fencing. Use DMP as the I/O fencing policy for coordinator disks that have either a single hardware path or multiple hardware paths to nodes.

For more information, refer to the *Installation Guide* and *Administrator's Guide*.

Auto deport of disabled FSS disk groups

In this release, when all the storage of a shared disk group is lost and the disk group has become disabled, the CVMVoldg agent in VCS will try to deport the disk group during clean and offline entry points. However, note that if there are any pending I/Os or open counts on volumes, then the disk group will not be deported and will remain in the disabled state.

Prior to this feature, the disk group used to remain in the disabled state and you had to manually deport it.

See the SFRAC or SFCFS admin guides for more information.

SFCache agent

SFCache agent is a new agent introduced in this release. The SFCache agent enables, disables, and monitors cache. In case of a cache fault, the application still runs without any issues on the very same system, but with degraded I/O performance. Considering this, the SFCache agent provides an attribute to control the agent behavior. You can either choose to receive "IGNORE" or initiate "FAILOVER" in case of cache fault.

For more information, see *Bundled Agents Reference Guide*.

Layered volume enhancements for recovery and snapshots

In this release, a new enhancement is done for layered volumes so that when storage disconnection and subsequent reconnection happen, only inconsistent regions in the affected sub-volume are synchronized using the FastResync feature. In case of a storage failure, the mirror of the sub-volume on that storage will be detached and the future IOs on the sub-volume will be tracked by the DCO associated with the parent volume. When such a detached mirror is reattached after restoring storage connectivity, only regions that are inconsistent in the mirror would be synchronized using the FastResync feature.

Prior to this release, for a layered volume, if the storage within a mirror of a sub-volume became inaccessible, it led to full synchronization of that mirror when the storage was reconnected.

For more information about FastResync, see the *Symantec Storage Foundation for Oracle RAC Administrator's Guide*.

Changes in array names for Fusion-io devices

Prior to this release, the generic array name, fusionio, was used for all Fusion-io devices. Starting in this release, the array name indicates the type of Fusion-io card. For example, the ioDrive cards display names such as fiodrive0_0.

Use the `vxdisk list` command to display the array name.

For example:

```
# vxdisk list

fiodrive0_0  auto:cdsdisk    -            -            online ssdtrim
fiodrive0_1  auto:cdsdisk    -            -            online ssdtrim
```

Read policy enhancement

In this release, to optimize the read performance, changes have been made in the plex read policies on VxVM volumes. When there are more than one mirror available to serve the read IO, VxVM will select the set of mirrors that will provide the optimal performance and round robin between those. In selecting the set of mirrors, the internal logic will take into account various factors such as site locality, disk connectivity, media type, layout(striping), etc. You can override the logic and set any plex as the preferred mirror or set a round-robin read policy to round robin between all the mirrors of a volume.

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

The `vxattachd` daemon added as a VCS resource

In this release, the automatic site reattachment daemon, `vxattachd`, has been added in the list of resources monitored by VCS. The ProcessOnOnly agent in VCS will now monitor the `vxattachd` daemon. If the `vxattachd` process is not running, then in the next monitor cycle this agent will detect and restart it.

For more information about the `vxattachd` daemon, see the *Administrator's Guide*.

VOM integration with FSS

The Flexible Storage Sharing (FSS) feature in VxVM has been integrated with the Veritas Operations Manager (VOM) version 6.1. All the FSS operations can be done through the VOM console.

See the *Veritas™ Operations Manager Management Server 6.1 User Guide* for details.

FSS Performance improvements

In this release, the remote I/O performance on the volumes of Flexible Storage Sharing (FSS) disk groups has been improved. As there are remote devices in the FSS environment, the I/Os destined on remote disks used to give less throughput. In this release, changes have been done in the Low Latency Transport (LLT) component and Cluster Volume Manager (CVM) to remove bottlenecks at different stages. With these changes, the I/Os on remote disks give improved performance, and can reach a maximum of 90% of the local disks' performance.

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

Disk Support for FSS

In this release, in addition to the disks supported in the HCL, those disks that have the capability of generating unique IDs, have been extended support in the FSS environment. The VxVM operations such as exporting a disk, adding a disk to an FSS disk group, and creating an FSS disk group have been enhanced to internally check the compatibility of disks involved in these operations. You can use the standalone utility `vxddladm checkfss` *`disksname`* to check whether the disk is supported for FSS.

For more information about FSS disk support, see the *Administrator's Guide*.

Changes related to SFDB tools

The following sections describe the changes related to Storage Foundation for Databases (SFDB) tools in 6.2.

Support for multitenant databases

SFDB tools support operations on Oracle 12c multitenant databases. The SFDB tools do not support operations on individual Pluggable Databases (PDB).

Managing OEM using the Symantec Storage plug-in

Symantec Storage plug-in provides a graphical interface to efficiently manage and view your Storage Foundation and VCS objects through Oracle Enterprise Manager 12c (OEM).

The plug-in has the following three tabs:

- SmartIO - provides a gateway to manage the objects that use Storage Foundation's SmartIO feature, which is an advanced caching solution.
- Snapshot - enables you to apply the SFDB's point-in-time copy technologies to the selected database objects, such as datafiles, tablespaces.
- Cluster - extracts various configuration-specific information from the Symantec Cluster Server and manifests them in a tabular format.

For details on downloading and using the plug-in, visit

https://www-secure.symantec.com/connect/downloads/sfha-solutions-62-symantec-storage-plug-oem-12c

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:

- `Raw` disk policy for I/O fencing
- Oracle RAC 11g Release 1 Clusterware
- PrivNIC and MultiPrivNIC agents for Oracle RAC 11.2.0.2 and later versions
- 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.

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/TECH225259
- 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.

For basic clusters See Table 1-3 on page 24.

For campus clusters See Table 1-4 on page 25.

Table 1-3 Hardware requirements for basic clusters

Item	Description
SF Oracle RAC systems	Two to sixteen systems with two or more CPUs. **Note:** Flexible Storage Sharing (FSS) supports upto 8-node cluster configurations only. 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 8 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. If you are using the SmartIO feature, check the network requirements in the *Symantec Storage Foundation and High Availability Solutions SmartIO for Solid State Drives Solutions Guide*. Oracle requires that all nodes use the IP addresses from the same subnet.
Fiber Channel or SCSI host bus adapters	At least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

Table 1-4 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.

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

Item	Description
I/O fencing	I/O fencing requires placement of a third coordinator point at a third site. The DWDM can be extended to the third site or the iSCSI LUN at the third site can be used as the third coordination point. Alternatively Coordination Point Server can be deployed at the third remote site as an arbitration point.

Supported Linux operating systems

This section lists the supported operating systems for this release of 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	Supported level and kernel version
Red Hat Enterprise Linux 7	3.10.0-123.el7 **Note:** SF Oracle RAC has not yet announced support for RHEL 7. You may find information pertaining to RHEL 7 in the installation and administrator guides. Note that this information will become relevant only after SF Oracle RAC announces support when due certification efforts are complete. Refer to the following TechNote for the latest information on the supported operating systems and Oracle RAC database versions. http://www.symantec.com/docs/DOC4848.
Red Hat Enterprise Linux 6	Update 3 (2.6.32-279.el6) Update 4 (2.6.32-358.el6) Update 5 (2.6.32-431.el6) Update 6 (2.6.32-504.el6)
SUSE Linux Enterprise 11	SP2 (3.0.13-0.27.1) SP3 (3.0.76-0.11.1)

Table 1-5 Supported operating systems *(continued)*

Operating systems	Supported level and kernel version
Oracle Linux 6 (RHEL compatible mode)	Update 3 (2.6.32-279.el6) Update 4 (2.6.32-358.el6) Update 5 (2.6.32-431.el6) Update 6 (2.6.32-504.el6)
Oracle Linux 7 (RHEL compatible mode)	3.10.0-123.el7 **Note:** SF Oracle RAC has not yet announced support for Oracle Linux 7. You may find information pertaining to OL 7 in the installation and administrator guides. Note that this information will become relevant only after SF Oracle RAC announces support when due certification efforts are complete. Refer to the following TechNote for the latest information on the supported operating systems and Oracle RAC database versions. http://www.symantec.com/docs/DOC4848

Note: All subsequent kernel updates are supported, but you should check the Symantec Operations Readiness Tools (SORT) website for additional information that applies to the exact kernel version for which you plan to deploy.

