

Veritas Storage Foundation™ Cluster File System High Availability Release Notes

Solaris

6.0 Platform Release 1

Veritas Storage Foundation™ Cluster File System High Availability Release Notes

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

Product version: 6.0 PR1

Document version: 6.0PR1.3

Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America [supportolutions@symantec.com](mailto: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 Cluster File System High Availability Release Notes

This document includes the following topics:

- [About this document](#)
- [Component product release notes](#)
- [About Veritas Storage Foundation Cluster File System High Availability](#)
- [About Symantec Operations Readiness Tools](#)
- [Important release information](#)
- [Changes introduced in 6.0 PR1](#)
- [Changes introduced in 6.0](#)
- [No longer supported](#)
- [System requirements](#)
- [Known issues](#)
- [Software limitations](#)
- [Documentation errata](#)
- [Documentation](#)

About this document

This document provides important information about Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) version 6.0 PR1 for Solaris. Review this entire document before you install or upgrade SFCFSHA.

The information in the Release Notes supersedes the information provided in the product documents for SFCFSHA.

This is Document version: 6.0PR1.3 of the *Veritas Storage Foundation Cluster File System High Availability Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

<https://sort.symantec.com/documents>

Component product release notes

In addition to reading this Release Notes document, review the component product release notes before installing the product.

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

/product_name/docs

Symantec recommends copying the `docs` directory on the software media that contains the product guides to the `/opt/VRTS` directory on your system.

This release includes the following component product release notes:

- *Veritas Storage Foundation Release Notes* (6.0 PR1)
- *Veritas Cluster Server Release Notes* (6.0 PR1)

About Veritas Storage Foundation Cluster File System High Availability

Veritas Storage Foundation Cluster File System High Availability by Symantec extends Veritas Storage Foundation to support shared data in a storage area network (SAN) environment. Using Storage Foundation Cluster File System High Availability, multiple servers can concurrently access shared storage and files transparently to applications.

Veritas Storage Foundation Cluster File System High Availability also provides increased automation and intelligent management of availability and performance.

Veritas Storage Foundation Cluster File System High Availability includes Veritas Cluster Server, which adds high availability functionality to the product.

To install the product, follow the instructions in the *Veritas Storage Foundation Cluster File System High Availability Installation Guide*.

For information on high availability environments, read the Veritas Cluster Server documentation.

About Symantec Operations Readiness Tools

[Symantec Operations Readiness Tools \(SORT\)](#) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Symantec products.

SORT can help you do the following:

- | | |
|---|--|
| Prepare for your next installation or upgrade | <ul style="list-style-type: none">■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture.■ Analyze systems to determine if they are ready to install or upgrade Symantec products.■ Download the latest patches, documentation, and high availability agents from a central repository.■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems. |
| Manage risks | <ul style="list-style-type: none">■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDDs), and high availability agents from a central repository.■ Identify and mitigate system and environmental risks.■ Display descriptions and solutions for hundreds of Symantec error codes. |
| Improve efficiency | <ul style="list-style-type: none">■ Find and download patches based on product version and platform.■ List installed Symantec products and license keys.■ Tune and optimize your environment. |

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

To access SORT, go to:

<https://sort.symantec.com>

Important release information

- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:
<http://www.symantec.com/docs/TECH164885>
- For the latest patches available for this release, go to:
<http://sort.symantec.com/>
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:
<http://www.symantec.com/docs/TECH170013>
Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

Changes introduced in 6.0 PR1

This section lists the changes in Veritas Storage Foundation Cluster File System High Availability 6.0 PR1.

Changes related to installation

The product installer includes the following changes.

New packaging framework in Solaris 11

Solaris 11 introduces the new IPS packaging framework. All of the Symantec packages have been modified to IPS packages.

Changes introduced in 6.0

This section lists the changes in Veritas Storage Foundation Cluster File System High Availability 6.0.

Changes related to Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)

SFCFSHA includes the following changes in 6.0 PR1:

Changes to the Intelligent Monitoring Framework feature

In this release, the Intelligent Monitoring Framework (IMF) feature is enabled by default. In previous releases, IMF was disabled by default.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for more information on enabling and disabling IMF.

See the *Veritas Cluster Server Installation Guide* and the *Veritas Cluster Server Administrator's Guide* for detailed information about IMF.

Availability of shared disk group configuration copies

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

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

Enhancements to CVM detach policies

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

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

Enhancements to master node selection for failover

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

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

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

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

Node join with DGDISABLED disk groups

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

Entering and displaying values in human-friendly units

Storage Foundation now supports reporting and inputting values in human-friendly units.

The following commands were modified to display human-friendly units:

- `diskusg`
- `ff`
- `fsadm`
- `fsckptadm`
- `fsvoladm`
- `vxdg free`
- `vxdisk list`
- `vxdisk -o thin list`
- `vxdisk -o thin,fssize list`
- `vxdmpadm iostat show`
- `vxedquota`
- `vxmemstat`
- `vxprint`
- `vxquot`
- `vxquota`
- `vxrepquota`
- `vxstat`
- `vxtune`

See the manual pages for more information.

Added nodes into the Clustered NFS cluster

The `cfsshare` command has the ability to add a node in the Clustered NFS (CNFS) cluster.

See the *Veritas Storage Foundation Cluster File System Administrator's High Availability Guide*.

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

Displaying SFCFSA information with `vxlist`

The `vxlist` command is a new display command that provides a consolidated view of the SFCFSA configuration. The `vxlist` command consolidates information from Veritas Volume Manager (VxVM) and Veritas File System (VxFS). The `vxlist` command provides various options to display information. For example, use the following form of the command to display file system information including information about the volume, disk group, and so on. In previous releases, you needed to run at least two commands to retrieve the following information.

```
# /opt/VRTSsfmh/bin/vxlist fs
```

TY	FS	FSTYPE	SIZE	FREE	%USED	DEVICE_PATH	MOUNT_POINT
fs	/	ext3	65.20g	51.70g	17%	/dev/sda1	/
fs	mnt	vxfs	19.84g	9.96g	49%	/dev/vx/dsk/bardg/voll	/mnt

For help on the `vxlist` command, enter the following command:

```
# vxlist -H
```

See the `vxlist(1m)` manual page.

Recovery for synchronization tasks

In this release, VxVM tracks the plex synchronization for the following commands: `vxplex att`, `vxassist mirror`, `vxsnap addmir`, `vxsnap reattach`, and `vxsnap restore`. If the system crashes or the `vxconfigd` daemon fails, VxVM provides automatic recovery for the synchronization task. When the system is recovered, VxVM restarts the synchronization from the point where it failed. The synchronization occurs in the background, so the volume is available without delay.

Secure deletion of Veritas Volume Manager disks

When you decommission a disk that contained sensitive data, you may need to destroy any remaining data on the disk. In this release, VxVM provides the ability to shred the data on the disk to minimize the chance that the data is recoverable.

When you specify the disk shred operation, VxVM shreds the entire disk, including any existing disk labels. After the shred operation, VxVM writes a new empty label on the disk to prevent the disk from going to the error state. The VxVM shred operation overwrites all of the addressable blocks with a digital pattern in one, three, or seven passes.

Caution: All data in the volume will be lost when you shred it. Make sure that the information has been backed up onto another storage medium and verified, or that it is no longer needed.

For more information on shredding disks, see the *Veritas Storage Foundation Administrator's Guide*.

Changes related to Veritas File System

Veritas File System includes the following changes in 6.0 PR1:

Default disk layout Version is now 9

In this release, disk layout Version 9 is now the default version, which enables support for the following features:

- File compression
- Data deduplication

See the *Administrator's Guide*.

Data deduplication on Solaris SPARC

On Solaris SPARC, you can run post-process periodic deduplication in a file system, which eliminates duplicate data without any continuous cost. This feature requires an Enterprise license.

See the *Administrator's Guide*.

File compression

You can compress files to reduce the space used, while retaining the accessibility of the files and having the compression be transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files.

See the *Administrator's Guide*.

Multi-threaded Thin Reclamation

You can perform multi-threaded Thin Reclamation operations for improved performance.

See the `fsadm_vxfs(1M)` and `vxfs_ts_reclaim(3)` manual pages.

Storage Checkpoints

The following changes were made to Storage Checkpoints:

- You can tune Veritas File System (VxFS) file systems to create removable Storage Checkpoints by default.
See the `vxtunefs(1M)` manual page.
- VxFS now attempts to remove removable Storage Checkpoints if the file system does not have enough space instead of failing the operation.
- Storage Checkpoints have improved visibility to the file system. With the `ckptautomnt` mount option, all Storage Checkpoints are made accessible automatically through a directory in the root directory of the file system that has the special name `.checkpoint`, which does not appear in directory listings. Inside this directory is a directory for each Storage Checkpoint in the file system. Each of these directories behave as a mount of the corresponding Storage Checkpoint with some exceptions.
See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

Partitioned directories

Normally, a large volume of parallel threads performing access and updates on a directory that commonly exist in an file system suffers from exponentially longer wait times for the threads. This feature creates partitioned directories to improve the directory performance of file systems. When any directory crosses the tunable threshold, this feature takes an exclusive lock on the directory inode and redistributes the entries into various respective hash directories. These hash directories are not visible in the name-space view of the user or operating system. For every new create, delete, or lookup thread, this feature performs a lookup for the respective hashed directory (depending on the target name) and performs the operation in that directory. This leaves the parent directory inode and its other hash directories unobstructed for access, which vastly improves file system performance.

See the *Administrator's Guide*.

