

Veritas Storage Foundation™ Release Notes

AIX

6.0

Veritas Storage Foundation 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

Document version: 6.0.4

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 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 Release Notes

This document includes the following topics:

- [About this document](#)
- [Component product release notes](#)
- [About Veritas Storage Foundation](#)
- [About Symantec Operations Readiness Tools](#)
- [Important release information](#)
- [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 (Storage Foundation) version 6.0 for AIX. Review this entire document before you install or upgrade Storage Foundation.

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

This is Document version: 6.0.4 of the *Veritas Storage Foundation 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.

About Veritas Storage Foundation

Veritas Storage Foundation by Symantec includes Veritas File System (VxFS) and Veritas Volume Manager (VxVM).

Veritas File System is a high performance journaling file system that provides easy management and quick-recovery for applications. Veritas File System delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity.

Veritas Volume Manager removes the physical limitations of disk storage. You can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Veritas Volume Manager also provides easy-to-use, online storage management tools to reduce downtime.

VxFS and VxVM are included in all Veritas Storage Foundation products. If you have purchased a Veritas Storage Foundation product, VxFS and VxVM are installed and updated as part of that product. Do not install or update them as individual components.

Veritas Storage Foundation includes the dynamic multi-pathing functionality.

The Veritas Replicator option, which replicates data to remote locations over an IP network, can also be licensed with this product.

Before you install the product, read the *Veritas Storage Foundation Release Notes*.

To install the product, follow the instructions in the *Veritas Storage Foundation Installation Guide*.

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/DDLs), and high availability agents from a central repository.■ Identify and mitigate system and environmental risks.■ Display descriptions and solutions for hundreds of Symantec error codes. |
| Improve efficiency | <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

This section lists the changes in Veritas Storage Foundation 6.0.

Changes related to Veritas Storage Foundation (SF)

Veritas Storage Foundation includes the following changes in 6.0:

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`
- `vx dg free`
- `vx disk list`
- `vx disk -o thin list`
- `vx disk -o thin, fssize list`

- `vxmpadm iostat show`
- `vxedquota`
- `vxmemstat`
- `vxprint`
- `vxquot`
- `vxquota`
- `vxrepquota`
- `vxstat`
- `vxtune`

See the manual pages for more information.

Displaying Storage Foundation information with `vxlist`

The `vxlist` command is a new display command that provides a consolidated view of the Storage Foundation 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.

Discovering renamed devices on AIX

Starting with AIX 6.1TL6, AIX provides a feature to rename a device using the `rendev` command. You can now specify user-defined names instead of the traditional `hdisk` name.

Veritas Dynamic Multi-Pathing (DMP) now can discover the renamed devices. DMP supports device renaming for both enclosure-based naming (EBN) and

operating system naming (OSN). Before renaming a device, remove the DMP node from VxVM/DMP control.

You can use the `vxddmpadm` command to enable and disable the renamed path.

The following features are not supported with renamed devices:

- Enabling rootability
- Migrating LVM to VxVM using the `vxconvert` command.
- Hot relocation

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 `vxddmpadm config` commands.

See the `vxddmpadm(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.

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.

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 support for LVM rootvg using DMP ODM definitions

This release of Dynamic Multi-Pathing supports LVM root volume group (rootvg) using DMP ODM definitions, which enables LVM commands such as `lspv` and `lsvg -p rootvg` to show the DMP device name rather than the native (hdisk) names.

When `dmp_native_support` feature is turned on, DMP supports online configuration for the LVM rootvg, that is, adding a DMP device to the LVM rootvg using the `extendvg` command, or removing a DMP device using the `reducevg` command. In previous releases, this required a reboot.

DMP enhancements

The following DMP enhancements have been made in this release:

- The `vxdmadm enable` command and the `vxdmadm 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 `vxddmpadm stat error` command and the `vxddmpadm stat restored` command are deprecated.
To see status for the restore tasks, use the `vxddmpadm 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:

```
vxddmpadm exclude dmp  
vxddmpadm include dmp
```

`vxddiskadm: 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 events` source is added to show the status of the `vxesd` daemon
- `vxscsiinq` diagnostic utility is enhanced to take hexadecimal page numbers as arguments.

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*.

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.

Veritas Volume Manager co-existence with Oracle Automatic Storage Management (ASM) disks

ASM disks are the disks used by Oracle Automatic Storage Management software. Veritas Volume Manager (VxVM) co-exists with Oracle ASM disks, by recognizing the disks as the type Oracle ASM. VxVM protects ASM disks from any operations that may overwrite the disk. VxVM classifies and displays the ASM disks as ASM format disks. You cannot initialize an ASM disk, or perform any VxVM operations that may overwrite the disk.