Note: Only 64-bit operating systems are supported on the AMD Opteron or the Intel Xeon EM64T (x86_64) Processor line.

Note: Configuring LLT over RDMA is not supported with Oracle Linux Unbreakable Enterprise Kernel 2 that is 2.6.39-400.17.1.el6uek.x86_64.

Note: For the latest information on supported hardware, visit the following hardware compatibility list URL:

http://www.symantec.com/docs/TECH211575

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

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

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

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

Required Linux RPMs for SF Oracle RAC

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

Note: Some required RHEL RPMs have different version numbers between RHEL update versions.

Table 1-6 lists the RPMs that SF Oracle RAC requires for a given Linux operating system.

Table 1-6 Required RPMs

Operating system	Required RPMs
RHEL 7 **Note:** Symantec recommends that you install RHEL 7 as the operating system of Server GUI.	bc -1.06.95-13.el7.x86_64 coreutils-8.22-11.el7.x86_64 ed-1.9-4.el7.x86_64 findutils-4.5.11-3.el7.x86_64 gcc-c++-4.8.2-16.el7.x86_64 gcc-4.8.2-16.el7.x86_64 glibc-2.17-55.el7.i686 glibc-2.17-55.el7.x86_64 glibc-headers-2.17-55.el7.x86_64 glib-networking-2.36.2-3.el7.x86_64 glibmm24-2.36.2-4.el7.x86_64 glibc-common-2.17-55.el7.x86_64 glibc-devel-2.17-55.el7.x86_64 glibc-devel-2.17-55.el7.i686 glib2-2.36.3-5.el7.x86_64 glibc-utils-2.17-55.el7.x86_64 kmod-14-9.el7.x86_64 ksh-20120801-19.el7.x86_64

Table 1-6 Required RPMs *(continued)*

Operating system	Required RPMs
RHEL 7 (continued)	libacl-2.2.51-12.el7.i686
	libacl-2.2.51-12.el7.x86_64
	libaio-devel-0.3.109-12.el7.x86_64
	libaio-devel-0.3.109-12.el7.i686
	libaio-0.3.109-12.el7.i686
	libaio-0.3.109-12.el7.x86_64
	libgcc-4.8.2-16.el7.i686
	libgcc-4.8.2-16.el7.x86_64
	libstdc++-4.8.2-16.el7.i686
	libstdc++-4.8.2-16.el7.x86_64
	lsof-4.87-4.el7.x86_64
	ncompress-4.2.4.4-3.el7.x86_64
	ncurses-libs-5.9-13.20130511.el7.x86_64
	nss-softokn-freebl-3.15.4-2.el7.i686
	pam-1.1.8-9.el7.i686
	parted-3.1-17.el7.x86_64
	pcre-8.32-12.el7.i686 (pcre(x86-32))
	policycoreutils-2.2.5-11.el7.x86_64
	prelink-0.5.0-6.el7.x86_64
	screen-4.1.0-0.19.20120314git3c2946.el7.x86_64
	systemd-libs-208-11.el7.i686
	systemd-libs-208-11.el7.x86_64
	xz-libs-5.1.2-8alpha.el7.i686 (xz-libs(x86-32))

Table 1-6 Required RPMs *(continued)*

Operating system	Required RPMs
OL 6	coreutils-8.4-19.el6.x86_64.rpm
	ed-1.1-3.3.el6.x86_64.rpm
	findutils-4.4.2-6.el6.x86_64.rpm
	glibc-2.12-1.80.el6.i686.rpm
	glibc-2.12-1.80.el6.x86_64.rpm
	ksh-20100621-16.el6.x86_64.rpm
	libacl-2.2.49-6.el6.x86_64.rpm
	libgcc-4.4.6-4.el6.i686.rpm
	libgcc-4.4.6-4.el6.x86_64.rpm
	libstdc++-4.4.6-4.el6.i686.rpm
	libstdc++-4.4.6-4.el6.x86_64.rpm
	mksh-39-7.el6.x86_64.rpm
	module-init-tools-3.9-20.0.1.el6.x86_64.rpm
	ncurses-libs-5.7-3.20090208.el6.x86_64.rpm
	nss-softokn-freebl-3.12.9-11.el6.i686.rpm
	openssl-1.0.0-20.el6_2.5.x86_64.rpm
	pam-1.1.1-10.el6_2.1.i686.rpm
	parted-2.1-18.el6.x86_64.rpm
	perl-5.10.1-127.el6.x86_64.rpm
	policycoreutils-2.0.83-19.24.0.1.el6.x86_64.rpm
	readline-6.0-4.el6.x86_64.rpm

Table 1-6 Required RPMs *(continued)*

Operating system	Required RPMs
RHEL 6	coreutils-8.4-19.el6.x86_64.rpm
	ed-1.1-3.3.el6.x86_64.rpm
	findutils-4.4.2-6.el6.x86_64.rpm
	glibc-2.12-1.80.el6.i686.rpm
	glibc-2.12-1.80.el6.x86_64.rpm
	ksh-20100621-16.el6.x86_64.rpm
	libacl-2.2.49-6.el6.x86_64.rpm
	libgcc-4.4.6-4.el6.i686.rpm
	libgcc-4.4.6-4.el6.x86_64.rpm
	libstdc++-4.4.6-4.el6.i686.rpm
	libstdc++-4.4.6-4.el6.x86_64.rpm
	mksh-39-7.el6.x86_64.rpm
	module-init-tools-3.9-20.el6.x86_64.rpm
	ncurses-libs-5.7-3.20090208.el6.x86_64.rpm
	nss-softokn-freebl-3.12.9-11.el6.i686.rpm
	openssl-1.0.0-20.el6_2.5.x86_64.rpm
	pam-1.1.1-10.el6_2.1.i686.rpm
	parted-2.1-18.el6.x86_64.rpm
	policycoreutils-2.0.83-19.24.el6.x86_64.rpm
	readline-6.0-4.el6.x86_64.rpm
	zlib-1.2.3-27.el6.x86_64.rpm

Table 1-6 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 11 SP2	coreutils-8.12-6.19.1.x86_64.rpm
	ed-0.2-1001.30.1.x86_64.rpm
	findutils-4.4.0-38.26.1.x86_64.rpm
	glibc-2.11.3-17.31.1.x86_64.rpm
	glibc-32bit-2.11.3-17.31.1.x86_64.rpm
	ksh-93u-0.6.1.x86_64.rpm
	libacl-2.2.47-30.34.29.x86_64.rpm
	libacl-32bit-2.2.47-30.34.29.x86_64.rpm
	libgcc46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm
	libgcc46-4.6.1_20110701-0.13.9.x86_64.rpm
	libncurses5-5.6-90.55.x86_64.rpm
	libstdc++46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm
	libstdc++46-4.6.1_20110701-0.13.9.x86_64.rpm
	module-init-tools-3.11.1-1.21.1.x86_64.rpm
	pam-32bit-1.1.5-0.10.17.x86_64.rpm
	parted-2.3-10.21.18.x86_64.rpm
	zlib-1.2.3-106.34.x86_64.rpm
	zlib-32bit-1.2.3-106.34.x86_64.rpm

