

Veritas Storage Foundation™ and High Availability Solutions 6.0

What's new in this release

AIX, HP-UX, Linux, Solaris

What's New In This Release

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: 6.0.1

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec, the Symantec 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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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 Web site 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 Web site.

<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

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>

What's new in this release

This document includes the following topics:

- [Introduction](#)
- [Featured product enhancements in 6.0](#)
- [Changes related to installation and upgrades](#)
- [Changes related to Veritas Dynamic Multi-Pathing \(DMP\)](#)
- [Changes related to Veritas Storage Foundation \(SF\)](#)
- [Changes related to Veritas Cluster Server \(VCS\)](#)
- [Changes related to Veritas Storage Foundation and High Availability \(SFHA\)](#)
- [Changes related to Veritas Storage Foundation Cluster File System High Availability \(SFCFSHA\)](#)
- [Changes related to Symantec VirtualStore \(SVS\)](#)
- [Changes related to Veritas Storage Foundation for Oracle RAC \(SF Oracle RAC\)](#)
- [Changes related to Veritas Storage Foundation for Sybase ASE CE \(SF Sybase CE\)](#)
- [Licensing changes in the SFHA Solutions 6.0 release](#)
- [Changes related to product documentation](#)

Introduction

This document covers the major new features, enhancements, and changes that are introduced in 6.0 for the Veritas Storage Foundation and High Availability

Solutions products. See the *Release Notes* for your product for a list of all changes in 6.0.

Note: The new features and enhancements listed in this document apply to AIX, HP-UX, Linux, and Solaris unless mentioned otherwise.

Featured product enhancements in 6.0

Veritas Storage Foundation and High Availability Solutions products 6.0 include the following major new features:

- Resilient Elastic Storage using:
 - Storage Foundation file compression
See “[File compression](#)” on page 8.
 - Storage Foundation data deduplication
See “[Data deduplication](#)” on page 9.

- Resilient Business Services using:
 - Virtual Business Services (Veritas Cluster Server)

- Virtualization support:
 - Symantec ApplicationHA 6.0
See “[Support for Symantec ApplicationHA 6.0](#)” on page 10.
 - Support for KVM on Linux
See “[Support for Kernel-based Virtual Machines \(KVM\) on Linux](#)” on page 10.

Video presentations for these features are available on the Symantec Connect Web site at:

<http://www.symantec.com/connect/6>

File compression

Compressing files reduces the space used, while retaining the accessibility of the files and being transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files. Reads cause data to be uncompressed in memory, only; the on-disk copy of the file remains compressed. In contrast, after a write, the new data is uncompressed on disk.

Only user data is compressible. You cannot compress Veritas File System (VxFS) metadata.

After you compress a file, the inode number does not change, and file descriptors opened before the compression are still valid after the compression.

Compression is a property of a file. Thus, if you compress all files in a directory, for example, any files that you later copy into that directory do not automatically get compressed. You can compress the new files at any time by compressing the files in the directory again.

To compress files, you must have VxFS file systems with disk layout Version 9 or later.

This feature is supported on Linux and Solaris.

See the *Administrator's Guide*.

Data deduplication

The data deduplication feature eliminates duplicate blocks used by your data by comparing blocks across the file system. When the data deduplication feature finds a duplicate block, it removes the space used and instead creates a pointer to the common block. If you change the duplicate file, thus making the files no longer share the same block, then that changed block is saved to disk instead of the pointer. You can perform post-process periodic deduplication in a file system to eliminate duplicate data without any continuous cost in CPU overhead. You can verify whether data is duplicated on demand, and then efficiently and securely eliminate the duplicates. The deduplication process performs the following tasks:

- Scans the file system for changes
- Fingerprints the data
- Identifies duplicates
- Eliminates duplicates after verifying the duplicates

The amount of space savings that you get from deduplicating depends on your data. Deduplicating different data gives different space savings.

Deduplication requires an Enterprise license.

See the *Administrator's Guide*.