For more information about VxVM co-existence with ASM disks, see the *Veritas Storage Foundation Administrator's Guide*.

Changing VxVM tunables

The `vxtune` command is used to display or modify the values of Veritas Volume Manager tunable parameters. In this release, the `vxtune` command is extended and enhanced. The `vxtune` command has the following new functionality:

- manages an extended list of Veritas Volume Manager tunable parameters, including Veritas Volume Replicator and Cluster Volume Manager tunable parameters.
- provides a template format for tuning parameters. The template feature enables you to export the list of tunable parameters into a file, modify the values as necessary, then reload the tunables with an import command.
- enhanced command output. The output now displays the current value, the default value, and whether a reboot is required for the new value to take effect. Optionally, the output displays a description of the tunable parameters.
- makes the tunable values persistent across reboots.
- categorizes the tunable parameters by VxVM component. Specify the component to list or export the tunable parameters in that category. The components are the following:
 - basevm
Basic core VxVM functionality.
 - fmr
FlashSnap functionality.
 - cvm
Cluster Volume Manager.
 - vvr
Veritas Volume Replicator.

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.

Online Migration of native LVM volumes to VxVM volumes

In this release, Veritas Volume Manager (VxVM) provides a feature to migrate volumes under native LVM control to VxVM volumes, with a limited application downtime.

This migrates source LVM volume data to target VxVM volumes on new storage, with the flexibility of different storage and layouts. Once the migration is set up, the application can be resumed, while data synchronization from source LVM to target VxVM volumes continues in the background.

The migration configuration is set up such that the application does not require immediate reconfiguration to the new VxVM device paths.

You can also choose the point of committing the migration, when data synchronization is complete for all required volumes. In case of errors, it provides a way to abort the migration and safely revert to the original LVM configuration.

This feature is also integrated with VCS to provide online migration in a VCS HA environment. During the migration process, VCS monitors and maintains high availability of the updated configuration.

A new CLI `vxmigadm` is provided, to administer online migration.

For more details, refer to *Veritas™ Storage Foundation and High Availability Solutions Solutions Guide*.

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 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
- vxdmpadm
- vxddladm
- vxvol
- vxcdsconvert
- vxresize
- vxdctl
- 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.

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:

- Data deduplication

See the *Administrator's Guide*.

Data deduplication

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*.

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 Administrator's Guide*.

Partitioned directories

Normally, a large volume of parallel threads performing access and updates on a directory that commonly exist in a 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*.

vxfscnvert can upgrade additional Veritas File System disk layout versions

The `vxfscnvert` command can upgrade the VxFS disk layout Version 4.

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 replication

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

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.

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*.

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.

Enhancements to Cached ODM Advisor (dbed_codm_adm)

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

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

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

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

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

Single CLI for different point-in-time copy operations

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

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

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

Support for file-level snapshots for database cloning

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

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

Enhanced authentication support

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

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

Licensing changes in the SFHA Solutions 6.0 release

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

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

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

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

The following products are included in this release:

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

- Storage Foundation Enterprise HA/DR for Oracle RAC
- Storage Foundation Enterprise for Sybase ASE CE
- Storage Foundation Enterprise HA/DR for Sybase CE

HA: High Availability

HA/DR: High Availability and Disaster Recovery

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

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

Changes related to installation and upgrades

The product installer includes the following changes in 6.0.

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 filesets 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.

Allow Response files to change tuning parameters

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

See the *Installation Guide* for more information.

Packaging updates

The following lists the package changes in this release.

- New `VRTSsfcp160` fileset for product installer scripts
The `VRTSsfcp160` fileset is introduced in this release. The `VRTSsfcp160` fileset contains the installer scripts and libraries that the installer uses to install, configure and upgrade Veritas products.
- New `VRTSfsadv` fileset for product data deduplication
The `VRTSfsadv` fileset is introduced in this release. The `VRTSfsadv` fileset 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.

[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> .
<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> .