Table 1-6 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 11 SP3	coreutils-8.12-6.25.27.1.x86_64.rpm
	ed-0.2-1001.30.1.x86_64.rpm
	findutils-4.4.0-38.26.1.x86_64.rpm
	glibc-2.11.3-17.54.1.x86_64.rpm
	glibc-32bit-2.11.3-17.54.1.x86_64.rpm
	ksh-93u-0.18.1.x86_64.rpm
	libacl-2.2.47-30.34.29.x86_64.rpm
	libacl-32bit-2.2.47-30.34.29.x86_64.rpm
	libgcc_s1-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libgcc_s1-4.7.2_20130108-0.15.45.x86_64.rpm
	libncurses5-5.6-90.55.x86_64.rpm
	libstdc++6-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libstdc++6-4.7.2_20130108-0.15.45.x86_64.rpm
	module-init-tools-3.11.1-1.28.5.x86_64.rpm
	pam-32bit-1.1.5-0.10.17.x86_64.rpm
	parted-2.3-10.38.16.x86_64.rpm
	zlib-1.2.7-0.10.128.x86_64.rpm
	zlib-32bit-1.2.7-0.10.128.x86_64.rpm

Additional RPMs required for Veritas Volume Manager

You must install the 32-bit `libudev` RPM before you install Veritas Volume Manager.

Table 1-7 lists the required RPMs.

Table 1-7 Additional RPMs required for Veritas Volume Manager

Operating system	Required RPMs
RHEL 6 Update 5	libudev-147-2.51.el6.i686.rpm
RHEL 6 Update 4	libudev-147-2.46.el6.i686.rpm
RHEL 6 Update 3	libudev-147-2.41.el6.i686.rpm
SLES 11 SP3	libudev0-32bit-147-0.84.1.x86_64.rpm

Table 1-7 Additional RPMs required for Veritas Volume Manager *(continued)*

Operating system	Required RPMs
SLES 11 SP2	libudev0-32bit-147-0.47.2.x86_64.rpm

Supported database software

For information on supported database versions, see the Software Compatibility List (SCL):

http://www.symantec.com/docs/TECH213121

For information on certified Oracle database versions, see the following technical note:

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

Support for minor database versions is also documented in the technical note.

Note: It is possible that an Oracle version listed in the SCL is not present in the support matrix. This is because the support matrix is updated for Oracle versions only after completing required Oracle certifications. The certification process usually take a few months to complete after the product release.

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

https://support.oracle.com

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.

Installation and upgrades fixed issues

This section describes the incidents that are fixed related to installation and upgrades in this release.

Table 1-8 Fixed issues related to installation and upgrades

Incident	Description
3326196	Rolling upgrade may encounter a problem if open volumes from different disk groups have the same name.
3326639	CP server service group fails to come online with the default database path after the CP server is upgraded from 6.0 to 6.2 on a multi-node cluster.
3442070	If you select rolling upgrade task from the Install Bundles menu, the Installer exits with an error.

Issues fixed in SF Oracle RAC 6.2

Table 1-9 lists the issues fixed in SF Oracle RAC 6.2.

Table 1-9 Issues fixed in SF Oracle RAC 6.2

Incident	Description
3321004	Installation of Oracle Clusterware using Oracle response file fails.
3348812	During HAIP configuration, the SF Oracle RAC web installer fails to update the /etc/hosts file with the HAIP alias.
3442070	If you select rolling upgrade task from the Install Bundles menu, the common product installer exits with an error.

Cluster Volume Manager fixed issues

This section describes the incidents that are fixed in Cluster Volume Manager (CVM) in this release.

Table 1-10 Cluster Volume Manager fixed issues

Incident	Description
640213	The node panics in case of overlapping reconfigurations due to race conditions.
3592783	For the FSS partially shared storage configuration, after you run `hastop -all` and then `hastart`, remote disks, which are not part of the disk group, are not visible in `vxdisk -o alldgs list`.
3582470	Data change object (DCO) volume gets created using the same node's disks for both plexes.
3573908	Multiple `cbrbk.tmp$$` files in the /var/tmp folder on each node do not clean up properly.
3552008	The `vxconfigrestore` (vxvol resync) operation hangs on the master node while recovering a stripe-mirror volume.
3551050	The `vx*` commands are not able to connect to vxconfigd.
3538683	After a master panic, the new master does not start plex resync when the older master comes online and joins the cluster.
3537519	The `vxdisk unexport` command hangs and then `vxconfigd` gets fault in an asymmetric array connection setup.
3523731	VxVM command check for device compliance prior to doing FSS operations on disks.
3508390	DCO object is unnecessarily marked in the BADLOG mode in cascade failure scenarios, and it results in requiring a full-recovery and can result in lost snapshots as well.
3505017	When the `da name` is the same as the existing `dm name`, the `vxdg addisk` operation from slave fails.
3496673	On a Flexible Storage Sharing (FSS) disk group, the read I/O performance is impacted.
3495811	When you create a disk group, the disk shows LMISSING state when the SCSI PGR operation fails.
3489167	Plex(es) on remote disks goes to DISABLED state because of a plex I/O error encountered after slave node reboot in Cluster Volume Manger (CVM).

Table 1-10 Cluster Volume Manager fixed issues *(continued)*

Incident	Description
3484570	Accessing CVM messages after decrementing the reference count causes a panic.
3482026	The `vxattachd`(1M) daemon reattaches plexes of manually detached site.
3430256	Space allocation for Volume: Single DAS disk in a disk group takes more preference than shared storage in that disk group.
3422704	Unique prefix generation algorithm redesign for forming cluster wide consistent name (CCN).
3411438	The `preferred` and `round-robin` read policy settings should be honoured irrespective of local connectivity to the plexes.
3394940	Anomaly numbers are displayed in the `vxstat` output.
3383625	When a cluster node that contributes the storage to the Flexible Storage Sharing (FSS) disk group rejoins the cluster, the local disks brought back by that node do not get reattached.
3373747	Adding new nodes to the 22-node cluster causes Cluster File System (CFS) failures after CVM deletes 2 nodes.
3368361	When siteconsistency is configured within a private disk group and Cluster Volumne Manager (CVM) is up, then the reattach operation of a detached site fails.
3329603	The `vxconfigd` related error messages are observed in system log files on every node for large cluster setups.
3300418	VxVM volume operations on shared volumes cause unnecessary read I/Os.
3286030	Vxattachd debug messages get displayed on the console during a reboot.
3283418	Writes from the source node hang due to heavy workload on the target node.
3281160	When autoreminor is set off, no error is thrown when you import a disk group having the same minor number as that of the existing imported disk group.

Table 1-10 Cluster Volume Manager fixed issues *(continued)*

Incident	Description
3152304	When connectivity to some of the plexes of a volume is lost from all nodes, an I/O hang occurs.
2705055	The `preferred` and `round-robin` read policy settings should be honoured irrespective of local connectivity to the plexes.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues

Table 1-11 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-11 SFDB tools fixed issues

Incident	Description
2869266	Checkpoint clone fails if the archive log destination is same as the datafiles destination.
3313775	SmartIO options are not restored after Reverse Resync Commit operation is performed.
3615735	During a Reverse Resync Begin operation, a mismatch in database control file version is observed.
3615745	For thin storage setups, the snapshot operation reports that the diskgroup cannot be split.
3615764	The flashSnap operation fails to create a symlink on a Symantec Volume Replicator (VVR) secondary site.