About Virtual Business Services

Virtual Business Services provide continuous high availability and reduce frequency and duration of service disruptions for multi-tier business applications running on heterogeneous operating systems and virtualization technologies. A Virtual Business Service represents the multi-tier application as a single consolidated entity and builds on the high availability and disaster recovery provided for the individual tiers by Symantec products such as Veritas Cluster Server and Symantec ApplicationHA. Additionally, a Virtual Business Service can also represent all the assets used by the service such as arrays, hosts, and file

systems, though they are not migrated between server tiers. A Virtual Business Service provides a single consolidated entity that represents a multi-tier business service in its entirety. Application components that are managed by Veritas Cluster Server or Symantec ApplicationHA can be actively managed through a Virtual Business Service.

You can configure and manage Virtual Business Services created in Veritas Operations Manager by using the Veritas Operations Manager Add-on for Virtual Business Services. Besides providing all the functionality that was earlier available through the Business Entity Operations Add-on, the Add-on for Virtual Business Services provides the additional ability to configure fault dependencies between the components of the multi-tier application.

Note: All the Application Entities that were created using Veritas Operations Manager Business Entity Operations Add-on versions 3.1 and 4.0 are available as Virtual Business Services after you deploy the Add-on for Virtual Business Services in Veritas Operations Manager 4.1. Veritas Operations Manager 4.1 is a prerequisite for running Virtual Business Services.

Support for Symantec ApplicationHA 6.0

Storage Foundation and High Availability Solutions products 6.0 support Symantec ApplicationHA 6.0.

Symantec ApplicationHA 6.0 provides high availability for business-critical applications through application visibility and control in VMware, Oracle VM Server for SPARC (formerly known as LDOMs), IBM Logical Partition (LPAR), and Red Hat Kernel-based Virtual Machines (KVM) environments.

Key benefits include the following:

- Full visibility and control over applications with the ability to start, stop, and monitor applications running inside virtual machines.
- Specialized Application Maintenance mode, in which ApplicationHA allows you to intentionally take an application out of its purview for maintenance or troubleshooting.

See the Symantec ApplicationHA product documentation for more information.

Support for Kernel-based Virtual Machines (KVM) on Linux

Veritas Storage Foundation High Availability Solutions (SFHA Solutions) products support various configurations in the Kernel-based Virtual Machine (KVM) environment. Veritas Storage Foundation High Availability Solutions 6.0 is certified on the Red Hat Enterprise Linux (RHEL) 6.1 distribution.

Storage Foundation and High Availability Solutions provide the following functionality for KVM guest virtual machines:

- Storage visibility
- Storage management
- Replication support
- High availability
- Cluster failover

The configurations profiled in the table below are the minimum required to achieve the storage and availability objectives listed. You can mix and match the use of SFHA Solutions products as needed to achieve the desired level of storage visibility, management, replication support, availability, and cluster failover for your KVM hosts and guest virtual machines.

Table 1-1 Storage Foundation and High Availability Solutions features in guest and host

Objective	Recommended SFHA Solutions product configuration
Storage visibility for KVM guest virtual machines	Dynamic Multi-Pathing (DMP) in the KVM guest virtual machines
Storage visibility for KVM hosts	DMP in the KVM hosts
Storage management features and replication support for KVM guest virtual machines	Storage Foundation (SF) in the KVM guest virtual machines
Advanced storage management features and replication support for KVM hosts	Storage Foundation Cluster File System (SFCFSHA) in the KVM hosts
End-to-end storage visibility in KVM hosts and guest virtual machines	DMP in the KVM host and guest virtual machines
Storage management features and replication support in the KVM guest virtual machines and storage visibility in the KVM host	DMP in the KVM host and SF in the KVM guest virtual machines
Application monitoring and availability for KVM guest virtual machines	Symantec ApplicationHA in the KVM guest virtual machines
Virtual machine monitoring and failover for KVM hosts	Veritas Cluster Server (VCS) in the KVM hosts

Table 1-1 Storage Foundation and High Availability Solutions features in guest and host (*continued*)

Objective	Recommended SFHA Solutions product configuration
Application failover for KVM guest virtual machines	VCS in the KVM guest virtual machines
Application availability and virtual machine availability	Veritas Application HA in the KVM guest virtual machine and VCS in the KVM host
Application failover across KVM guest virtual machines and physical hosts	VCS in a cluster across KVM guest virtual machines and KVM physical host machines

Each configuration has specific advantages and limitations.