Table 1-2 Deprecated documents (*continued*)

Deprecated documents	Notes
<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 File System Programmer's Reference Guide</i>

No longer supported

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

- Several documents are deprecated in this release.
See [“Changes related to product documentation”](#) on page 27.
- Disk layout Version 4 is no longer supported. You cannot create nor mount a file system with disk layout Version 4. 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

System requirements

The following topics describe the system requirements for this release:

Supported AIX operating systems

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

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

Table 1-4 Supported AIX operating systems

Operating systems	Levels	Chipsets
AIX 7.1	TL0 or TL1	Any chipset that the operating system supports
AIX 6.1	TL5	Power 7, Power 6, or earlier

For Dynamic Multi-Pathing (DMP), Veritas Storage Foundation (SF), and Veritas Storage Foundation and High Availability Solutions (SFHA), install IBM APAR for AIX 6.1 TL6 and TL7, or AIX 7.1 TL0 and TL1. Contact IBM to get the necessary APAR for your level.

AIX 7.1 support for virtual processors

Veritas Storage Foundation supports up to 1024 virtual processors on AIX 7.1.

Hardware compatibility list (HCL)

The hardware compatibility list contains information about supported hardware and is updated regularly. 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.

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

<http://www.symantec.com/docs/TECH170013>

For information on specific High Availability setup requirements, see the *Veritas Cluster Server Installation Guide*.

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
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 Storage Foundation 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 Storage Foundation memory requirements

Symantec recommends 2 GB of memory over the minimum requirement for the operating system.

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 64.

Issues related to installation

This section describes the known issues during installation and upgrade.

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

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

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

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

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

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

To unfreeze the service groups manually

- 1 List all the frozen service groups

```
# hagrps -list Frozen=1
```

- 2 Unfreeze all the frozen service groups:

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

EULA changes (2161557)

The locations for all EULAs have changed.

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

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

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

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

NetBackup 6.5 or older version is installed on a VxFS file system (2056282)

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to Veritas Storage Foundation (SF) 6.0, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (*/usr/opensv*), then while upgrading to SF 6.0, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure filesets *VRTSspbx*, *VRTSat*, and *VRTSicisco*. This causes NetBackup to stop working.

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the */usr/opensv/netbackup/bin/version* file and */usr/opensv/netbackup/version* file to the */tmp* directory. If you have clustered NetBackup installed, you must also copy the */usr/opensv/netbackup/bin/cluster/NBU_RSP* file to the */tmp* directory. After you unmount the NetBackup file system, manually copy these two version files from */tmp* to their original directories. If you have clustered NetBackup installed, you must also copy the */usr/opensv/netbackup/bin/cluster/NBU_RSP* file from */tmp* to its original directory.

If the *version* files' directories do not exist, create the directories:

```
# mkdir -p /usr/opensv/netbackup/bin  
# mkdir -p /usr/opensv/netbackup/bin
```

Run the installer to finish the upgrade process. After upgrade process completes, remove the two version files and their directories.

If your system is already affected by this issue, then you must manually install the `VRTSspbx`, `VRTSat`, and `VRTSicisco` filesets after the upgrade process completes.

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

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

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

The VRTSaclib fileset is deprecated (2032052)

The VRTSaclib fileset is deprecated. For installation, uninstallation, and upgrades, note the following:

- Fresh installs: Do not install VRTSaclib.
- Upgrade: Uninstall old VRTSaclib and install new VRTSaclib.
- Uninstall: Ignore VRTSaclib.

The VRTSvxvm fileset fails to install on a few cluster nodes because the template file is corrupted (2348780)

The installer debug log displays the failure of the `errupdate` command as following:
`errupdate -f /usr/lpp/VRTSvxvm/inst_root/VRTSvxvm.err`. The `errupdate` command gets invoked through `/usr/lib/instl/install` by the operating system. The command also fails for the `VRTSvxfs`, `VRTSglm`, and `VRTSgms` packages.

The `errupdate` command generally creates a `*.undo.err` file to remove entries from the Error Record Template Repository in case of failed installation or cleanup. However, in this case the `*.undo.err` file does not get generated as the `errupdate` command fails. Also, it is not possible to manually remove entries from the Error Record Template Repository in order to undo the changes made by the failed installation, because the file is corrupted.

Workaround: Save a copy of the `/var/adm/ras/errtmpl` and `/etc/trcfmt` files before you install the product. Replace `/var/adm/ras/errtmpl` and `/etc/trcfmt` files with the ones that you saved, when the installation fails because the template file is corrupted. Uninstall all the packages you installed and reinstall.