Delayed allocation for extending writes

Performance of extending writes on local mounts is improved using the delayed allocation feature, which is turned on by default for all applicable writes.

See the *Administrator's Guide*.

Migrating a source file system to the VxFS file system over NFS v3

NFS is one of the most commonly used file systems in network-attached storage solutions and is one of the standard file sharing mechanisms used in UNIX environments. This feature enables you to migrate a source file system to the VxFS file system over your existing NFS v3 solution.

See the *Veritas Storage Foundation and High Availability Solutions Solutions Guide*.

vxfsconvert can upgrade additional Veritas File System disk layout versions

The `vxfsconvert` command can upgrade the VxFS disk layout Version 4 and 5.

FileSnap creation over Network File System

You can create a FileSnap over Network File System (NFS) by creating a hard link from an existing file to a new file with the extension “:snap:vxfs:”.

See the *Administrator's Guide*.

Free space defragmentation

You can now specify the `-C` option with the `fsadm` command to minimize file system free space fragmentation. This attempts to generate bigger chunks of free space in the specified device.

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 6.0 PR1:

Creating a volume of maximum size

In previous releases, Veritas Volume Manager provided a two-step approach to creating a volume of the maximum size. You had to run the `vxassist maxsize` command to find the maximum size of the volume to be created with the given constraints. Then, you had to run the `vxassist make` command and specify the volume size as the maximum determined by the `vxassist maxsize` command.

In this release, you can create a maximum sized volume with a single command. Specify the `vxassist make` command with the `maxsize` keyword. The `vxassist` command creates the maximum sized volume possible, taking into consideration any other allocation attributes that you specify.

Changes to the instant snapshot (version 20) data change object (DCO) volume layout

In this release, the volume layout of the data change object (DCO) has been changed to improve the I/O performance and scalability of instant snapshots. The change in layout does not alter how you administer instant snapshots. The only visible effect is in improved I/O performance and in some cases, increased size of DCO volume. As with previous releases, you create DCOs for instant snapshots using "vxsnap prepare" or by specifying "logtype=dco dconversion=20" while creating volume with "vxassist make".

The instant snapshot DCO (previously known as a version 20 DCO) now uses dynamic creation of maps on the preallocated storage.

Veritas Volume Manager throttling of administrative I/O

In this release, Veritas Volume Manager (VxVM) provides throttling of administrative I/O. During heavy I/O loads, VxVM throttles I/O that it creates to do administrative operations. This behavior ensures that the administrative I/Os do not affect the application I/O performance. When the application I/O load is lighter, VxVM increases the bandwidth usage for administrative I/O operations.

VxVM automatically manages the I/O throttling for administrative tasks, based on its perceived load on the storage. Currently, I/O throttling is supported for the copy operations which use ATOMIC_COPY and involve one destination mirror. The I/O throttling is transparent, and does not change the command usage or output. The following commands are supported:

- vxassist mirror
- vxassist snapcreate
- vxevac
- vxplex att
- vxplex cp
- vxplex mv
- vxprint
- vxsnap addmir
- vxsnap reattach
- vxsd mv
- vxtune

The administrative I/O operations allocate memory for I/O from a separate memory pool. You can tune the maximum size of this pool with the tunable parameter, `vol_max_adminio_poolsz`.

See the *Veritas Storage Foundation Administrator's Guide* for information about tuning the `vol_max_adminio_poolsz` parameter.

Command completion for Veritas commands

Veritas Storage Foundation Cluster File System High Availability now supports command completion for Veritas Volume Manager (VxVM) commands and Dynamic Multi-Pathing (DMP) commands. In this release, command completion is supported only on the bash shell. The shell must be bash version 2.4 or later.

To use this feature, press **Tab** while entering a supported VxVM or DMP command. The command is completed as far as possible. When there is a choice, the command completion displays the next valid options for the command. Enter one of the displayed values. A value in brackets indicates a user-specified value.

Note: Platform-specific options are not supported with command completion in this release.

The following commands support command completion:

- `vxassist`
- `vxdisk`
- `vxplex`
- `vxprint`
- `vxsnap`
- `vxstat`
- `vxtune`
- `vxcache`
- `vxconfigd`
- `vxtask`
- `vxreattach`
- `vxddmpadm`
- `vxddladm`
- `vxvol`

- vxcdsconvert
- vxresize
- vxctl
- vxsd
- vxdisksetup
- vxdiskunsetup
- vxrecover
- vxedit
- vxdg
- vxclustadm

vxdisk -o thin list command now shows the disk space used by a VxFS file system

The `vxdisk -o thin list` command now shows the disk space used by a VxFS file system.

Changes related to Veritas Dynamic Multi-Pathing (DMP)

The following sections describe changes in this release related to Veritas Dynamic Multi-Pathing (DMP).

DMP detects "persist through power loss" storage device server capability

In this release, DMP detects when a storage device server has the capability "persist through power loss". Certain arrays, such as Oracle's Sun Storage 7310, use this capability to preserve the persistent reservation and registrations across power cycles, controller reboots, and other similar operations.

If DMP detects that the device supports this capability, then DMP sets the APTPL (Activate Persist Through Power Loss) bit to 1 in the PERSISTENT RESERVE OUT parameter data sent with a REGISTER, REGISTER AND IGNORE EXISTING KEY service action, according to SPC-3 specifications.

When APTPL is set to 1, the persistent reservation (PR) keys are preserved during array controller takeover or failback operations.

Tuning Dynamic Multi-Pathing with templates

Veritas Dynamic Multi-Pathing (DMP) has multiple tunable parameters and attributes that you can configure for optimal performance. In this release, DMP introduces a template method to update several tunable parameters and attributes

with a single operation. The template represents a full or partial DMP configuration, showing the values of the parameters and attributes of the host.

To view and work with the tunable parameters, you can dump the configuration values of the DMP tunable parameters to a file. Edit the parameters and attributes, if required. Then, load the template file to a host to update all of the values in a single operation.

For more information about tuning DMP with templates, see the *Veritas Dynamic Multi-Pathing Administrator's Guide*.

You can manage the DMP configuration file with the `vxdmpadm config` commands.

See the `vxdmpadm(1M)` man page.

Changes to DMP support for ALUA arrays

In this release, DMP has improved support for ALUA arrays. DMP now efficiently handles most implementations of the ALUA standard. The enhancements include the following:

- DMP now detects whether an ALUA array is A/A-A, A/A or A/P-F.
- DMP handles the array state correctly, when a node is taken out of the cluster. The enclosure level attribute failoverpolicy is now set internally.
- DMP handles Standby and unavailable LUN states for ALUA arrays.
- DMP monitors LUN ownership changes. DMP can shift the I/O load depending on the current state of the LUN.

Dynamic Multi-Pathing (DMP) detects and reports extended attributes from Veritas Operations Manager

If you have Veritas Operations Manager (VOM), and you have configured a central Management Server, the Device Discovery layer (DDL) of DMP can obtain extended attributes for managed hosts. DDL obtains these additional attributes out of band from the VOM database. DMP displays these attributes as output of the `vxdisk -p list` command.

See the *Administrator's Guide*.

DMP tunable parameter `dmp_enable_restore` renamed to `dmp_restore_state`

The DMP tunable parameter `dmp_enable_restore` has been renamed to `dmp_restore_state`. The `dmp_restore_state` tunable can have the following values:

- `enabled`
Enables and starts the DMP path restoration thread.
- `disabled`

Stops and disables the DMP path restoration thread.

- stopped

Stops the DMP path restoration thread until the next device discovery cycle.

Command completion for DMP commands

Veritas Dynamic Multi-Pathing (DMP) now supports command completion for DMP commands. In this release, command completion is supported only on the bash shell. The shell must be bash version 2.4 or later.

To use this feature, press **Tab** while entering a supported VxVM or DMP command. The command is completed as far as possible. When there is a choice, the command completion displays the next valid options for the command. Enter one of the displayed values. A value in brackets indicates a user-specified value.

Note: Platform-specific options are not supported with command completion in this release.

The following commands support command completion:

- `vxdisk`
- `vxdlpadm`
- `vxddladm`

DMP enhancements

The following DMP enhancements have been made in this release:

- The `vxdlpadm enable` command and the `vxdlpadm disable` command now accept multiple controllers on the command line.
- In addition, you can now enable or disable paths between a given controller and a port-id pair. If you specify both an HBA controller and an array port, DMP disables I/O on the specific portion of the Storage Area Network (SAN).
- The `vxdlpadm stat error` command and the `vxdlpadm stat restored` command are deprecated.
To see status for the restore tasks, use the `vxdlpadm gettune` command.
- Excluding or including paths from DMP is deprecated.
Excluding paths from DMP but not from VxVM can lead to unsupported configurations. The command operations to exclude or include paths from DMP are now deprecated. You can exclude or include paths from VxVM. The deprecated commands are as follows:

```
vxdlpadm exclude dmp
```

```
vxdlpadm include dmp  
vxdiskadm: DMP options under Suppressing or including devices for  
VxVM
```

- `vxddladm list devices` command now displays the name of the ASL even if the device is skipped.
- `vxddladm status eventsource` is added to show the status of the `vxesd` daemon
- `vxscsiinq` diagnostic utility is enhanced to take hexadecimal page numbers as arguments.