See [“Support for Symantec ApplicationHA 6.0”](#) on page 10.

For implementation information:

See the *Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide for Linux*.

Changes related to installation and upgrades

The product installer includes the following changes in 6.0.

Support for product upgrades using the Network Installation Manager Alternate Disk Migration utility on AIX

You can now use the Network Installation Manager Alternate Disk Migration (nimadm) utility to upgrade the operating system and the product.

See the *Installation Guide* for more information.

Support for product installation using Ignite-UX on HP-UX

The installer now supports the `-ignite` option to create Software Distributor (SD) bundles. You can use these SDs to install the Storage Foundation and High Availability Suite of products using HP-UX Ignite installation tool.

See the *Installation Guide* for more information.

Support for product installation using yum on Linux

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

See the *Installation Guide* for more information.

Support for product deployment using Flash Archives on Solaris 10

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

See the *Installation Guide* for more information.

Performing a rolling upgrade

You can now use the script- and Web-based installers to perform rolling upgrades on clusters. Rolling upgrades require less downtime than other upgrade procedures.

In 6.0, the rolling upgrade support is added for the products on HP-UX. This support is already available on other platforms since 5.1 SP1.

Changes related to Veritas Dynamic Multi-Pathing (DMP)

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

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 `vxdmptadm` 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 *Storage Foundation and High Availability Solutions Tuning Guide*.

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*.

Changes related to Veritas Storage Foundation (SF)

Veritas Storage Foundation includes the following changes in 6.0:

Changes related to Veritas Volume Manager

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.

In 6.0, this support is added to the products on AIX, Linux, and Solaris. This feature is already supported on HP-UX.

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. The size of the DRL (Dirty region logging) map does not depend on volume size and can be configured using `drlmapsz` while creating DCO volume.

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 feature is supported on AIX and HP-UX.

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.

Online migration can utilize SmartMove capability for the data synchronization, if VxFS is configured on source LVM volumes.

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*.

Changes related to Veritas 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:

- File compression
(on Linux and Solaris)
- Data deduplication
- File replication
(on Linux)

See the *Administrator's Guide*.

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 *product Administrator's Guide*.

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

Partitioned directories

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

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.

This feature is supported on AIX and Solaris.

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

Parallel Direct I/O on Linux

On VxFS, each `iovec` is performed synchronously for the `readv(2)` call and `writev(2)` call. For both `readv(2)` and `writev(2)`, the Single Unix Specification states, "The `readv/writev()` function shall always fill an area completely before proceeding to the next." However, for direct I/O, Linux ignores this requirement and submits a number of `iovecs` in parallel before waiting for completion. In this release, VxFS now performs parallel direct I/O for both reads and writes, which improves VxFS performance. This support for parallel direct I/O can be enabled by setting the VxFS module load tunable `vx_parallel_dio`.

See the *Administrator's Guide*.

vxfsconvert can upgrade additional Veritas File System disk layout versions

On AIX and Linux:

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

On HP-UX and Solaris:

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

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.

FileSnap creation over Network File System

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

See the *Administrator's Guide*.

Changes related to replication

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

File Level Replication on Linux

Veritas File Replicator (VFR) supports file-level replication of application data, tracks all updates to the File System and periodically replicates these updates at the end of a configured time interval. VFR leverages Veritas File System (VxFS) data deduplication and will not replicate data that is already on the destination. VFR also supports VxFS compression and compressed files will be replicated as such. In addition, VFR also supports reversible data transfers. VFR is available as an option to Storage Foundation, included in the Veritas Replicator (new name for Veritas Volume Replicator) license and also in Symantec VirtualStore 6.0.

See the *Storage Foundation and High Availability Solutions Replication Administrator's Guide* for more details.

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*.

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*.

ASM to VxFS migration

You can use the ASM2VxFS utility to perform in-place migration of Oracle Database server running on Oracle's Automatic Storage Management (ASM) to Veritas File System (VxFS).