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

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

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

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

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

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

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

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

Workaround:

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

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

When you uninstall CommandCentral Storage Managed Host from a system where Veritas Storage Foundation 6.0 is installed, SF 6.0 reconfiguration or uninstallation fails (2631486)

On a system where Veritas Storage Foundation (SF) 6.0 is installed, if you uninstall CommandCentral Storage (CCS) Managed Host (MH) using the installer script

from the CCS media, the installer script removes the contents of `/opt/VRTSperl`. As a result, SF 6.0 reconfiguration or uninstallation using `/opt/VRTS/install/install_sf_product_name` or `/opt/VRTS/install/uninstall_sf_product_name` fails, because the installer script removed the contents of `/opt/VRTSperl`.

Workaround: To uninstall CCS MH from a system where SF 6.0 is installed, before you perform the uninstallation, perform the procedure in the following CCS TechNote:

<http://www.symantec.com/business/support/index?page=content&id=HOWTO36496>

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

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

Workaround:

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

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

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

Workaround: Do one of the following:

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

Veritas Storage Foundation known issues

This section describes the known issues in this release of Veritas Storage Foundation (SF).

AT Server crashes when authenticating unixpwd user multiple times (1705860)

There is a known issue in the AIX kernel code that causes 'getgrent_r' function to corrupt the heap. This issue is present in AIX 5.3 and AIX 6.1 Refer to IBM's Web site for more information:

<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52585>

AT uses `getgrent_r` function to get the groups of the authenticated user.

IBM has released the fix as a patch to `fileset bos.rte.libc`. There are different patches available for different version of `bos.rte.libc`. You need to check the version of `bos.rte.libc` (For example: `lslpp -l | grep bos.rte.libc`) and apply the appropriate IBM patch:

- For version 6.1.3.1:
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52959>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52959/>
- For version 6.1.2.4:
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52720>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52720/>
- For version 6.1.2.5 :
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52975>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52975/>

There are IBM patches for only certain version of `bos.rte.libc` that are available. If your system has a different `bos.rte.libc` version, you may have to upgrade to a higher version where the fix is available. If your version is not available, you may have to contact IBM.

Not all the objects are visible in the VOM GUI (1821803)

After upgrading SF stack from 5.0MP3RP2 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the SFM GUI.

Workaround:

To resolve this known issue

- ◆ On each manage host where `VRTSsfmh 2.1` is installed, run:

```
# /opt/VRTSsfmh/adm/dclisetup.sh -U
```

A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880622)

A volume's placement class tags are not visible in the Veritas Enterprise Administrator (VEA) GUI when you are creating a SmartTier placement policy if you do not tag the volume with the placement classes prior to constructing a volume set for the volume.

Workaround: To see the placement class tags in the VEA GUI, you must tag the volumes prior to constructing the volume set. If you already constructed the volume set before tagging the volumes, restart `vxsvc` to make the tags visible in the GUI.

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.

The cluster may hang if a node goes down (1835718)

The cluster may hang if a node goes down while one array is disabled or offline in a mirror=enclosure configuration.

This may occur, if a node panics or loses power while one array of a mirror=enclosure configuration is offline or disabled, then the cluster, fencing, I/O loads, and VxVM transactions hang.

Workaround: There is no workaround for this issue.

vxconvert failures if PowerPath disks are formatted as simple disks (857504)

If a PowerPath disk is formatted as a simple disk (a foreign device), then the `vxconvert` utility may fail during conversion of LVM to VxVM. To view the format of the disk, use the `vxdisk list` command. This issue may also occur if the

/etc/vx/darecs file contains an hdiskpower disk entry. This entry may be present if PowerPath disks were configured as foreign disks in Storage Foundation 4.0, and the entry was not changed after subsequent upgrades.

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

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

```
Cannot setup space
```

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.

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

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


```
# vxassist -g dgname maxsize
```

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

Workaround:

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

Hardware paths for operating system paths have changed in DMP 6.0 (2410716)

In DMP 6.0, the hardware paths for operating system paths have changed. After upgrading to DMP 6.0, path attributes are reset to the default values. You must reconfigure any path-level attributes that were defined in the `/etc/vx/dmppolicy.info` file.

Workaround:**To configure path-level attributes**

- 1 Remove the path entries from the `/etc/vx/dmppolicy.info` file.
- 2 Reset the path attributes.

Upgrading Storage Foundation to version 6.0 marks vSCSI disks as cloned disks (2434444)

This issue is seen when you upgrade from a previous version of Storage Foundation which has vSCSI disks included in a disk group. After upgrading Storage Foundation to 6.0, the vSCSI disks that were included in a disk group are marked as cloned disks.

Workaround:

Use the following procedure to clear the clone disk flag.