LLT, GAB, and I/O fencing fixed issues

Table 1-12 lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-12 LLT, GAB, and I/O fencing fixed issues

Incident	Description
3335137	Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups.

Table 1-12 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
3473104	When virtual NICs are configured under LLT without specifying the MTU size 1500 in lltab, cluster does not function properly. For example, VCS engine commands may hang and print below message in the engine logs: `VCS CRITICAL V-16-1-51135 GlobalCounter not updated`
3031216	The dash (-) in a disk group name causes vxfentsthdw(1M) and Vxfenswap(1M) utilities to fail.
3471571	Cluster nodes may panic if you stop the HAD process by force on a node and reboot that node.
3410309	LLT driver fails to load and logs the following message in the syslog when a mismatch is observed in the RDMA-specific symbols. llt: disagrees about version of symbol rdma_connect llt: Unknown symbol rdma_connect llt: disagrees about version of symbol rdma_destroy_id llt: Unknown symbol rdma_destroy_id
3532859	The Coordpoint agent monitor fails if the cluster has a large number of coordination points.

Known issues

This section covers the known issues in this release.

For Oracle RAC issues:

See “Oracle RAC issues” on page 40.

For SF Oracle RAC issues:

See “SF Oracle RAC issues” on page 41.

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

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.

`installer -requirements` does not list RHEL 6 Update 6 and Oracle Linux 7 as supported platforms (3657260)

The `installer -requirements` command does not list RHEL 6 Update 6 and Oracle Linux 7 as supported platforms though they are qualified with version 6.2.

Workaround: The correct supported list is mentioned in the latest version of the product Release Notes. See the latest Release Notes on the Symantec website for the updated list.

https://sort.symantec.com/documents

Installer reports incorrect minimal version for several required Oracle Linux 7 RPMs (3653382)

During installation, the product installer reports incorrect minimum version for the following required Oracle Linux 7 RPMs:

```
systemd-libs-208-11.el7.i686
coreutils-8.22-11.el7.x86_64
policycoreutils-2.2.5-11.el7.x86_64
```

The correct minimum version required for the RPMs is as follows:

```
systemd-libs-208-11.0.1.el7.i686.rpm
coreutils-8.22-11.0.1.el7.x86_64
policycoreutils-2.2.5-11.0.1.el7.x86_64
```

Workaround: Install the required operating system RPMs using native methods, such as yum, or install them manually.

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
```

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 the first phase of a rolling upgrade, make sure the CVM is online on all nodes without errors (2595441)

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

The uninstaller does not remove all scripts (2696033)

After removing SF Oracle RAC, some of the RC scripts remain in the `/etc/rc*.d/` folder. This is due to an issue with the chkconfig rpm in RHEL6 and updates. You can manually remove the scripts from the `/etc/rc*.d/` folder after removing the VxVM RPMs.

Workaround: Install the chkconfig-1.3.49.3-1 chkconfig rpm from the RedHat portal. Refer to the following links:

http://grokbase.com/t/centos/centos/117pfhe4zz/centos-6-0-chkconfig-strange-behavior

http://rhn.redhat.com/errata/RHBA-2012-0415.html

LLT known issues

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

Dependency may get overruled when uninstalling multiple RPMs in a single command [3563254]

When performing uninstallation of multiple RPMs through a single comment, the system identifies and follows the specified dependency among the RPMs as the uninstallation progresses. However, if the pre-uninstallation script fails for any of

the RPMs, the system does not abort the task but instead uninstalls the remaining RPMs.

For example, if you run `rpm -e VRTSllt VRTSgab VRTSvxfen` where the RPMs have a dependency between each other, the system bypasses the dependency if the pre-uninstallation script fails for any RPM.

Workaround: Uninstall the RPMs independently.

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

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

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

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

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

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

LLT port stats sometimes shows recvcnt larger than recvbytes (1907228)

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

LLT may incorrectly declare port-level connection for nodes in large cluster configurations [1810217]

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

If you manually re-plumb (change) the IP address on a network interface card (NIC) which is used by LLT, then LLT may experience heartbeat loss and the node may panic (3188950)

With the LLT interfaces up, if you manually re-plumb the IP address on the NIC, then the LLT link goes down and LLT may experience heartbeat loss. This situation may cause the node to panic.

Workaround: Do not re-plumb the IP address on the NIC that is currently used for LLT operations. Take down the stack before you re-plumb the IP address for the LLT interface.

A network restart of the network interfaces may cause heartbeat loss for the NIC interfaces used by LLT

A network restart may cause heartbeat loss of the network interfaces configured LLT. LLT configured for UDP or LLT configured for RDMA may experience loss of heartbeat between the interfaces, which may cause the node to panic.

Workaround: Recommendations before you restart the network:

- Assess the effect of a network restart on a running cluster that is using LLT over RDMA or LLT over UDP.
- Do not use the network restart functionality to add or configure a new NIC to the system.
- If you are using the network restart functionality, make sure that the LLT interfaces are not affected.
- Increase the `llt-peerinact` time to a higher value to allow network restart to complete within that time.

  Run the `# lltconfig -T peerinact:6000` command to increase the peerinact time to 1 minute.

When you execute the /etc/init.d/llt start script to load the LLT module, the syslog file may record messages related to kernel symbols associated with Infiniband (3136418)

When you execute /etc/init.d/llt start to start the LLT module on some Linux kernel versions, the syslog file may display the following messages for multiple such symbols:

```
kernel: llt: disagrees about version of symbol ib_create_cq
kernel: llt: Unknown symbol ib_create_cq
```

The LLT module is shipped with multiple module *.ko files, which are built against different kernel versions. If the kernel version on the node does not match the kernel version against which the LLT module is built, the LLT module fails to load and logs

RDMA-related messages in the syslog file. In this case, the kernel logs these messages. The modinst script loads the compatible module on the system and starts LLT without any issues.

Workaround: Rearrange the kernel versions in the `/opt/VRTSllt/kvers.lst` file such that the first line displays the kernel version that is most likely to be compatible with the kernel version on the node. This rearrangement allows the modinst script to load the best possible kernel module first. Therefore, the warning message is less likely to appear.

GAB known issues

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

While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373)

After you unregister the gtx port with `-nodeinit` option, the `gabconfig -C` command shows refcount as 1. But when forceful `deinit` option (`gabdebug -R GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinited on user request
```

The `refcount` value is incremented by 1 internally. However, the refcount value is shown as 2 which conflicts with the `gabconfig -C` command output.

Workaround: There is no workaround for this issue.

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.

I/O fencing known issues

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

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.

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.