This utility is currently tested with Storage Foundation 5.1SP1. Its coverage will be extended to include 6.0 in future. At this time, technical support will be provided to customers looking to migrate from ASM to VxFS.

See *Veritas Storage Foundation: In-place Migration of Oracle Database from ASM to Veritas File System*.

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*.

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.

Displaying product information with vxlist

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

Changes related to Veritas Cluster Server (VCS)

The following sections contain changes related to VCS kernel components such as LLT, GAB, and I/O fencing, and clusters in secure mode.

For more information on changes related to VCS, see the *Veritas Cluster Server Release Notes*.

Changes to the VCS engine

IMF is enabled by default and has new agent support

The Intelligent Monitoring Framework (IMF) functionality is enabled by default for all agents that can leverage IMF.

The following agents are IMF-aware in VCS 6.0:

- Zone
- WPAR
- DB2udb (provides only PRON IMF support)
- Sybase
- SybaseBk

Enable and disable IMF by automated script for agents

VCS provides the `/opt/VRTSvcs/bin/haimfconfig` script to enable and disable IMF for agents when VCS is either running or stopped. You can use this script to disable IMF for all IMF-aware agents, including bundled agents, enterprise agents, and custom agents. You must run the script once on each node of the cluster.

Note: The automated script restarts the agent, if it is running on the current node, to enable the IMF after confirming with the user.

VCS support of CPU interrupt disablement on AIX

HAD process runs at a high priority on the system; however, interrupts are always run at the highest priority. In case lot of interrupts are generated on the system and the CPU where HAD process is running gets to service them, HAD does not get a chance to heartbeat with GAB and perform its task of making applications highly available.

To overcome this issue, VCS allows you to run HAD on a specific processor and mask off all interrupts on that processor. you can configure HAD to run on a specific processor by setting the `CPUBinding` attribute. For more information on CPU binding of HAD, refer to the *Administrator's Guide*.

Support multiple children in service group dependency

A service group can have dependency on multiple child service groups. All child dependencies must be satisfied for the parent service group to go online. The dependency types which are not supported in a multiple child configuration are "online local hard" and "offline local". The `-propagate` option cannot be used if the dependency tree contains global and/or remote dependency.

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

Single command line option to online or offline service groups in a service group dependency

If you have a parent service group in a service group dependency, you can online the entire dependency tree bottom-up with a single command, without having to online each group manually. For online operation, the command starts to online the service groups from the lowermost service group in dependency tree. Similarly, If you have a child service group in a service group dependency, you can offline the entire dependency tree top-down with a single command, without having to offline each group manually. For offline operation, the command offlines the service group from the top of the dependency tree.

You can use the following commands to online or(and) offline the service groups respectively:

- `hagr -online -propagate <grp_name> -sys <sys_name>`
- `hagr -offline -propagate <grp_name> -sys <sys_name>`

Note: The `-propagate` option cannot be used if the dependency tree contains global and/or remote dependency.

Using the `-propagate` option with `hagr -online`, all required child service groups are automatically brought online by VCS. Similarly using the `-propagate` option with `hagr -offline`, all the required parent groups are automatically brought offline by VCS.

Ability to send notifications to a wider audience

You can configure users, other than the owners of resources, resource types, service groups, systems, or clusters, as recipients of notifications about events related to a resource, resource type, service group, system, or cluster.

Use the following attributes to configure recipients of notifications:

- `ResourceRecipients`

- TypeRecipients
- GroupRecipients
- SystemRecipients
- ClusterRecipients

The registered recipients get notifications about the events that have a severity level that is equal to or greater than the level specified in the attribute. For more information, see *Veritas Cluster Server Administrator's Guide*.

VCS now supports clusters with mixed operating system versions of AIX

All nodes in the cluster must run the same VCS version. Each node in the cluster may run a different version of the AIX operating system, as long as the operating system version is supported by the VCS in the cluster.

Caution: While using mixed operating systems (OS) for an extended period of time, you may see issues with migrating OS-specific applications between major versions of the OS.