To clear the clone disk flag

- 1 Remove the vSCSI devices that are in error state (`ibm_vscsi#_#`) using the following command:

```
# vxdisk rm device_name
```

- 2 Deport the disk group.

```
# vxvg deport dg_name
```

- 3 Re-import the disk group with a new udid.

```
# vxdg -o updateid import dg_name
```

- 4 Display the devices that are part of the disk group.

```
# vxdisk -g dg_name list
```

- 5 Clear the clone_disk tag from these devices.

```
# vxdisk set device_name clone=off
```

The vxsnap print command shows incorrect value for percentage dirty (2360780)

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the `%dirty`. In Storage Foundation 6.0, if this command is run while the volumes are online and being actively used, the shown `%dirty` may lag from actual percentage dirty for instant snap data change object (DCO) volumes. That is, the command output may show less `%dirty` than actual.

Encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path names are not the same (2607706)

The encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path name are not the same.

For example:

Dmpnode:sdh

Paths: sda sdb

Work-around:

Before running the encapsulation command (`vxencap`), run the following command:

```
# vxddladm assign names
```

Recovery and rollback to original configuration may not succeed if the system reboots while the online migration setup is in partial state (2611423)

During online migration from LVM to VxVM volumes, if there is a system reboot when the migration setup is in partial state, that is, the start operation has not completed successfully, then the recover and abort operations might not be able to recover and rollback the configuration.

Workaround: This needs manual intervention for cleanup, depending on the state, to restore the original configuration.

Array controller reboot on CLARiiON storage in failovermode 1 or 4 on AIX 6.1 (2418875)

On an AIX 6.1 host, when you reboot the array controller on a CLARiiON array that is configured in failovermode 1 or 4, the dmpnode may go in failed state, resulting in I/O failures on the LUN.

Veritas File System known issues

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

Cannot use some commands from inside an automounted Storage Checkpoint (2490709)

If your current work directory is inside an automounted Storage Checkpoint, for example `/mnt1/.checkpoint/clone1`, some commands display the following error:

```
can't find current directory
```

This issue is verified with the following commands:

- `cp -r`
- `du`

However, this issue might occur with other commands.

Workaround: Run the command from a different directory.

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.

Performance on a VxFS file system can be slower than on a JFS file system (2511432)

There are two causes for the performance degradation:

- There are unnecessary page faults that set the write permissions on mapped pages.
- The entire file flushes.

The issue of unnecessary page faults has been fixed, which has improved the performance considerably. However, the performance on a VxFS file system can still be slower than on a JFS file sometimes due to the entire file flushing.

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.

You are unable to unmount the NFS exported file system on the server if you run the fsmigadm command on the client (2355258)

Unmounting the NFS-exported file system on the server fails with the "Device busy" error when you use the `fsmigadm` command on the NFS client.

Workaround: Unexport the file system prior to unmounting.

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.

Replication known issues

This section describes the replication known issues in this release of Veritas Storage Foundation.

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 vxrlink 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 bunkerdgrname 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 `vxrvg -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 -V 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 -V vxfs /dev/vx/dsk/dg/data_volume
```

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:

```
vradmin 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 `vradmin` on all the cluster nodes using the following commands:

```
# /lib/svc/method/vras-vradmin.sh stop
# /lib/svc/method/vras-vradmin.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
```

vxassist relayout removes the DCM (2162522)

If you perform a relayout 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 `vradmin 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 `vradm 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 `vradm 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.

Veritas Storage Foundation for Databases (SFDB) tools known issues

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

Database Storage Checkpoints created by using `dbed_ckptcreate` may not be visible after upgrading to 6.0 (2626248)

After upgrading from a 5.0 release to 6.0, the Database Storage Checkpoints created earlier using `dbed_ckptcreate` may not be migrated.

Workaround

Perform the following steps to make the old Database Storage Checkpoints visible.

To resolve the issue

- 1 Remove the new repository.
 - Examine the contents of the `/var/vx/vxdba/rep_locfile` to determine the location of the 6.0 repository.

- Remove the `.sfae` directory specified as the `location` attribute.
- 2 Remove the repository location file: `/var/vx/vxdba/rep_loc`.
- 3 Create a symlink `/var/vx/vxdba/<SID>/sfdb_rept` pointing to the `.sfdb_rept` directory created in the same location as the `.sfae` directory removed earlier.

```
$ ln -s <location>/sfdb_rept /var/vx/vxdba/<SID>/sfdb_rept
```

This step creates a symlink to the old repository.