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

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

```
# cpsadm -s 10.209.125.200 -a ping_cps
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using
environment variable
CPS_NODEID
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

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

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

the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

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

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.

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).

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 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 RPM is not removed from the system, the `cpsadm` command loads the old security libraries present on the system. As the installer runs the `cpsadm` command on the CP server to add or upgrade the SF Oracle RAC cluster (application cluster), the installer also fails.

Workaround: Perform the following procedure on all of the nodes of the CP server.

To resolve this issue

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.

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.

CoordPoint agent does not report the addition of new disks to a Coordinator disk group [2727672]

The LevelTwo monitoring of the CoordPoint agent does not report a fault even if the constituent of a coordinator disk group changes due to addition of new disks in the cooridnator disk group

Workaround: There is no workaround for this issue.

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.

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 `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 the VRTSvxfen RPM, 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.

VCS fails to take virtual machines offline while restarting a physical host in RHEV and KVM environments (3320988)

In RHEV and KVM environments, the virtualization daemons `vdsmd` and `libvirtd`required to operate virtual machines are stopped before VCS is stopped during a reboot of the physical host. In this scenario, VCS cannot take the virtual machine resource offline and therefore the resource fails to stop. As a result , LLT, GAB and fencing fail to stop. However, the virtual network bridge is removed leading to the loss of cluster interconnects and causing a split-brain situation.

Workaround: If the virtual network bridge is not assigned to any virtual machine, remove the virtual bridge and configure LLT to use the physical interface. Alternatively, before initiating a reboot of the physical host, stop VCS by issuing the `hastop -local` command. The `-evacuate` option can be used to evacuate the virtual machines to another physical host.

Fencing may panic the node while shut down or restart when LLT network interfaces are under Network Manager control [3627749]

When the LLT network interfaces are under Network Manager control, then shutting down or restarting a node may cause fencing race resulting in a panic. On RHEL, VCS requires that LLT network interfaces are not put under Network Manager control, as it might cause problems when a node is shut down or restarted. During shutdown, the Network Manager service might stop before the VCS shutdown scripts are called. As a result, fencing race is triggered and the losing sub-cluster panics.

Workaround: Either exclude the network interfaces to be used by LLT from Network Manager control or disable the Network Manager service before configuring LLT. Please refer to the Red Hat documentation to do the same.

ASM disk groups configured with normal or high redundancy are dismounted if the CVM master panics due to network failure in FSS environment or if CVM I/O shipping is enabled (3600155)

Disk-level remote write operations are paused during reconfiguration for longer than the default ASM heartbeat I/O wait time in the following scenarios:

- CVM master node panics
- Private network failure

As a result, the ASM disk groups get dismounted.

Workaround: See to the Oracle metalink document: 1581684.1

ODM linking changes in Oracle RAC 12.1.0.2 [3719159]

Oracle has changed the location of ODM library used for loading during database startup in version 12.1.0.2. As a result, the Veritas ODM library is not used by the Oracle database even after linking the Oracle database with Veritas ODM.

Note: The product installer will display the following warning message during the post-check operation. You can ignore the message.

```
CPI WARNING V-9-40-4999 Libraries are not linked properly on sys1.
```

Perform the following manual steps to link or unlink Veritas ODM.

To link the Veritas ODM with Oracle RAC 12.1.0.2 binaries

1. Log in as the Oracle user.
2. If the Oracle database is running, then shut down the Oracle database. Run the following on one node:

   ```
   $ srvctl stop database -db db_name
   ```

3. Perform the following step before you link the Veritas ODM library with the Oracle binaries:
   - Export the ORACLE_HOME environment variable, where Oracle database binaries are installed.
   - For Oracle RAC database, disable direct NFS:

     ```
     $ cd $ORACLE_HOME/rdbms/lib
     $ make -f ins_rdbms.mk dnfs_off
     ```

     where ORACLE_HOME is the location where Oracle database binaries are installed.
   - Change the present working directory:

     ```
     $ cd $ORACLE_HOME/rdbms/lib/odm
     ```

   - Link the Veritas ODM library with the Oracle binaries:

     ```
     $ ln -s /opt/VRTSodm/lib64/libodm.so libodm12.so
     ```

   - Start the database. Run the following on one node:

     ```
     $ srvctl start database -db db_name
     ```

4. To confirm that the Oracle database starts with Veritas Extension for ODM, check the alert log for the following text:

   ```
   Veritas <version> ODM Library
   ```

   The alert log location depends on the Oracle version in use. For more information on the exact location of the alert log, see the Oracle documentation.

To unlink the Veritas ODM library

1. Log in as the Oracle user.
2. Export the ORACLE_HOME environment variable, where Oracle database binaries are installed.

3 If the Oracle database is running, then shut down the Oracle database. Run the following on one node:

```
$ srvctl stop database -db db_name
```

4 Unlink Veritas ODM:

```
$ cd $ORACLE_HOME/lib
$ ln -sf $ORACLE_HOME/lib/libodmd12.so libodm12.so
$ cd $ORACLE_HOME/rdbms/lib/odm
$ rm libodm12.so
$ cd $ORACLE_HOME/rdbms/lib
$ make -f ins_rdbms.mk dnfs_on
```

5 Start the database. Run the following on one node:

```
$ srvctl start database -db db_name
```

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
```

A volume remains in DETACHED state even after storage nodes join back to the cluster (3628933)

This issue occurs in an FSS configuration, in the following scenario:

1. The volume is part of an FSS disk group with storage from only a subset of the nodes.
2. The storage nodes fail or are rebooted while I/O is in progress on all the nodes.
3. The node in the cluster that is not contributing to storage becomes the master.
4. The storage nodes come up and join the cluster.

The issue is that the volume remains in a detached state even after the storage nodes rejoin the cluster. Trying to start the volume manually with the following command generates an error:

```
# vxvol start -g dg_name vol_name
VxVM vxvol ERROR V-5-1-10128  DCO experienced
IO errors during the operation.
Re-run the operation after ensuring that DCO is accessible
```

Workaround:

Deport the disk group and then import the disk group.

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.

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

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

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

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

The vxconfigd daemon fails to start after machine reboot (3566713)

The `shutdown -r` command makes sure that the file contents on the OS file system are written properly to the disk before a reboot. The `volboot` file is created in the OS file system, and is used to bring up the `vxconfigd` daemon after the system reboot. If the machine reboots for any reason without proper shutdown, and the `volboot` file contents are not flushed to the disk, `vxconfigd` will not start after the system reboots.

Workaround:

You must rerun the `vxinstall` script to re-create the `volboot` file and to start the `vxconfigd` daemon and other daemons.

Restarting the vxconfigd daemon on the slave node after a disk is removed from all nodes may cause the disk groups to be disabled on the slave node (3591019)

The issue occurs if the storage connectivity of a disk is removed from all the nodes of the cluster and the `vxconfigd` daemon is restarted on the slave node before the disk is detached from the slave. All the disk groups are in the `dgdisabled` state on the slave node, but show as `enabled` on the other nodes.

If the disk was detached before the `vxconfigd` daemon is restarted, the issue does not occur.

In a Flexible Storage Sharing (FSS) environment, removing the storage connectivity on a node that contributes DAS storage to a shared disk group results in global connectivity loss because the storage is not connected elsewhere.

Workaround:

To prevent this issue:

Before restarting the `vxconfigd` daemon, if a disk in a shared disk group has lost connectivity to all nodes in the cluster, make sure that the disk is in the `detached` state. If a disk needs to be detached, use the following command:

```
# vxdisk check diskname
```

To resolve the issue after it has occurred:

If `vxconfigd` is restarted before the disks got detached, remove the node from the cluster and rejoin the node to the cluster.

Health check monitoring fails with policy-managed databases (3609349)

The health check option of the Symantec Cluster Server agent for Oracle fails to determine the status of the Oracle resource in policy-managed database environments. This is because the database SID is dynamically created during the time of the health check as a result of which the correct SID is not available to retrieve the resource status.

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 when CVMDeportOnOffline is set to 1

When CVMDeportOnOffline is set to 1, 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 and the disk group is required to be deported during offline, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

Rolling upgrade not supported for upgrades from SF Oracle RAC 5.1 SP1 with fencing configured in `dmp`mode.