Changes related to replication

Veritas Storage Foundation and High Availability Solutions includes the following changes related to replication in 6.0 PR1:

vvrcheck configuration utility

There is now a configuration utility, `/etc/vx/diag.d/vvrcheck`, that displays current replication status, detects and reports configuration anomalies, and creates statistics files that can be used by display tools. The `vvrcheck` also runs diagnostic checks for missing daemons, valid licenses, and checks on the remote hosts on the network. For more information, see the `vvrcheck(1M)` man page.

SmartMove for VVR

The initial sync between the Primary and Secondary is performed using the autosync option. The autosync to sync the volume now uses the SmartMove API from VxFS and provides the data only sync between the Primary and Secondary. This increases the initial autosync performance, which is dependent on the file system usage in the volume. This feature also helps thin provision LUNs configured on the Secondary site to use storage space only for data.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Thin provisioning and reclamation support for VVR

Storage reclamation is now possible on VVR volumes with VxFS file system on it. The storage corresponding to the volumes on the Secondary RVG is automatically reclaimed when the Primary volumes are reclaimed. The existing `vxdisk reclaim` or `fsadm -R` commands function for reclaiming VVR objects as well. For storage reclamation to work, the volumes on the Primary RVG must be mounted.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Enable compression with VVR

VVR compression lets you send data over the network in a compressed format from a Primary to one or more Secondary hosts. Compression reduces network bandwidth consumption and is useful in scenarios where there is low available bandwidth or where the bandwidth is shared among several applications. The compression option can be enabled on a per system or per Secondary basis using the CLI.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Replication performance improvement

Replication performance is improved by introducing Secondary logging (logging the I/O on the Secondary SRL before writing to the data volume). The primary requirement for this feature to work is to have the same size SRL on both the Secondary and Primary. The Secondary SRL is used for staging the I/O from the Primary, and parallelize the data volume write. This improves the replication performance both in VVR and CVR. By default, this feature is enabled in 6.0.

There are other replication-specific tunables that may be increased to obtain the maximum replication performance.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Support for 8-node cluster applications

In a shared disk group environment, VVR supports replication of 8-node cluster applications. In previous releases, support was limited to 4-node cluster applications.

The following improvements enable scalability to 8-node support:

- Improved message processing allows the logowner to process more messages per second, resulting in improved application throughput
- Secondary logging feature improves replication performance
- Improved CPU usage provides more CPU cycles to the logowner to process requests from other nodes in a cluster
- Increased limit on max outstanding I/Os with VVR

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Changes related to SFDB tools

The following sections describe the changes related to Storage Foundation for Databases (SFDB) tools in 6.0 PR1.

Support for space-optimized snapshots for database cloning

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

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

Enhancements to Cached ODM Advisor (dbed_codm_adm)

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

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

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

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

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

Single CLI for different point-in-time copy operations

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

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

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

Support for file-level snapshots for database cloning

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

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

Enhanced authentication support

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

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

SmartTier integration with OEM

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

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

Changes to SFCFSHA clusters in secure mode

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

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

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

Changes to LLT

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

- Faster detection of LLT link failures

LLT now has the capability to detect the LLT link failures immediately using the operating system's link-down notification feature.

- LLT now supports VLAN tagging (IEEE 802.1Q).
- The `lltconfig` command includes the following new options:
 - `-Q`
You can use this option to configure a link as "hidden". LLT needs this type of link when you enable faster detection of link failure.
 - `-j`
You can use this option to set the LLT link failure detection level.
 - The command also includes a new timer parameter `linkstable`. This parameter specifies the amount of time to wait before LLT processes the link-down event for any link of the local node when you have enabled faster link failure detection. This `linkstable` prevents a flaky link from causing unnecessary membership changes.
 - `-N`
You can use this option to list all the used cluster IDs.
 - `-M`
You can use this option to display the currently loaded LLT module version information.

See the `lltconfig` manual page for more information.

The `llttab` file includes the following changes:

- `hidden`
The new keyword `hidden` indicates that the link is a hidden link which LLT uses when you enable faster detection of link failure.
- `set-linkfaildetectlevel`
You can use this new command option in the `/etc/llttab` file to set the LLT link failure detection level.
- Link utilization statistics are enhanced that help in the root cause analysis of performance related issues.
- Periodic flushing of ARP cache is disabled.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* and the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

Changes to GAB

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

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

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

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

See the Veritas Storage Foundation Cluster File System High Availability Administrator's Guide for more details.

GAB can now recognize clients with names in addition to ports

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

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

The `gabconfig` command has new `-C` option

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

Changes to I/O fencing

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

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

At the time of a network partition, the VxFEN module elects the lowest node in each sub-cluster as the racer node to race for the coordination points on behalf

of the sub-cluster. The other spectator nodes wait on the racer node to do the fencing.

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

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

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

See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

Support for Non-SCSI3 fencing

In environments that do not support SCSI-3 PR, non-SCSI-3 fencing provides reasonable data protection by causing the winning side to delay by a configurable amount (`loser_exit_delay`, default 55). Additionally, Symantec has enhanced the fencing component to help panic the losing side quickly. Together, these enhancements help narrow down the window of potential data corruption drastically.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* and the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

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

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

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

- If a number of nodes in a cluster are not up, GAB port (port a) still comes up in all the member-nodes in the cluster.

- If the coordination points do not have keys from any non-member nodes, I/O fencing (GAB port b) also comes up.

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

See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

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

Installer support to migrate between fencing configurations in an online cluster

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

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

See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

Support for multiple virtual IP addresses in CP servers

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

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* and the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

Support for Quorum agent in CP servers

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

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* and the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

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

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

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

Licensing changes in the SFHA Solutions 6.0 release

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

- The Cluster File System license is deprecated. CFS customers are entitled to the Storage Foundation Cluster File System High Availability (SFCFS HA) functionality.
- The VVR Option is renamed as Veritas Replicator Option. This option includes VVR (volume-based replication) and the new file-based replication solution.
- The VVR Enterprise license is deprecated; you can use Storage Foundation Enterprise and add Veritas Replicator Option to get this functionality. VVR Enterprise customers are entitled to Storage Foundation Enterprise with Replicator Option.
- The VCS license enables full cluster functionality as well as the limited start/stop functionality.
- Storage Foundation Enterprise CFS for Oracle RAC (Linux/x64) customers are entitled to Storage Foundation Enterprise for Oracle RAC (Linux/x64.)

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

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

The following products are included in this release:

- Dynamic Multi-Pathing

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

HA: High Availability

HA/DR: High Availability and Disaster Recovery

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

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

Changes related to installation and upgrades

The product installer includes the following changes in 6.0 PR1.

Support for product deployment using Flash Archives on Solaris 10

Use a Flash archive to install the Storage Foundation and High Availability Suite of products and the Solaris 10 operating system in conjunction with JumpStart.

See the *Installation Guide* for more information.

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

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

Using the installer for Veritas Dynamic Multi-pathing (DMP)

You can use the script- or Web-based installer to install, configure, and uninstall Veritas Dynamic Multi-pathing. You can enable DMP using the DMP license or using any Storage Foundation license key.

Unencapsulation not required for some upgrade paths

Unencapsulation is no longer required for certain upgrade paths.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for more details.

Web-based installer supports configuring SFCFSHA cluster in secure mode

You can now configure the SFCFSHA cluster in secure mode using the Web-based installer.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for more details.

Web-based installer supports configuring disk-based fencing for SFCFSHA

You can now configure disk-based fencing for the SFCFSHA cluster using the Web-based installer.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for more details.

The installer can automatically detect and configure LLT links

The installer detects link connection status among all cluster nodes and chooses the most suitable links for LLT communication. It then can set the priority of the LLT private heartbeat links based on their media speed. Aggregated and bonded NICs are supported.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for more details.

The Web-based installer supports adding nodes

The Web-based installer has increased parity with the script-based installer. It now supports the ability to add nodes to a cluster. It also supports configuring secure clusters and fencing configuration.

The installer provides automated, password-less SSH configuration

When you use the installer, it enables SSH or RSH communication among nodes. It creates SSH keys and adds them to the authorization files. After a successful completion, the installer removes the keys and system names from the appropriate files.

When you use the installer for SSH communications, meet the following prerequisites:

- The SSH (or RSH) daemon must be running for auto-detection.
- You need the superuser passwords for the systems where you plan to install VCS.

The installer can check product versions and hotfixes

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

You can discover the following information with these commands:

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

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

Using the installer's postcheck option

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

Rolling upgrade improvements

The rolling upgrade procedure has been streamlined and simplified.

Packaging updates

The following lists the package changes in this release.

- **New `VRTSsfcp60` package for product installer scripts**
The `VRTSsfcp60` package is introduced in this release. The `VRTSsfcp60` package contains the installer scripts and libraries that the installer uses to install, configure and upgrade Veritas products.
- **New `VRTSfsadv` package for product data deduplication on Solaris SPARC**
The `VRTSfsadv` package is introduced in this release on Solaris SPARC. The `VRTSfsadv` package contains the libraries for the data deduplication feature.

For more information, see the *Installation Guide*.

Enhancements to collecting a VxExplorer troubleshooting archive

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

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

To learn more about using the data collector to collect a VxExplorer archive, see: www.symantec.com/docs/HOWTO32575

Changes related to product documentation

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

[Table 1-1](#) lists the documents introduced in this release.

Table 1-1 New documents

New documents	Notes
<i>Veritas Storage Foundation Installation Guide</i>	Installation and upgrade information for Storage Veritas Foundation.
<i>Veritas Storage Foundation Administrator's Guide</i>	Administration information for Veritas Storage Foundation.

Table 1-1 New documents (*continued*)

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

[Table 1-2](#) lists the documents that are deprecated in this release.