- 4 Import repository data by running the `dbed_update` command.

This step imports the data from the old repository.

The old Database Storage Checkpoints are now visible.

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

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

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

Workaround

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

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

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

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

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

The error messages are unclear.

Workaround

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

FlashSnap validate reports snapshot unsplittable (2534422)

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

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

Workaround

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

Attempt to use SmartTier commands fails (2332973)

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

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

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

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

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

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

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

Workaround

You can use one of the following workarounds:

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

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

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

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

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

- BALANCE
- CHECKPOINT
- METADATA

Workaround

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

User authentication fails (2579929)

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

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

Reattempting the operation fails with the following error message:

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

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

Workaround

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

To set `umask` to a less strict value

- ◆ Use the command:

```
# umask 022
```

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

To resolve the problem if you have already done authentication setup

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

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

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

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

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

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

Workaround

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

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

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

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

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

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

Workaround

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

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

The `sfua_rept_migrate` command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 or 5.0MP3 to SFHA or SF for Oracle RAC 6.0.

When upgrading from Storage Foundation version 5.0 or 5.0MP3 to Storage Foundation 6.0 the `S*vxdbms3` startup script is renamed to `NO_S*vxdbms3`. The `S*vxdbms3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbms3` startup script and gives the error message:

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

Workaround

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

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

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

Workaround

There is no workaround for this issue.

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 64.

Veritas File System software limitations

The following are software limitations in the 6.0 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 shell cannot handle 64-bit inode numbers inside the .checkpoint directory when uniqueino is enabled

Due to a limitation with the AIX operating system, the shell cannot handle the 64-bit inode numbers inside the `.checkpoint` directory when the `uniqueino` mount option is enabled. Some shell functions such as auto-complete and globs, for example `rm *`, do not function properly in the `.checkpoint` directory. This also affects 32-bit applications that try to read the contents of the `.checkpoint` directory or any of its subdirectories. This does not affect any 64-bit applications.

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.

Veritas Volume Manager software limitations

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

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, set the following DMP tunables:

Table 1-6

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

The change is persistent across reboots.

To change the tunable parameters

- 1 Issue the following commands:

```
# 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
```

Limitation with device renaming on AIX 6.1TL6

If you rename an operating system (OS) path with the `rendev` command on AIX 6.1TL6, the operation might remove the paths from DMP control. DMP cannot discover these paths.

DMP support in AIX virtualization environment (2138060)

DMP does not support exporting paths to the same LUN through both vSCSI and NPIV interfaces.

DMP treats the same LUN seen through vSCSI and NPIV interfaces as two separate LUNs, because the behavior of the LUN at the VIOC level is different due to the intermediate SCSI interface at the VIOS level for vSCSI devices.

Replication software limitations

The following are replication software limitations in this release of Veritas Storage Foundation.

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.

Veritas Storage Foundation for Databases (SFDB) tools software limitations

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

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: 6.0.4 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 64.

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

Veritas Storage Foundation Administrator's Guide

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

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

Replace the beginning of the sentence that begins with, "For the 64-bit kernel versions of AIX 5.2 and 5.3..." with "For the 64-bit kernel versions of AIX 6.1 and 7.1...".

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 1-7](#) lists the documentation for Veritas Storage Foundation.

Table 1-7 Veritas Storage Foundation documentation

Document title	File name
<i>Veritas Storage Foundation Release Notes</i>	sf_notes_60_aix.pdf
<i>Veritas Storage Foundation Installation Guide</i>	sf_install_60_aix.pdf
<i>Veritas Storage Foundation Administrator's Guide</i>	sf_admin_60_aix.pdf
<i>Veritas Storage Foundation: Storage and Availability Management for Oracle Databases</i>	sf_adv_ora_60_aix.pdf
<i>Veritas File System Programmer's Reference Guide</i>	vxfs_ref_60_aix.pdf

[Table 1-8](#) lists the documentation for Veritas Storage Foundation and High Availability Solutions products.

Table 1-8 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_60_aix.pdf
<i>Veritas Storage Foundation and High Availability Solutions Virtualization Guide</i>	sfha_virtualization_60_aix.pdf
<i>Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide</i>	sf_replication_admin_60_aix.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.

Index

C

Configuring DMP
using templates 12

D

DMP
tuning with templates 12

K

Known issues
SF 36
VxFS 43

N

No longer supported 29
no longer supported 30

T

Tuning DMP
using templates 12

V

VxFS Limitations
software 60