Changes related to VCS triggers

This release includes the following changes related to VCS triggers:

- VCS can execute trigger scripts specific to a service group and/or resource. Thus, trigger scripts for multiple objects need not be merged into single trigger script.
- You can enable the `nofailover`, `postonline` and `postoffline` triggers for each system by using the `TriggersEnabled` attribute.
- You can execute multiple scripts for a trigger. The trigger scripts must be installed inside the trigger directory using the `T<num>` nomenclature. VCS executes the trigger scripts in `T<num>` order.
For example: If the `preonline` directory contains the scripts `T00preonline`, `T01preonline`, `T02preonline`, then the script `T00preonline` is executed first, then `T01preonline` and finally, `T02preonline` is executed.
- The agent restarts a faulted resource if the `RestartLimit` is set. Whenever the agent restarts a resource, VCS invokes the `resrestart` trigger if the `TriggerResRestart` attribute is set to 1 or if `RESRESTART` is specified in the `TriggersEnabled` attribute. Otherwise, VCS invokes the `resstatechange` trigger. See the *Veritas Cluster Server Administrator's Guide* for more information.

Caution: Use of `resstatechange` to indicate restart is being deprecated. In later releases, you must use only the `resrestart` trigger to indicate restarting of resources.

Changes to LLT

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

- **Faster detection of LLT link failures**
 LLT now has the capability to detect the LLT link failures immediately using the operating system's link-down notification feature.
 This feature is supported on HP-UX and Solaris.
- **LLT now supports VLAN tagging (IEEE 802.1Q).**

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

Changes to GAB

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

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

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

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

See the Administrator's Guide for more details.

GAB can now recognize clients with names in addition to ports

When kernel clients initialize GAB API, they can now define a client name string. GAB now adds a client name which enables GAB to track the client even before GAB port is registered. GAB also passes the client name information to LLT when

registering the LLT port. The `lltstat -p` command also displays the GAB client names when providing the status details of the ports in use.

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

Changes to I/O fencing

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

Support for Non-SCSI3 fencing

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

In 6.0, this support is added to VCS, SFHA, SFCFSHA, and SVS on Solaris. This feature is already supported on AIX, HP-UX, and Linux.

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

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

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

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

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

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

See the *Administrator's Guide* for more details.

Support for multiple virtual IP addresses in CP servers

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

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

Support for Quorum agent in CP servers

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

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

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

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

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

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

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

See the *Administrator's Guide* for more details.

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

Installer support to migrate between fencing configurations in an online cluster

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

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

See the *Administrator's Guide* for more details.

Changes to VCS bundled agents

This section describes changes to the bundled agents for VCS.

See the *Veritas Cluster Server Administrator's Guide* and *Veritas Cluster Server Bundled Agents Reference Guide* for more information.

Support for Windows DNS server

The DNS agent now supports Windows DNS server in its configuration. A new attribute `UseGSSAPI` is added to DNS agent configuration for this functionality.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on using this attribute and additional requirements for configuring DNS agent with Windows DNS server.

RVGPrimary agent starts replication from new primary post role change to the other secondaries in the RDS

After a successful migration or takeover of a Secondary RVG, the RVGPrimary agent ensures to automatically start the replication from the new Primary to the additional Secondary(s) that exists in the RDS, if any.

Refer to the *SF Replication Administrator's Guide* for information about how the RVGPrimary agent works in a multiple secondary setup.

Improved LVMVG agent performance on AIX

The LVMVG monitor is now C++ based and uses LVM API. Therefore, the CPU utilization is low even with large number of LVMVG resources.

LVMVG agent on AIX supports synchronization of ODM entries

LVMVG agent on AIX now supports synchronization of the ODM entries of all the nodes in a cluster, if the disk configuration of the volume group is changed on a node.

If physical volumes are added, deleted or replaced in a volume group, some or all of the ODM entries on other nodes become stale. This may cause the volume group resource online failure on those nodes. The updated disk information can now be propagated to all nodes in the cluster using “updatepv” action entry point. See the *Bundled Agent Reference Guide* for more information about the updatepv action entry point.

First failure data capture for VCS agent entry point times out

To debug entry point timeout, a new attribute AdvDbg is introduced. The attribute helps VCS agents capture some information like process stack, process tree and core on entry point timeout. This is helpful in troubleshooting entry point timeout scenarios.