Table 1-2 Deprecated documents

Deprecated documents	Notes
<i>Veritas File System Administrator's Guide</i>	Content now appears in the <i>Veritas Storage Foundation Administrator's Guide</i> and in the <i>Veritas Storage Foundation Cluster File System High Availability Administrator's Guide</i> .
<i>Veritas Volume Manager Administrator's Guide</i>	Content now appears in the <i>Veritas Storage Foundation Administrator's Guide</i> and in the <i>Veritas Storage Foundation Cluster File System High Availability Administrator's Guide</i> .

Table 1-2 Deprecated documents (*continued*)

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

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

Table 1-3 Online documents

Document
<i>Veritas Cluster Server Agent Developer's Guide</i>
<i>Veritas Cluster Server Application Note: Dynamic Reconfiguration of Oracle Sun Servers</i>
<i>Veritas File System Programmer's Reference Guide</i>

No longer supported

The following features are not supported in this release of SFCFSHA products:

- Several documents are deprecated in this release.
 See [“Changes related to product documentation”](#) on page 34.
- Disk layout Version 4 and 5 are no longer supported. You cannot create nor mount a file system with disk layout Version 4 or 5. You can use the `vxfsconvert` utility to upgrade the disk layout to Version 7 or later after installing this release.
 See the `vxfsconvert(1M)` manual page.

- Disk layout Version 6 is deprecated. You can only local mount a file system with disk layout Version 6, and the only operation that you can perform is to upgrade the disk layout to a supported version by using the `vxupgrade` utility. Symantec recommends that you upgrade from Version 6 to the latest default disk layout version. You cannot create new file systems with disk layout Version 6. If you upgrade a file system from disk layout Version 6 to a later version, once the upgrade operation finishes, you must unmount the file system cleanly, then re-mount the file system.
See the `vxupgrade(1M)` manual page.

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

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

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

Cluster Volume Manager (CVM) no longer supported for Sun Clusters

In this release, Cluster Volume Manager (CVM) no longer supports Sun Clusters.

Solaris 8 and Solaris 9 branded zones are no longer supported for SFHA configurations

Solaris 8 and Solaris 9 branded zones are no longer supported for SFHA configurations in this release.

System requirements

The topics in this section describe the system requirements for this release.

- [Supported operating systems](#)
- [Veritas Storage Foundation for Database features supported in database environments](#)

Supported operating systems

[Table 1-4](#) shows the supported operating systems for this release.

Table 1-4 Supported operating systems

Operating systems	Levels	Chipsets
Solaris 11	SRU 11.0.x.y.z	SPARC
Solaris 11	SRU 11.0.x.y.z	x64

Note: Supported Repository Update (SRU) digits in the table represent Release.Update.SRU.Build.Respin.

Storage Foundation Cluster File System High Availability memory requirements

2 GB of memory is required.

Storage Foundation Cluster File System High Availability CPU requirements

A minimum of 2 CPUs is required.

Veritas Storage Foundation Cluster File System High Availability node requirements

Veritas Storage Foundation Cluster File System High Availability supports mixed cluster environments with Solaris 10 SPARC operating systems as long as all the nodes in the cluster have the same CPU architecture.

Veritas Storage Foundation for Database features supported in database environments

Veritas Storage Foundation for Database (SFDB) product features are supported for the following database environments:

Table 1-5 SFDB features supported in database environments

SFDB feature	DB2	Oracle	Sybase
Oracle Disk Manager, Cached Oracle Disk Manager	No	Yes	No
Quick I/O, Cached Quick I/O	Yes	Yes	Yes
Concurrent I/O	Yes	Yes	Yes

Table 1-5 SFDB features supported in database environments (*continued*)

SFDB feature	DB2	Oracle	Sybase
Storage Checkpoints	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes
Database Storage Checkpoints	No	Yes	No
Database Flashsnap	No	Yes	No
SmartTier for Oracle	No	Yes	No

For the most current information on SFCFSA and single instance Oracle versions supported, see:

<http://www.symantec.com/docs/DOC4039>

Review current documentation for your database to confirm the compatibility of your hardware and software.

Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x8000. When you install the Veritas File System package, `VRTSvxfs`, the `VRTSvxfs` packaging scripts check the values of these variables in the kernel. If the values are less than the required values, the installer gives a warning. You need to set the values and reboot the system.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System.

Use the following commands to check the values of the variables:

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

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

If the values shown are less than 8000, you can expect a reboot after installation.

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

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

```
set lwp_default_stksize=0x8000
set rpcmod:svc_default_stksize=0x8000
```

Disk space requirements

Before installing any of the Veritas Storage Foundation products, confirm that your system has enough free disk space.

Use the "Perform a Preinstallation Check" (P) menu or the `-precheck` option of the product installer to determine whether there is sufficient space.

```
# ./installer -precheck
```

Number of nodes supported

SFCFSHA supports cluster configurations with up to 64 nodes.

Known issues

This section covers the known issues in this release.

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

See ["Documentation"](#) on page 83.

Issues related to installation

This section describes the known issues during installation and upgrade.

Uninstalling the SFCFSHA removes the install log files (2721542)

There is a known issue if you uninstall the SFCFSHA it removes the install log files.

Workaround: Before uninstalling the SFCFSHA, move the install log files to another location to preserve them.

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

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

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

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

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

You can split a cluster into two and reconfigure SFCFSHA on the two clusters using the installer. For example, you can split a cluster *clus1* into *clus1A* and *clus1B*.

However, if you use the installer to reconfigure the SFCFSHA, the installer retains the same cluster UUID of *clus1* in both *clus1A* and *clus1B*. If both *clus1A* and *clus1B* use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: None.

After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon you change the locale of nodes that use it. The vxconfig daemon starts at boot. If you have changed locale, you need to restart the daemon.

Workaround:

Refer to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for the section, "vxconfig daemon recovery."

Adding a node to a cluster fails if you did not set up passwordless ssh

Adding a node to a cluster fails if you did not set up passwordless `ssh` prior to running the `installsfcfsha -addnode` command.

Workaround: You must set up passwordless `ssh`, and then run the `installsfcfsha -addnode` command.

Unable to stop some SFCFSHA processes (2329580)

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

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

For example, use `installsfcfsha` to re-configure SFCFSHA rather than using `installvcs`.

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

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

Workaround:

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

On Solaris 11, uninstalling an IPS package does not always remove the package's SMF services from the SMF database

On Solaris 11, after uninstalling an IPS package, some of the SMF services owned by the package that was uninstalled continue to exist in the SMF database. This issue is due to a defect in the Solaris 11 operating system.

For more information, refer to the Oracle bug ID#7104300 (Removing an external dependency manifest and its manifest the instance is masked).

This issue sometimes prevents you from reinstalling the same package. For example, if you uninstall the `VRTSvxxvm` package, reinstalling the package completes, but some SMF services fail to get added to the repository, which leads to VxVM

drivers not getting loaded, as well as `vxconfigd` and other VxVM daemons not starting.

Workaround: Update your Solaris 11 systems to Solaris 11 Support Repository Update 1.

On Solaris 11, uninstalling an IPS package moves some files to the `$IMAGE_META/lost+found` directory

After uninstalling an IPS package, some of the files related to the package get moved to the `$IMAGE_META/lost+found` directory. The value of `$IMAGE_META` generally is `/var/pkg`. These files are not part of the package, but reside in the directories created by the package.

One example is the log files. In the `VRTSvxvm` package, the CBR log files are kept in `/var/adm/vx`. During the uninstallation of the `VRTSvxvm` package, the following warning message displays:

```
The following unexpected or editable files and directories were
salvaged while executing the requested package operation; they
have been moved to the displayed location in the image:
```

```
var/adm/vx -> /var/pkg/lost+found/var/adm/vx-20120221T114835Z
usr/lib/vxvm/bin -> /var/pkg/lost+found/usr/lib/vxvm/bin-20120221T114835Z
etc/vx/vras -> /var/pkg/lost+found/etc/vx/vras-20120221T114835Z
```

Workaround: There is no workaround for this issue. This is Solaris 11 IPS behaviour.

Veritas Storage Foundation Cluster File System High Availability known issues

This section describes the known issues in this release of Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

AMF driver may panic under very heavy CFS load (2713877)

The AMF driver may panic under very heavy CFS load with the following similar stack trace:

```
> ::status
debugging crash dump vmcore.0 (64-bit) from fssolspr02-v03
operating system: 5.11 11.0 (sun4v)
image uuid: b973f16f-b5d1-c92c-9312-d4b2350460eb
panic message: BAD TRAP: type=31 rp=2a1037af000 addr=18 mmu_fsr=0 occurred
```

```
module "unix" due to a NULL pointer dereference
dump content: kernel pages only
WARNING: /dev/kmem written to, use ::kmemstatus to view log.
> ::stack
mutex_owner_running+0x14(30055026b38, 18, 6001db5ccf8, ffffffff,
30055026b38, 30055026b38)
amf_lock+4(3004c1db1d8, 600174d3478, 2a1037af460, 30054e8e638, 3004e15c541,
30000220008)
amf_s_wait_events+0xe0c(600174d3440, 3004d4e8f98, 3004c1db180, 2a1037af468,
30054e8e638, 0)
amf_reaper_notify+0xbac(3004e15c500, 0, 6001d3ee080, 3004d4e8f98, 600174d3440,
ffbffa68)
amfioc1+0x864(40046106, 28bb0, 0, 3004e15c500, 1, 0)
amf_ioc1+0x54(15600000000, 40046106, 28bb0, 100003, 3004da8fc08, 0)
fop_ioc1+0xc8(3004d052400, 40046106, 28bb0, 100003, 3004da8fc08, 2a1037afacc)
ioc1+0x16c(4, 40046106, 28bb0, 4, ff319e74, ff30e1f4)
syscall_trap32+0xcc(4, 40046106, 28bb0, 0, ff319e74, ff30e1f4)
```

Workaround: There is no workaround for this issue.