Rolling upgrade is not supported if you are upgrading from SF Oracle RAC 5.1 SP1 with fencing configured in `dmp`mode. This is because fencing fails to start after the

system reboots during an operating system upgrade prior to upgrading SF Oracle RAC.

The following message is displayed:

```
VxVM V-0-0-0 Received message has a different protocol version
```

Workaround: Perform a full upgrade if you are upgrading from SF Oracle RAC 5.1 SP1 with fencing configured in `dmp`mode.

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

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

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

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

Workaround: There is no workaround for this issue.

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.

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.

vxdisk resize from slave nodes fails with "Command is not supported for command shipping" error (3140314)

When running the `vxdisk resize` command from a slave node for a local disk, the command may fail with the following error message:

```
VxVM vxdisk ERROR V-5-1-15861 Command is not supported for command
shipping.
Operation must be executed on master
```

Workaround: Switch the master to the node to which the disk is locally connected and run the `vxdisk resize` on that node.

vxconfigbackup fails on Flexible Storage Sharing disk groups (3079819)

The `vxconfigbackup` command fails on disk groups with remote disks that have the Flexible Storage Sharing attribute set with the following error messages:

```
VxVM vxconfigbackup ERROR V-5-2-3719 Unable to get the disk serial number.
VxVM vxconfigbackup ERROR V-5-2-6144 configbackup cannot proceed without
uuid.
    FAILED:EXEC /usr/lib/vxvm/bin/vxconfigbackup
```

Workaround: There is no workaround for this issue.

CVR configurations are not supported for Flexible Storage Sharing (3155726)

Cluster Volume Replicator (CVR) configurations are not supported in a Flexible Storage Sharing environment.

Default volume layout with DAS disks spans across different disks for data plexes and DCO plexes (3087867)

The default volume layout for Flexible Storage Sharing disk groups is a two-way mirrored volume with a DCO log. The DCO log plexes may be allocated using host class instances that may be different from the ones used for the data plexes.

Workaround: Use the command line `alloc` attributes to explicitly set allocation requirements. For example, you can specify `alloc=host:host1,host:host2` during volume creation on an FSS disk group, and allocate DCO plexes on the same host (failure domain) as the data plexes.

SG_IO ioctl hang causes disk group creation, CVM node joins, and storage connects/disconnects, and vxconfigd to hang in the kernel (3193119)

In RHEL 5.x, the `SG_IO ioctl` process hangs in the kernel. This causes disk group creation and CVM node joins to hang. The `vxconfigd` thread hangs in the kernel during storage connects/disconnects and is unresponsive.

Workaround: This issue is fixed in RHEL 6.3. Upgrade to RHEL 6.3.

Disk group remains in deported state for remote disks after the destroying the disk group from a node that is not exporting the disks (3117153)

When a disk group is destroyed, the headers on some of the disks that are not locally connected are not cleared, and have stale configurations. The disk configuration displays that the disk belongs to the deported disk group.

Workaround: Reinitialize the disks that have stale configurations:

```
# vxdisk -f init diskname format=cdsdisk
```

Preserving Flexible Storage Sharing attributes with vxassist grow and vxresize commands is not supported (3225318)

Preservation of FSS attributes using `vxasssist grow` and `vxresize` is not supported. FSS attributes include the attributes that are specified on the command line as well as the attributes implicitly assumed for FSS disk groups. These attributes are not reused with further `vxassist` orpations on the volume such as the `growby` and the `vxresize` commands.

Workaround:

There is no workaround for this issue.

vxdg adddisk operation fails when adding nodes containing disks with the same name (3301085)

On a slave node, when using the `vxdg adddisk` command to add a disk to a disk group, and if the device name already exists in the disk group as disk name (disk media name), the operation fails with the following message:

```
VxVM vxdg ERROR V-5-1-599 Disk disk_1: Name is already used.
```

Workaround: Explicitly specify the disk media name, which is different from the existing disk media name in the disk group, when running the `vxdg adddisk` command on the slave node.

For example:

```
# vxdg -g diskgroup adddisk dm1=diskname1 dm2=diskname2 dm3=diskname3
```

Flexible Storage Sharing export operation fails when nodes in the cluster are joined in parallel (3327028)

When two or more nodes join the cluster in parallel in an FSS environment, the remote disk creation on some nodes may fail with the following message in the syslog:

```
vxvm:vxconfigd: V-5-1-12143 CVM_VOLD_JOINOVER command received for node(s) 1
vxvm:vxconfigd: V-5-1-3866 node 1: vxconfigd not ready
vxvm:vxconfigd: V-5-1-3866 node 1: vxconfigd not ready
vxvm:vxconfigd: V-5-1-18321 Export operation failed : Slave not joined
...
vxvm:vxconfigd: V-5-1-4123 cluster established successfully
```

The automatic reattach of subdisks and plexes may not occur, causing some resources to remain in the offline or faulted state. User intervention is required to remove the fault and bring the resources online.

Workaround:

Manually reattach the disks from the node that has connectivity to the disks:

```
# vxreattach diskname
```

If the resources are faulted, clear the fault and online the service group:

```
# hagrp -clear service_group
```

```
# hagrp -online service_group -any
```

In a Flexible Storage Sharing disk group, the default volume layout is not mirror when creating a volume with the mediatype:ssd attribute (3209064)

As per current VxVM behavior, specifying a subset of disks, such as mediatype:ssd, as a command line argument during volume creation, takes precedence over internal FSS attributes. VxVM does not implicitly apply by default the mirrored volume layout for a FSS volume.

Workaround: Explicitly specify the `layout=mirror` attribute during volume creation.

```
# vxassist -g diskgroup make volume size mediatype:ssd layout=mirror
```

Remote writes hang due to heavy sync workload on the target node in FSS environments (3283418)

With the default Completely Fair Queuing (CFQ) I/O scheduler in Linux, local reads/writes which are sync workload are given priority over async remote writes. Whereas remote reads happen at par with local I/Os. The reads are always considered as sync type, but remote writes are only considered as async. As a result, sync I/Os are given preference over the async writes and are dispatched after a long time.

Workaround:

The VxVM recommended scheduler is Deadline. With the Deadline scheduler, I/Os work fine. This issue doesn't occur in non-rotating disk media like solid-state disk (SSD) devices as they don't have any I/O scheduler.

For more information, see the TechNote for Symantec recommended Linux scheduler:

http://www.symantec.com/business/support/index?page=content&id=TECH181220

This TechNote is for Redhat releases, but it also applies to SLES 11 OS. Please refer to your OS vendor's administrator guidelines and best case practices before proceeding with this change.

FSS Disk group creation with 510 exported disks from master fails with Transaction locks timed out error (3311250)

Flexible Storage Sharing (FSS) Disk group creation for local disks that are exported may fail if the number of disks used for disk group creation is greater than 150, with the following error message:

```
VxVM vxdg ERROR V-5-1-585 Disk group test_dg: cannot create: Transaction
    locks timed out
```

A similar error can be seen while adding more that 150 locally exported disks (with `vxdg adddisk`) to the FSS disk group, with the following error message:

```
VxVM vxdg ERROR V-5-1-10127 associating disk-media emc0_0839 with emc0_0839:
        Transaction locks timed out
```

Workaround:

Create an FSS disk group using 150 or less locally exported disks and then do an incremental disk addition to the disk group with 150 or less locally exported disks at a time.

vxconfigrestore is unable to restore FSS cache objects in the pre-commit stage (3461928)

While restoring a Flexible Storage Sharing (FSS) disk group configuration that has cache objects configured, the following error messages may display during the pre-commit phase of the restoration:

```
VxVM vxcache ERROR V-5-1-10128  Cache object meta-data update error
VxVM vxcache ERROR V-5-1-10128  Cache object meta-data update error
VxVM vxvol WARNING V-5-1-10364 Could not start cache object
VxVM vxvol ERROR V-5-1-11802 Volume volume_name cannot be started
VxVM vxvol ERROR V-5-1-13386 Cache object on which Volume volume_name
 is constructed is not enabled
VxVM vxvol ERROR V-5-1-13386 Cache object on which Volume volume_name
 is constructed is not enabled
```

The error messages are harmless and do not have any impact on restoration. After committing the disk group configuration, the cache object and the volume that is constructed on the cache object are enabled.