CFS commands might hang when run by non-root (2403263)

The CFS commands might hang when run by non-root.

Workaround

To resolve this issue

- ◆ Use `halogin` command to save the authentication information before running any CFS commands on a non-root sessions.

When you run the `halogin` command, VCS stores encrypted authentication information in the user's home directory.

Miscalculated file set usage (2123429)

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

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

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

Workaround

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

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

The cfsmntadm add command may fail with no errors (2169538)

The `cfsmntadm add` command fails, if one host name is a substring of another host name in the list.

Note: VOM is affected by this issue when adding a CFS mount to a cluster that has systems with host names that are substrings of each other.

Workaround

Run the `cfsmntadm` command with the `"all="` option on one of the nodes in the CFS cluster to add the cfsmounts to all nodes.

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

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

Workaround

Create a resource dependency between the various CFSmount resources.

CVMVolDg agent may fail to deport CVM disk group

The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the `CVMDeportOnOffline` attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.

Workaround: If multiple CVMVolDg resources are configured for a shared disk group, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

NFS issues with VxFS Storage Checkpoint (1974020)

NFS clients mounting VxFS Storage Checkpoints that are NFS-exported by SFCFS cluster nodes using a Virtual IP may receive the following error message upon Virtual IP failover:

```
Stale NFS file handle
```

This is a result of major numbers of VxFS Storage Checkpoints not necessarily being the same on all SFCFS cluster nodes.

Workaround: There is no workaround at this time.

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.

Panic due to null pointer de-reference in vx_bmap_lookup() (2582232)

A null pointer dereference in the `vx_bmap_lookup()` call can cause a panic.

Workaround: Resize the file system with the `fsadm` command from the primary node of the cluster.

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

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

Workaround: There is no workaround for this issue.

"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 SFCPSHA cluster that has `CFSMountresources`:

```
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.

Veritas Volume Manager known issues

The following are the Veritas Volume Manager known issues for this release.

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

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

Work-around:

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

Veritas Volume Manager (VxVM) might report false serial split brain under certain scenarios (1834513)

VxVM might detect and report a false serial split brain when all of the following conditions are met:

- One or more arrays that provide the shared storage for the cluster are being powered off
- At the same time when the arrays are being powered off, an operation that requires an internal transaction is initiated (such as VxVM configuration commands)

In such a scenario, disk group import will fail with a split brain error and the `vxsplitlines` output will show 0 or 1 pools.

Workaround:

To recover from this situation

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

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

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

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

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

vxdisk -f init can overwrite some of the public region contents (1190117)

If a disk was initialized by a previous VxVM version or defined with a smaller private region than the new default of 32 MB, then the public region data will be overridden.

Workaround:

Specify explicitly the length of privoffset, puboffset, publen, and privlen while initializing the disk.

The relayout operation fails when there are too many disks in the disk group. (2015135)

The attempted relayout operation on a disk group containing approximately more than 300 LUNs or disks may fail with the following error:

```
Cannot setup space
```

Expanding a LUN to a size greater than 1 TB fails to show correct expanded size (2123677)

This issue occurs when you perform a Dynamic LUN Expansion for a LUN that is smaller than 1 TB and increase the size to greater than 1 Tb. After the expansion, Veritas Volume Manager (VxVM) fails ongoing I/O, and the public region size is reset to original size. After you run the `vxdisk scandisks` command, VxVM does not show the correct expanded size of the LUN. The issue is due to underlying Solaris issues. Refer to Sun Bug Id 6929449 and Sun Bug Id 6912703.

Workaround: There is no workaround for this issue.

Co-existence check might fail for CDS disks

In Veritas Volume Manager (VxVM) 5.1 SP1, VxVM introduces the ability to support Cross-platform Data Sharing (CDS) on disks larger than 1 TB. VxVM uses the SUN VTOC Table to initialize the cdsdisk layout on devices up to 1 TB. VxVM uses the GUID Partition Table (GPT) to initialize the cdsdisk layout on devices larger than 1 TB.

In layouts where SUN VTOC Table is used for initialization (typically, when the disk size has never exceeded 1 TB), the AIX co-existence label can be found at sector 7 and VxVM ID block (also known as HP co-existence label) can be found at sector 16.

In layouts where GPT is used for initialization (typically, when the disk size is currently greater than or had earlier exceeded 1 TB), the AIX co-existence label is placed at sector 55 and VxVM ID block (also known as HP co-existence label) is placed at sector 64. Consequently, AIX utilities would not be able to recognize a cdsdisk initialized using GPT to be a valid VxVM disk. Symantec is working with IBM and third party OEMs to enhance the co-existence check in these utilities.

Workaround: There is no workaround for this issue.

I/O fails on some paths after array connectivity is restored, due to high restore daemon interval (2091619)

If a path loses connectivity to the array, the path is marked as suspected to fail and hence is not used for I/O. After the connectivity is restored, the restore daemon detects that the path is restored when the restore daemon probes the paths. The restore daemon makes the path available for I/O. The restore daemon probes the paths at the interval set with the tunable parameter `dmp_restore_interval`. If you set the `dmp_restore_interval` parameter to a high value, the paths are not available for I/O until the next interval.

Changes in enclosure attributes are not persistent after an upgrade to VxVM 6.0 PR1 (2082414)

The Veritas Volume Manager (VxVM) 6.0 PR1 includes several array names that differ from the array names in releases prior to release 5.1SP1. Therefore, if you upgrade from a previous release to VxVM 6.0 PR1, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0 PR1. Manually reconfigure the enclosure attributes to resolve the issue.

[Table 1-6](#) shows the Hitachi arrays that have new array names.

Table 1-6 Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC
TagmaStoreUSPV	Hitachi_USP-V
TagmaStoreUSPVM	Hitachi_USP-VM
<New Addition>	Hitachi_R700
Hitachi AMS2300 Series arrays	New array names are based on the Model Number 8x. For example, AMS_100, AMS_2100, AMS_2300, AMS_2500, etc.

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0 PR1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

vxsnap addmir command sometimes fails under heavy I/O load (2441283)

The `vxsnap addmir` command sometimes fails under heavy I/O load and produces multiple errors.

Workaround: Rerun the `vxsnap addmir` command.

The "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set (235456)

In Cluster Volume Manager environment, "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set. This can happen on master/slave nodes.

Workaround:

Administrator has to run "vxdisk scandisks" or "vxdisk -o alldgs list" followed by "vxdg listclone" to get all the disks containing "clone_disk" or "udid_mismatch" flag on respective host.

Failback to primary paths does not occur if the node that initiated the failover leaves the cluster (1856723)

When CVM is configured on non-A/A storage, if a node loses access to the storage through all the primary paths, then all the nodes in the cluster switches to the secondary paths. If the node which raised the protocol leaves the cluster and if all the rest of the nodes in the cluster are seeing the primary paths as healthy, then failback to primary paths never happens.

Issues if the storage connectivity to data disks is lost on a CVM slave node while vxconfigd was not running on the node (2562889)

If storage connectivity to data disks is lost on a CVM slave node while vxconfigd was not running on the node, this may result in following issues when vxconfigd comes up on this node:

- The shared disk groups on the disconnected storage are marked as dgdisabled on the slave node only.
- The shared disk groups are available to rest of the cluster nodes but no transactions, such as VxVM configuration changes, are possible on any shared disk group.
- Attempts to deport such shared disk groups will fail.

Work-arounds:

Use one of the following work-arounds:

- Remove the faulty slave node out of CVM cluster, restore storage connectivity, and rejoin the node to the cluster.
- Restart vxconfigd on the CVM master node.

The vxcdsconvert utility is supported only on the master node (2616422)

The vxcdsconvert utility should be run only from the master node, not from the slave nodes of the cluster.

Re-enabling connectivity if the disks are in local failed (lfailed) state (2425977)

In a Cluster Volume Manager (CVM) cluster, you can disable connectivity to the disks at the controller or enclosure level with the `vxdmpadm disable` command. In this case, CVM may place the disks into the `lfailed` state. When you restore connectivity with the `vxdmpadm enable` command, CVM may not automatically clear the `lfailed` state. After enabling the controller or enclosure, you must run disk discovery to clear the locally failed state.

To run disk discovery

- ◆ Run the following command:

```
# vxdisk scandisks
```

Issues with the disk state on the CVM slave node when vxconfigd is restarted on all nodes (2615680)

When a CVM master node and a slave node have lost storage access, and `vxconfigd` is restarted on all nodes, the disk state on the CVM slave node shows as invalid.

Work-around:

To work around this issue

- 1 Restore storage connectivity.
- 2 Deport the disk group.
- 3 Import the disk group.

After changing the preferred path from the array side, the secondary path becomes active (2490012)

For EVA arrays, DMP requires that the prefer bit is static. If the prefer bit is not static, issues like the following may occur. After changing the prefer path of LUN from the array side, and performing a disk discovery (`vxdisk scandisks`) from the host, the secondary path becomes active for the LUN.

Work-around:

To work around this issue

- 1 Set the pref bit for the LUN.
- 2 Perform disk discovery again:

```
# vxdisk scandisks
```

The `vxdisksetup` command fails to initialize disks in `cdsdisk` format for LDOM disks greater than 1 TB (2557072)

The `vxdisksetup` command fails to initialize disks in `cdsdisk` format for LDOM disks greater than 1 TB. This issue is due to an LDOM operating system command which fails when the number of partitions in the GUID partition table (GPT) label is greater than 9. The `cdsdisk` format requires at least 128 partitions to be compatible with Linux systems.

Workaround: There is no workaround for this issue.

Veritas File System known issues

This section describes the known issues in this release of Veritas File System (VxFS).

Taking a FileSnap over NFS multiple times with the same target name can result in the 'File exists' error (2353352)

The "File exists" error occurs as a result of the caching behavior of the NFS client. Because the link operation is successful, the NFS client assumes that a file with the specified target name, such as `file2::snap:vxfs:`, was created. As a result, the NFS client caches a file with this name.

Workaround: Remove the target file after a snapshot is created. This forces the NFS client to remove the name from the cache. For example:

```
# ln file1 file2::snap:vxfs:
# rm file2::snap:vxfs:
```

Enabling delayed allocation on a small file system sometimes disables the file system (2389318)

When you enable delayed allocation on a small file system, such as around 100 MB, the file system can get disabled. In this case, the following error message displays in the system console log:

```
mesg 001: V-2-1: vx_nospace - file_system file system full
(size block extent)
```

Workaround: Use the `vxtunefs` command to turn off delayed allocation for the file system.

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system nears 100% usage even if other volumes have free space (2438368)

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system is nearing 100% usage even if other volumes in the file system have free space.

Workaround: After sufficient space is freed from the volume, delayed allocation automatically resumes.

A mutex contention in vx_worklist_lk() can use up to 100% of a single CPU (2086902)

A mutex contention in the vx_worklist_lk() call can use up to 100% of a single CPU.

Workaround: There is no workaround for this issue.

Upgrading from disk layout Version 8 to 9 on a file system with partitioned directories and Storage Checkpoints can return with a read-only file system error message (2583201)

Upgrading from disk layout Version 8 to 9 on a file system with partitioned directories and Storage Checkpoints can return with a read-only file system error message. The issue with partitioned directories occurs because disk layout Version 9 has a new hash function. The issue with Storage Checkpoints occurs because the Storage Checkpoints are marked as read-only during the upgrade.

Workaround: Before upgrading a VxFS file system with disk layout Version 8 to Version 9, use the following procedure to avoid this error message.

To avoid the system error message

- 1** Disable the partitioned directories feature if the feature is enabled by setting the pdir_enable tunable to 0.
See the vxtunefs(1M) manual page.
- 2** Remove all Storage Checkpoints before the upgrade.
See the fsckptadm(1M) manual page.

Using cross-platform data sharing to convert a file system that has more than 32k nlinks does not update the vx_maxlink and maxlink_enable tunables (2655788)

If you use cross-platform data sharing to convert a file system that has more than 32k nlinks, the conversion process does not update the `vx_maxlink` and `maxlink_enable` tunables on the target file system.

Workaround: After the cross-platform data sharing conversion completes, validate the values of the `vx_maxlink` and `maxlink_enable` tunables. If the file system had more than 32k nlinks before the conversion, ensure that these tunables are updated on the target file system before mounting the file system.

Deduplication can fail with error 110 (2591473)

In some cases, data deduplication fails with a message similar to the following example:

```

Saving      Status      Node           Type           Filesystem
-----
00%         FAILED      node01         MANUAL         /data/fs1
                2011/10/26 01:38:58 End full scan with error

```

In addition, the deduplication log contains an error similar to the following example:

```
2011/10/26 01:35:09 DEDUP_ERROR AddBlock failed. Error = 110
```

These errors indicate that the deduplication process is running low on space and needs more free space to complete.

Workaround: Make more space available on the file system.

vxresize fails while shrinking a file system with the "blocks are currently in use" error (2437138)

The `vxresize shrink` operation may fail when active I/Os are in progress on the file system and the file system is being shrunk to a size closer to its current usage. You see a message similar to the following example:

```

UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink /dev/vx/rdisk/dg1/voll -
blocks are currently in use.
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume
voll, in diskgroup dg1

```

Workaround: Rerun the shrink operation after stopping the I/Os.

System hang when using ls, du and find (2598356)

The system sometimes hangs when using the `ls`, `du`, or `find` commands. The hang occurs in the following stack:

```
schedule_timeout
vx_iget
vx_dirlook
vx_lookup
do_lookup
do_path_lookup
```

Workaround: There is no workaround for this issue.

Expanding a 100% full file system can cause a panic (2599590)

Expanding a 100% full file system can cause a panic with the following stack trace:

```
bad_kern_reference()
$cold_vfault()
vm_hndlr()
bubbledown()
vx_logflush()
vx_log_sync1()
vx_log_sync()
vx_worklist_thread()
kthread_daemon_startup()
```

Workaround: There is no workaround for this issue.

Warning message sometimes appear in the console during system startup (2354829)

During system startup, following messages sometimes appear in system console:

```
WARNING: couldn't allocate SDT table for module vxfs
WARNING: couldn't allocate FBT table for module vxfs
Loading smf(5) service descriptions: 2/2
```

These warnings indicate that the SDT and FBT DTrace probes might not be available for the VxFS module. The VxFS module still loads and works correctly. Dtrace SDT/FBT has limits on the size of module that it can support. Since the VxFS module exceeds the size that Dtrace can support, SDT and FBT Dtrace probes might not work for VxFS.

Workaround: There is no workaround for this issue.

Replication known issues

This section describes the replication known issues in this release of Veritas Storage Foundation Cluster File System High Availability.

In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon

Issue: After upgrading VVR to an IPv6-only environment in 6.0 release, `vradmin` commands may not work when a colon is specified in the RVG, data volume(s) and/or SRL name. It is also possible that after upgrading VVR to an IPv6-only environment, `vradmin createpri` may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

Workaround: Make sure that colons are not specified in the volume, SRL and RVG names in the VVR configuration

`vradmin` commands might fail on non-logowner node after logowner change (1810827)

When VVR is used for replicating shared disk groups in an SFCFS or SFRAC environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg` commands:

```
vradmind not reachable on cluster peer
```

In addition, all other `vradmin` commands (except `vradmin printvol`) fail with the error:

```
"VxVM VVR vradmin ERROR V-5-52-488 RDS has configuration error related to the master and logowner."
```

This is due to a defect in the internal communication sub-system, which will be resolved in a later release.

Workaround: Restart `vradmind` on all the cluster nodes using the following commands:

```
# /lib/svc/method/vras-vradmind.sh stop
# /lib/svc/method/vras-vradmind.sh start
```

While vradmin commands are running, vradmind may temporarily lose heart beats (2162625, 2275444)

This issue may occasionally occur when you use `vradmin` commands to administer VVR. While the `vradmin` commands run, `vradmind` may temporarily lose heartbeats, and the commands terminate with the following error message:

```
VxVM VVR vradmin ERROR V-5-52-803 Lost connection to host host;  
terminating command execution.
```

Workaround:

To resolve this issue

- 1 Depending on the application I/O workload and network environment, uncomment and increase the value of the `IPM_HEARTBEAT_TIMEOUT` variable in the `/etc/vx/vras/vras_env` on all the hosts of the RDS to a higher value. The following example increases the timeout value to 120 seconds.

```
export IPM_HEARTBEAT_TIMEOUT  
IPM_HEARTBEAT_TIMEOUT=120
```

- 2 Restart `vradmind` on all the hosts of the RDS to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following on all the hosts of the RDS:

```
# /etc/init.d/vras-vradmind.sh stop  
# /etc/init.d/vras-vradmind.sh start
```

vradmin syncvol command compatibility with IPv6 addresses (2075307)

The `vradmin syncvol` command does not work with the compressed form of IPv6 addresses. In IPv6 environments, if you run the `vradmin syncvol` command and identify the target host using compressed form of the IPv6 address, the command fails with following error message:

```
# vradmin -s -full syncvol vol1 fe80::221:5eff:fe49:ad10:dg1:vol1  
VxVM VVR vradmin ERROR V-5-52-420 Incorrect format for syncvol.
```

Also, if you run the `vradmin addsec` command and you specify the Secondary host using the compressed IPv6 address, the `vradmin syncvol` command also fails – even if you specify the target as `hostname`.

Workaround: When you use the `vradmin addsec` and `vradmin syncvol` commands, do not specify compressed IPv6 addresses; instead, use hostnames.

RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2054804)

The RVGPrimary agent initiated operation to start replication between the original Primary and the bunker fails during failback – when migrating back to the original Primary after disaster recovery – with the error message:

```
VxVM VVR vxlink ERROR V-5-1-5282 Error getting information from remote host. Internal Error.
```

The issue applies to global clustering with a bunker configuration, where the bunker replication is configured using storage protocol. It occurs when the Primary comes back even before the bunker disk group is imported on the bunker host to initialize the bunker replay by the RVGPrimary agent in the Secondary cluster.

Workaround:

To resolve this issue

- 1 Before failback, make sure that bunker replay is either completed or aborted.
- 2 After failback, deport and import the bunker disk group on the original Primary.
- 3 Try the start replication operation from outside of VCS control.

Bunker replay did not occur when the Application Service Group was configured on some of the systems in the Primary cluster, and ClusterFailoverPolicy is set to "AUTO" (2047724)

The time that it takes for a global cluster to fail over an application service group can sometimes be smaller than the time that it takes for VVR to detect the configuration change associated with the primary fault. This can occur in a bunkered, globally clustered configuration when the value of the `ClusterFailoverPolicy` attribute is `Auto` and the `AppGroup` is configured on a subset of nodes of the primary cluster.

This causes the RVGPrimary online at the failover site to fail. The following messages appear in the VCS engine log:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname could not be imported on bunker host hostname. Operation failed with error 256 and message VxVM VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote server unreachable... Timestamp VCS ERROR V-16-2-13066 (hostname) Agent is calling clean for resource(RVGPrimary) because the resource is not up even after online completed.
```

Workaround:

To resolve this issue

- ◆ When the configuration includes a bunker node, set the value of the `OnlineRetryLimit` attribute of the `RVGPrimary` resource to a non-zero value.