Change in naming scheme is not reflected on nodes in an FSS environment (3589272)

In a Flexible Storage Sharing (FSS) environment, if you change the naming scheme on a node that has local disks, the remote disk names are not reflected with the corresponding name change. If you change the naming scheme on a node where exported disks are present, to reflect the updated remote disk names, you must either export the disks again or restart the node where the remote disks are present

Workaround:

There is no workaround for this issue.

vxassist does not create data change logs on all mirrored disks, if an FSS volume is created using DM lists (3559362)

When a Flexible Storage Sharing (FSS) volume is created using DM lists, the `vxassist` command does not create data change logs on all the mirrored disks; the number of DCO mirrors is not equal to the number of data mirrors. The `vxassist` command creates a two-way DCO volume.

Workaround:

Manually add a DCO mirror using the `vxassist -g` *`diskgroup`* `mirror` *`dco_volume`* command.

Intel SSD cannot be initialized and exported (3584762)

Initializing an Intel SSD with the Flexible Storage Sharing (FSS) export option may fail with the following error message:

```
VxVM vxedpart ERROR V-5-1-10089 partition modification failed: Device
or resource busy
```

Workaround:

Initialize the private region of the SSD disk to 0 and retry the disk initialization operation.

For example:

```
# dd if=/dev/zero of=/dev/vx/dmp/intel_ssd0_0 bs=4096 count=1

# vxdisksetup -i intel_ssd0_0 export
```

VxVM may report false serial split brain under certain FSS scenarios (3565845)

In a Flexible Storage Sharing (FSS) cluster, as part of a restart of the master node, internal storage may become disabled before network service. Any VxVM objects on the master node's internal storage may receive I/O errors and trigger an internal transaction. As part of this internal transaction, VxVM increments serial split brain (SSB) ids for remaining attached disks, to detect any SSB. If you then disable the network service, the master leaves the cluster and this results in a master takeover. In such a scenario, the master takeover (disk group re-import) may fail with a false split brain error and the `vxsplitlines` output displays 0 or 1 pools.

For example:

```
Syslog: "vxvm:vxconfigd: V-5-1-9576 Split Brain. da id is 0.2,
while dm id is 0.3 for dm disk5mirr
```

Workaround:

To recover from this situation

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

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

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

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

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

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.

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SF Oracle RAC.

Workaround:

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:

```
$ 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.

Workaround: 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 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 6.2 (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 6.2.

When upgrading from SF Oracle RAC version 5.0 to SF Oracle RAC 6.2 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.

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

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

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

```
Retrieving snapshot information ...                          Done
Importing snapshot diskgroups ...                            Done
Mounting snapshot volumes ...                                Done
Preparing parameter file for clone database ...              Done
Mounting clone database ...
ORA-00845: MEMORY_TARGET not supported on this system


SFDB vxsfadm ERROR V-81-0612 Script
/opt/VRTSdbed/applications/oracle/flashsnap/pre_preclone.pl failed.
```

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

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

To remount the /dev/shm file system with sufficient available space

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

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

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

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

4. Start the database.

Clone fails with error "ORA-01513: invalid current time returned by operating system" with Oracle 11.2.0.3 (2804452)

While creating a clone database using any of the point-in-time copy services such as Flashsnap, SOS, Storage Checkpoint, or Filesnap, the clone fails. This problem appears to affect Oracle versions 11.2.0.2 as well as 11.2.0.3.

You might encounter an Oracle error such as the following:

```
/opt/VRTSdbed/bin/vxsfadm -s flashsnap -o clone
-a oracle -r dblxx64-16-v1 --flashsnap_name TEST11 --clone_path
/tmp/testRecoverdb --clone_name clone1
USERNAME:  oragrid
STDOUT:
Retrieving snapshot information ...                              Done
Importing snapshot diskgroups ...                                Done
Mounting snapshot volumes ...                                    Done

ORA-01513: invalid current time returned by operating system
```

This is a known Oracle bug documented in the following Oracle bug IDs:

- Bug 14102418: DATABASE DOESNT START DUE TO ORA-1513
- Bug 14036835: SEEING ORA-01513 INTERMITTENTLY

Workaround: Retry the cloning operation until it succeeds.

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.

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'
```

In the cloned database, the seed PDB remains in the mounted state (3599920)

In Oracle database version 12.1.0.2, when a container database (CDB) is cloned, the **PDB$SEED** pluggable database (PDB) remains in the mounted state. This

behavior is observed because of the missing datafiles in the cloned database for all point-in-time copies.

When you attempt to open the cloned seed database, the following error is reported:

```
"ORA-01173" oracle error.
...
SFDB vxsfadm ERROR V-81-0564 Oracle returned error.

Reason: ORA-01122: database file 15 failed verification check
ORA-01110: data file 15: '/tmp/test1/data/sfaedb/newtbs1.dbf'
ORA-01202: wrong incarnation of this file - wrong creation time
...
```

Workaround: There is no workaround for this issue.

Cloning of a container database may fail after a reverse resync commit operation is performed (3509778)

After a reverse resync operation is performed, the cloning of a container database may fail with the following error message:

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

Reason: ORA-01503: CREATE CONTROLFILE failed
ORA-01189: file is from a different RESETLOGS than previous files
ORA-01110: data file 6: '/tmp/testRecoverdb/data/sfaedb/users01.dbf'
```

Workaround: There is no workaround for this issue.

If one of the PDBs is in the read-write restricted state, then cloning of a CDB fails (3516634)

Cloning a container database (CDB) for point-in-time copies fails if some of the pluggable databases (PDBs) are open in the restricted mode. The failure occurs with the following error message:

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

Reason: ORA-65106: Pluggable database #3 (PDB1) is in an invalid state.
```

Workaround: There is no workaround for this issue.

Cloning of a CDB fails for point-in-time copies when one of the PDBs is in the read-only mode (3513432)

For Oracle version 12.1.0.1, cloning a container database (CDB) fails if one of the pluggable databases (PDBs) is in the read-only mode. The failure occurs with the following error message:

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

Reason: ORA-00376: file 9 cannot be read at this time
ORA-01111: name for data file 9 is unknown - rename to correct file
ORA-01110: data file 9: '/ora_base/db_home/dbs/MISSING00009'...
```

Workaround: There is no workaround for this issue.

If a CDB has a tablespace in the read-only mode, then the cloning fails (3512370)

For Oracle version 12.1.0.1, when a container database (CDB) has a tablespace in the read-only mode for all point-in-time copies, cloning of that CDB fails with the following error message:

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

Reason: ORA-01122: database file 15 failed verification check
ORA-01110: data file 15: '/tmp/test1/data/sfaedb/newtbs1.dbf'
ORA-01202: wrong incarnation of this file - wrong creation time
...
```

Workaround: There is no workaround for this issue.