The RVGPrimary agent may fail to bring the application service group online on the new Primary site because of a previous primary-elect operation not being run or not completing successfully (2043831)

In a primary-elect configuration, the `RVGPrimary` agent may fail to bring the application service groups online on the new Primary site, due to the existence of previously-created instant snapshots. This may happen if you do not run the `ElectPrimary` command to elect the new Primary or if the previous `ElectPrimary` command did not complete successfully.

Workaround: Destroy the instant snapshots manually using the `vrxvrg -g dg -P snap_prefix snapdestroy rvg` command. Clear the application service group and bring it back online manually.

A snapshot volume created on the Secondary, containing a VxFS file system may not mount in read-write mode and performing a read-write mount of the VxFS file systems on the new Primary after a global clustering site failover may fail (1558257)

Issue 1:

When the `vradmin ibc` command is used to take a snapshot of a replicated data volume containing a VxFS file system on the Secondary, mounting the snapshot volume in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/snapshot_volume  
is corrupted. needs checking
```

This happens because the file system may not be quiesced before running the `vradmin ibc` command and therefore, the snapshot volume containing the file system may not be fully consistent.

Issue 2:

After a global clustering site failover, mounting a replicated data volume containing a VxFS file system on the new Primary site in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/data_volume  
is corrupted. needs checking
```

This usually happens because the file system was not quiesced on the original Primary site prior to the global clustering site failover and therefore, the file systems on the new Primary site may not be fully consistent.

Workaround: The following workarounds resolve these issues.

For issue 1, run the `fsck` command on the snapshot volume on the Secondary, to restore the consistency of the file system residing on the snapshot.

For example:

```
# fsck -F vxfs /dev/vx/dsk/dg/snapshot_volume
```

For issue 2, run the `fsck` command on the replicated data volumes on the new Primary site, to restore the consistency of the file system residing on the data volume.

For example:

```
# fsck -F vxfs /dev/vx/dsk/dg/data_volume
```

vxassist relay layout removes the DCM (2162522)

If you perform a relay layout that adds a column to a striped volume that has a DCM, the DCM is removed. There is no message indicating that this has happened. To replace the DCM, enter the following:

```
#vxassist -g diskgroup addlog vol logtype=dcm
```

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

```
# vxrlink -g diskgroup status rlink
```

- 3 Stop the affected RVG. Enter the following:

```
# vxrvg -g diskgroup stop rvg
```

- 4 Disassociate the volumes from the RVG. Enter the following:

```
# vxvol -g diskgroup dis vol
```

- 5 Resize the volumes. In this example, the volume is increased to 10 GB. Enter the following:

```
# vxassist -g diskgroup growto vol 10G
```

- 6 Associate the data volumes to the RVG. Enter the following:

```
# vxvol -g diskgroup assoc rvg vol
```

- 7 Start the RVG. Enter the following:

```
# vxrvg -g diskgroup start rvg
```

- 8 Resume or start the applications.

Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (2478684)

Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (DCM), even if you have enough disk space.

Workaround: Add a LUN to the diskgroup before creating the primary diskgroup.

verifydata operation fails when replicating between versions 5.1 and 6.0 (2360713)

When replicating in a cross-version VVR environment consisting of hosts running Storage Foundation 5.1 and hosts running Storage Foundation 6.0, the `vradm` `verifydata` command fails with the following error:

```
VxVM VVR vxrsync ERROR V-5-52-2222 [from host]: VxVM in.vxrsyncd  
ERROR V-5-36-2125 Server volume access error during [assign volids]  
volume path: [/dev/vx/dsk/dg/snapshot_volume] reason: [this could be  
because a target volume is disabled or an rlink associated with a  
target volume is not detached during sync operation].
```

Workaround: There are two workarounds for this issue.

- Upgrade the hosts running Storage Foundation 5.1 to Storage Foundation 5.1SP1 or later and re-run the `vradmin verifydata` command.
- Follow the offline verification procedure in the "Verifying the data on the Secondary" section of the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*. This process requires ensuring that the secondary is up-to-date, pausing replication, and running the `vradmin syncrvg` command with the `-verify` option.

Replication hang when VVR logowner is on CVM slave node (2405943)

When VVR is used for asynchronous replication in shared disk group environment, one of the nodes of the cluster at the primary site is chosen as the logowner. When the logowner node is on a node which is a slave node for the underlying CVM cluster, in the presence of heavy I/O from a node that is not the logowner, it is possible to get into a replication hang. This is due to an internal defect which will be fixed in later releases.

Workaround: Enable the PreOnline trigger of the RVGLogOwner agent so that the VVR logowner will always reside on the CVM master node. For the detailed procedure, refer to the RVGLogowner agent notes section in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Cannot relayout data volumes in an RVG from concat to striped-mirror (2162537)

This issue occurs when you try a relayout operation on a data volume which is associated to an RVG, and the target layout is a striped-mirror.

Workaround:

To relayout a data volume in an RVG from concat to striped-mirror

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

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

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

```
# vxvol -g diskgroup dis vol
```

- 5 Relayout the volumes to striped-mirror. Enter the following:

```
# vxassist -g diskgroup relayout vol layout=stripe-mirror
```
- 6 Associate the data volumes to the RVG. Enter the following:

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

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

vradmin functionality may not work after a master switch operation (2163712)

In certain situations, if you switch the master role, `vradmin` functionality may not work. The following message displays:

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

Workaround:

To restore vradmin functionality after a master switch operation

- 1 Restart `vradmind` on all cluster nodes. Enter the following:

```
# /etc/init.d/vras-vradmind.sh stop  
# /etc/init.d/vras-vradmind.sh start
```
- 2 Re-enter the command that failed.

Issues related to LLT

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

LLT port stats sometimes shows recvcnt larger than recvbytes (1788315)

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 (1809827)

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.

Issues related to GAB

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 deinitd 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.

Cluster panics during reconfiguration (2590413)

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

Issues related to I/O fencing

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

Delay in rebooting Solaris 10 nodes due to vxfen service timeout issues (1897449)

When you reboot the nodes using the `shutdown -i6 -g0 -y` command, the following error messages may appear:

```
svc:/system/vxfen:default:Method or service exit
timed out. Killing contract 142
svc:/system/vxfen:default:Method "/lib/svc/method/vxfen stop"
failed due to signal Kill.
```

This error occurs because the vxfen client is still active when VCS attempts to stop I/O fencing. As a result, the vxfen stop service times out and delays the system reboot.

Workaround: Perform the following steps to avoid this vxfen stop service timeout error.

To avoid the vxfen stop service timeout error

- 1 Stop VCS. On any node in the cluster, run the following command:

```
# hastop -all
```

- 2 Reboot the systems:

```
# shutdown -i6 -g0 -y
```

CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

See the *Veritas Storage Foundation Cluster File System High Availability 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 `vx.fenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

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

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.

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

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

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

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@node1,
domaintype vx; not allowing action
```

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

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

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

The `vxferswap` utility runs the `vxfenconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxferswap` 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, `vxferswap` 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 `vxfsnswap` 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 `vxfsnswap` 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.

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

When you configure fencing in customized mode and do not 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 `vxcpsserv` process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: No workaround.

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

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

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

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

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

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

Workaround: You can use any of the following options:

- Before running the `configure_cps.pl` utility, change the default shell for root user to either KSH or bash.
- Perform the following steps after running the `configure_cps.pl` utility on each node of the cluster:

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

```
# rm -rf /var/VRTSvc/vcsauth/data/CPSEVER
```

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

```
# ln -s path_of_CP_server_credential_directory \
/var/VRTSvc/vcsauth/data/CPSEVER
```

- Start the CPSSG service group:

```
# hagrps -online CPSSG -any
```

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

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSats package is not removed from the system, the `cpsadm` command loads the old security libraries present on the system. As the installer runs the `cpsadm` command on the CP server to add or upgrade the SFCFSHA cluster (application cluster), the installer also fails.

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

- Rename `cpsadm` to `cpsadmbin`.

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

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

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

- Provide the following permissions to the new file:

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

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

Secondary-host related operations may fail in a pure IPv6 environment (2625708)

In a pure IPv6 configuration, using the `--secondary_host` option with `vxsfdm` or the equivalent `-t` option for `dbed_vmchecksnap` may fail with the following error message if name resolution is configured to use `/etc/hosts` rather than DNS:

```
SFDB vxsfadm ERROR V-81-0161 The given secondary host <hostname>  
is not reachable.
```

Action: Specify a different secondary host, or remove the `-t` option if the secondary host is not running at this time.

Workaround

Use one of the following workarounds:

- Configure name resolution to use DNS instead of a hosts file, if a DNS server is available.
- Do not use the `--secondary_host` option.
This option is not mandatory for off-host operations. Off-host operations are fully supported without specifying this option.

dbed_update or vxsfadm may fail in a pure IPv6 environment (2619958)

In a pure IPv6 environment, the `dbed_update` or the `vxsfadm` command may fail with the following error:

```
SFDB vxsfadm ERROR V-81-0685  
Oracle instance SIDB does not seem to be running
```

Workaround

Ensure that an IPv4 address is configured for localhost. This IP does not need to be reachable from other hosts.

To configure an IPv4 localhost address

- ◆ Run the following command as root:

```
# ifconfig lo0 plumb
```

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

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

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

Workaround

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

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

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

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

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

The error messages are unclear.

Workaround

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

FlashSnap validate reports snapshot unsplittable (2534422)

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

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

Workaround

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

Attempt to use SmartTier commands fails (2332973)

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

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

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

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

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

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

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

Workaround

You can use one of the following workarounds:

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

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

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

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

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

- BALANCE
- CHECKPOINT
- METADATA

Workaround

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

User authentication fails (2579929)

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

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

Reattempting the operation fails with the following error message:

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

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

Workaround

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

To set umask to a less strict value

- ◆ Use the command:

```
# umask 022
```

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

To resolve the problem if you have already done authentication setup

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

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

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

```
# chmod o+r /etc/vx/vxdbed/admin.properties
```

```
# chmod o+rx /var/vx/vxdba/auth/users
```

```
# find /opt/VRTSdbed/at-broker -type d -exec chmod o+rx {} \;
```

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

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

Workaround

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

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

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

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

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

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

Workaround

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

Clone command fails if PFILE entries have their values spread across multiple lines (1764885)

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

Workaround

There is no workaround for this issue.

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

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

Workaround: Set MonitorOption attribute for Oracle resource to 0.

Software limitations

This section covers the software limitations of this release.

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

See “[Documentation](#)” on page 83.

Veritas Storage Foundation Cluster File System High Availability software limitations

The following are software limitations in this release of Veritas Storage Foundation Cluster File System High Availability.

cfsmntadm command does not verify the mount options (2078634)

You must confirm the mount options are correct which are then passed to the `cfsmntadm` command. If the mount options are not correct, the mount fails and the CFSMount resource will not come online. You can check the VCS engine log file for any mount failure messages.

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

When all nodes present in the SFCFSHA 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 `vxfcntlclearpre` utility.

For more information on the `vxfcntlclearpre` utility, see the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

Veritas File System software limitations

The following are software limitations in the 6.0 PR1 release of Veritas Storage Foundation.

Recommended limit of number of files in a directory

To maximize VxFS performance, do not exceed 100,000 files in the same directory. Use multiple directories instead.

The vxlist command cannot correctly display numbers greater than or equal to 1 EB

The `vxlist` command and all of the other commands that use the same library as the `vxlist` command cannot correctly display numbers greater than or equal to 1 EB.

Limitations with delayed allocation for extending writes feature

The following limitations apply to the delayed allocation for extending writes feature:

- In the cases where the file data must be written to disk immediately, delayed allocation is disabled on that file. Examples of such cases include Direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.
- Delayed allocation is not supported on memory mapped files.
- Delayed allocation is not supported with BSD quotas. When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.
- Delayed allocation is not supported for shared mounts in a cluster file system.

FlashBackup in NetBackup 7.1 and prior does not support disk layout Version 8 and 9

The FlashBackup feature of NetBackup 7.1 or prior does not support a VxFS file system with disk layout Version 8 or 9.

Compressed files that are backed up using NetBackup 7.1 or prior become uncompressed when you restore the files

The NetBackup 7.1 release and prior does not support the file compression feature. If you back up compressed files using NetBackup 7.1 or a prior release, the files become uncompressed when you restore the files.

Veritas Volume Manager software limitations

The following are software limitations in this release of Veritas Volume Manager.

Converting a multi-pathed disk

When converting a multi-pathed disk that is smaller than 1 TB from a VTOC label to an EFI label, you must issue the `format -e` command for each path. For example,

if a node has two paths, `c1t2d0s2` and `c2t2d0s2`, you must run the `format -e` command on each of the two paths.

DMP settings for NetApp storage attached environment

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

[Table 1-7](#) describes the DMP tunable parameters and the new values.

Table 1-7 DMP settings for NetApp storage attached environment

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

The change is persistent across reboots.

To change the tunable parameters

- 1 Issue the following commands:

```
# vxddmpadm settune dmp_restore_internal=60

# vxddmpadm settune dmp_path_age=120
```

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

```
# vxddmpadm gettune dmp_restore_internal

# vxddmpadm gettune dmp_path_age
```

DMP and EMC PowerPath coexistence

EMC PowerPath (EMCPP) currently does not claim support for Solaris 11. Therefore, DMP and EMCPP coexistence is not supported in Solaris 11.

You cannot encapsulate the boot disk in Solaris 11

Encapsulation of the boot disk in Solaris 11 is not supported.

Replication software limitations

The following are replication software limitations in this release of Veritas Storage Foundation Cluster File System High Availability.

Replication in a shared environment

Currently, replication support is limited to 8-node cluster applications.

IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.
- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.
- A replication configuration in which an IPv6 address is specified for the `local_host` attribute of a primary RLINK and an IPv4 address is specified for the `remote_host` attribute of the same RLINK.
- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.
- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 6.0 and the prior major releases of Storage Foundation (5.1 and 5.1SP1). Replication between versions is supported for disk group versions 150, 160, and 170 only. Both the Primary and Secondary hosts must be using a supported disk group version.

Limitations related to I/O fencing

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

Preferred fencing limitation when VxFEN activates RACER node re-election

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

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *Veritas Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

In a cluster using SCSI-3 based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on both the data disks and coordinator disks. In a cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordination points and data disks to prevent possible difficulties with subsequent cluster startup.

Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordination points and data disks. Depending on the order of reboot and subsequent startup events, the cluster may warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown -r command on one node at a time and wait for each node to complete shutdown.

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

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

Veritas Storage Foundation for Databases (SFDB) tools software limitations

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

Oracle Data Guard in an Oracle RAC environment

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

Upgrading if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 6.0 PR1: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 6.0 PR1.

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.

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.

Documentation errata

The following sections cover additions or corrections for Document version: 6.0PR1.3 of the product documentation. These additions or corrections may be included in later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).

See the corresponding Release Notes for documentation errata related to that component or product.

See [“Documentation”](#) on page 83.

See [“About Symantec Operations Readiness Tools”](#) on page 9.

Veritas Storage Foundation Cluster File System High Availability Administrator's Guide

The following errata applies to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

"Requirements" section in the "Common Internet File System" chapter

Replace the list of requirements with the following list:

- For Solaris 9:
 - CIFS requires Samba version 3.2 or later and Oracle patch 112837-21.
 - If you are running Samba 3.2.x, CIFS requires you to add the Samba library path to the `LD_LIBRARY_PATH` variable in the `/opt/VRTSvcs/bin/vcsenv` file.
- For Solaris 10:
 - CIFS requires Samba version 3.2 or later.
 - If you are running Samba 3.2.x, CIFS requires you to add the Samba library path to the `LD_LIBRARY_PATH` variable in the `/opt/VRTSvcs/bin/vcsenv` file.
 - Prior knowledge of Samba is a prerequisite.

"VxFS Version 9 disk layout" section in the "Disk layout" appendix

The following text should be deleted:

The Version 8 disk layout supports group quotas.

See "About quota files on Veritas File System" on page x.

Documentation

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

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

<http://sort.symantec.com/documents>

Documentation set

[Table 1-8](#) lists the documentation for Veritas Storage Foundation Cluster File System High Availability.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability documentation

Document title	File name
<i>Veritas Storage Foundation Cluster File System High Availability Release Notes</i>	sfdfs_notes_60pr1_sol.pdf
<i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i>	sfdfs_install_60pr1_sol.pdf
<i>Veritas Storage Foundation Cluster File System High Availability Administrator's Guide</i>	sfdfs_admin_60pr1_sol.pdf

[Table 1-9](#) lists the documents for Veritas Cluster Server.

Table 1-9 Veritas Cluster Server documentation

Title	File name
<i>Veritas Cluster Server Installation Guide</i>	vcs_install_60pr1_sol.pdf
<i>Veritas Cluster Server Release Notes</i>	vcs_notes_60pr1_sol.pdf
<i>Veritas Cluster Server Administrator's Guide</i>	vcs_admin_60pr1_sol.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents_60pr1_sol.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev_60_unix.pdf
<i>Veritas Cluster Server Application Note: Dynamic Reconfiguration for Oracle Servers</i>	vcs_dynamic_reconfig_60pr1_sol.pdf
<i>Veritas Cluster Server Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_agent_60pr1_sol.pdf
<i>Veritas Cluster Server Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_agent_60pr1_sol.pdf
<i>Veritas Cluster Server Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_agent_60pr1_sol.pdf

[Table 1-10](#) lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-10 Veritas Storage Foundation and High Availability Solutions products documentation

Document title	File name
<i>Veritas Storage Foundation and High Availability Solutions Solutions Guide</i>	sfha_solutions_60pr1_sol.pdf
<i>Veritas Storage Foundation and High Availability Solutions Virtualization Guide</i>	sfha_virtualization_60pr1_sol.pdf
<i>Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide</i>	sf_replication_admin_60pr1_sol.pdf

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

<http://sort.symantec.com/documents>

Manual pages

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

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

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

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

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

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

See the `man(1)` manual page.