If any SFDB installation with authentication setup is upgraded to 6.2, the commands fail with an error (3644030)

The commands fail with the error message similar to the following:

```
SFDB vxsfadm ERROR V-81-0450 A remote or privileged command could not be
executed on prodhost

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

Action: Verify that the prodhost is reachable. If it is, verify
that
the vxdbd daemon is enabled and running using the [
```

```
/opt/VRTS/bin/sfae_config
status ] command, and enable/start vxdbd using the [
/opt/VRTS/bin/sfae_config
enable ] command if it is not enabled/running. Also make sure you are
authorized to run SFAE commands if running in secure mode.
```

Workaround: Set up the authentication for SFDB again. See *Storage and Availability Management for Oracle Databases* or *Storage and Availability Management for DB2 Databases*.

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 83.

SF Oracle RAC supports IPv4 addresses only for Oracle HAIP configurations (3656701)

SF Oracle RAC supports IPv4 addresses only for static IP addresses in Oracle high availability IP (HAIP) configurations.

Configurations that use IPv6 or a mix of IPv4 and IPv6 will fail internally though no error may be reported by the product installer.

Cloned disks operations not supported for FSS disk groups

In this release, the VxVM cloned disks operations are is not supported with FSS disk groups. If you clone a disk in the FSS disk groups, the cloned device cannot be imported. If you prefer to use hardware mirroring for disaster recovery purposes, you need to make sure that such devices should not be used to create FSS disk groups.

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

VRTSvxvm upgrade using native command fails (3384435)

If you upgrade the `VRTSvxvm` rpm using the native command, `rpm -Uvh VRTSvxvm`, the command fails.

Workaround:

1. Remove the older `VRTSlvmconv` rpm.

   ```
   # rpm -e VRTSlvmconv
   ```

2. Upgrade the `VRTSvxvm` rpm.

   ```
   # rpm -Uvh VRTSvxvm
   ```

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.

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.

SELinux supported in disabled and permissive modes only

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

Policy-managed databases not supported by CRSResource agent

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

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

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

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

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

Cached ODM not supported in SF Oracle RAC environments

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

Unsupported FSS scenarios

The following scenario is not supported with Flexible Storage Sharing (FSS):

Symantec NetBackup backup with FSS disk groups

The sfcache operations may display error messages in the caching log when the operation completed successfully (3611158)

The `sfcache` command calls other commands to perform the caching operations. If a command fails, additional commands may be called to complete the operation. For debugging purposes, the caching log includes all of the success messages and failure messages for the commands that are called.

If the `sfcache` command has completed successfully, you can safely ignore the error messages in the log file.

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

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 6.2, you must upgrade the Oracle binaries and database to version 10.2.0.5, before upgrading to 6.2.

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.

Symantec Storage Foundation for Oracle RAC documentation

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

Table 1-13 Symantec Storage Foundation for Oracle RAC documentation

Document title	File name	Description
Symantec Storage Foundation for Oracle RAC Release Notes	sfrac_notes_62_lin.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of the product.
Symantec Storage Foundation for Oracle RAC Installation and Configuration Guide	sfrac_install_62_lin.pdf	Provides information required to install and configure the product.
Symantec Storage Foundation for Oracle RAC Administrator's Guide	sfrac_admin_62_lin.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 1-17 on page 86.

Symantec Storage Foundation Cluster File System High Availability documentation

Table 1-14 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 1-17 on page 86.

Table 1-14 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_62_lin.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_62_lin.pdf	Provides information required to install the product.
Symantec Storage Foundation Cluster File System High Availability Administrator's Guide	sfcfs_admin_62_lin.pdf	Provides information required for administering the product.

Symantec Cluster Server documentation

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

Table 1-15 Symantec Cluster Server documentation

Title	File name	Description
Symantec Cluster Server Release Notes	vcs_notes_62_lin.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_62_lin.pdf	Provides information required to install the product.
Symantec Cluster Server Administrator's Guide	vcs_admin_62_lin.pdf	Provides information required for administering the product.

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

Title	File name	Description
Symantec High Availability Solution Guide for VMware	sha_solutions_62_vmware_lin.pdf	Provides information on how to install, configure, and administer Symantec Cluster Server in a VMware virtual environment, by using the VMware vSphere Client GUI.
Symantec Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents_62_lin.pdf	Provides information about bundled agents, their resources and attributes, and more related information.
Symantec Cluster Server Generic Application Agent Configuration Guide	vcs_gen_agent_62_lin.pdf	Provides notes for installing and configuring the generic Application agent.
Symantec Cluster Server Agent Developer's Guide (This document is available online only.)	vcs_agent_dev_62_unix.pdf	Provides information about the various Symantec agents and procedures for developing custom agents.
Symantec Cluster Server Agent for DB2 Installation and Configuration Guide	vcs_db2_agent_62_lin.pdf	Provides notes for installing and configuring the DB2 agent.
Symantec Cluster Server Agent for Oracle Installation and Configuration Guide	vcs_oracle_agent_62_lin.pdf	Provides notes for installing and configuring the Oracle agent.
Symantec Cluster Server Agent for Sybase Installation and Configuration Guide	vcs_sybase_agent_62_lin.pdf	Provides notes for installing and configuring the Sybase agent.

Symantec Storage Foundation documentation

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

Table 1-16 Symantec Storage Foundation documentation

Document title	File name	Description
Symantec Storage Foundation Release Notes	sf_notes_62_lin.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_62_lin.pdf	Provides information required to install the product.

Table 1-16 Symantec Storage Foundation documentation *(continued)*

Document title	File name	Description
Symantec Storage Foundation Administrator's Guide	sf_admin_62_lin.pdf	Provides information required for administering the product.
Symantec Storage Foundation: Storage and Availability Management for DB2 Databases	sfhas_db2_admin_62_unix.pdf	Provides information about the deployment and key use cases of the SFDB tools with Storage Foundation High Availability (SFHA) Solutions products in DB2 database environments. It is a supplemental guide to be used in conjunction with SFHA Solutions product guides.
Symantec Storage Foundation: Storage and Availability Management for Oracle Databases	sfhas_oracle_admin_62_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_62_lin.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 1-17 lists the documentation for Symantec Storage Foundation and High Availability Solutions products.

Table 1-17 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_62_unix.pdf	Provides information about the new features and enhancements in the release.

Table 1-17 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 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_62_lin.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_62_lin.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 SmartIO for Solid State Drives Solutions Guide	sfhas_smartio_solutions_62_lin.pdf	Provides information on using and administering SmartIO with SFHA solutions. Also includes troubleshooting and command reference sheet for SmartIO.
Symantec Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide (This document is available online.)	sfhas_dr_impl_62_lin.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_62_lin.pdf	Provides information on using Symantec Replicator Option for setting up an effective disaster recovery plan by maintaining a consistent copy of application data at one or more remote locations. Symantec Replicator Option provides the flexibility of block-based continuous replication with Symantec Volume Replicator Option (VVR) and file-based periodic replication with Symantec File Replicator Option (VFR).

Table 1-17 Symantec Storage Foundation and High Availability Solutions products documentation *(continued)*

Document title	File name	Description
Symantec Storage Foundation and High Availability Solutions Troubleshooting Guide	sfhas_tshoot_62_lin.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.

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

To edit the man(1) configuration file

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

   ```
   export LC_ALL=C
   ```

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

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

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

   where other man paths are specified in the configuration file.

3. Add new section numbers. Change the line:

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

   to

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

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

https://sort.symantec.com/documents