See the *Veritas Cluster Server Agent Developer's Guide* for more information.

Changes to database agents

Changes to DB2 agent

- The product agent for DB2 now supports intelligent resource monitoring for online Db2 processes in PRON mode for non-MPP and MPP configuration mode.
- Prior to product 6.0 release, in case of partition mobility from source node to the target node having high speed inter-connect/switch configuration, the switch name entry would not get updated in the db2nodes.cfg configuration file. In product 6.0 release, the DB2 agent ensures that the switch name gets updated correctly in the configuration file.
- Added IMF support for DB2 agent: The DB2 agents now leverage the IMF functionality for instantaneous detection of resource state change. This also reduces the CPU footprint of the agent by significantly reducing the periodic monitoring for resource states.

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

Changes to the Oracle agent

- The VCS agent for Oracle has two new attributes: DBName and ManagedBy.

Changes related to Veritas Cluster Server (VCS)

- The VCS agent for Oracle has additional options to the StartUpOpt and ShutDownOpt attributes:
 - The StartUpOpt attribute introduces SRVCTLSTART_RO as additional startup options.
 - The ShutDownOpt attribute introduces SRVCTLSTOP_TRANSACTION, SRVCTLSTOP_ABORT, and SRVCTLSTOP_IMMEDIATE as additional shut down options.
- The VCS agent for Oracle introduces support for policy managed database.
- The VCS agent for Oracle ASM instance introduces the following additional Startup options:
 - STARTUP_MOUNT
 - STARTUP_OPEN
 - SRVCTLSTART_MOUNT
 - SRVCTLSTART_OPEN
- The VCS agent for Oracle ASM instance introduces the following additional Shutdown option:
 - SRVCTLSTOP
- With Oracle version 11.2.0.2, the Oracle agent for VCS supports Startup and Shutdown options that use `srvctl` utility for Oracle restart configuration.

Changes to the Sybase agent

The Veritas Cluster Server agent for Sybase includes the following new or enhanced features:

- The VCS agents for Sybase and SybaseBk now support intelligent resource monitoring.
- The VCS agents for Sybase now supports intelligent resource monitoring.
- Intelligent monitoring framework (IMF) is enabled by default in product 6.0 release. The `haimgconfig` script can be used to enable/disable IMF for Sybase and SybaseBk agents. The Sybase agent now leverages the IMF functionality for instantaneous detection of resource state change. This also reduces the CPU footprint of the agent by significantly reducing the periodic monitoring for resource states.
See the [Administrator's Guide](#) and [Agent for Sybase Installation and Configuration Guide](#) for more information.
- The Sybase agent introduces the following new attributes:

- Quorum_dev
- interfaces_File
- ShutdownWaitLimit (default value 60)
- DelayAfterOnline (default value 10)
- DelayAfterOffline (default value 2)
- The SybaseBk agent introduces the following new attribute:
 - interfaces_File
- The default value of ToleranceLimit attribute is set to 1 (one) for Sybase agent.
- The DetailMonitor attribute is deprecated in VCS 6.0. Instead, LevelTwoMonitorFreq attribute of Sybase agent may be used. The default value of LevelTwoMonitorFreq attribute is 0 (zero).
- The long pathname limitation for \$SYBASE is resolved.
- With VCS 6.0 release using VCS Cluster Manager (Java Console), Sybase and SybaseBk agents encrypt the password by default. Sybase and SybaseBk agents supports both plain text and encrypted password. If required, the plain text value can be specified for agent attributes using the command line or by editing the configuration file.

With VCS 6.0 release using VCS Cluster Manager (Java Console), Sybase agents encrypt the password by default. Sybase agent supports both plain text and encrypted password. If required, the plain text value can be specified for agent attributes using the command line or by editing the configuration file.
- Sybase agent uses new timeout option during shutdown of Sybase dataserver used instead of "shutdown with nowait". For Sybase ASE Enterprise edition the timeout option for shutdown command is supported for versions 12.5.4 and 15.0.2 onwards.

Sybase agent uses new timeout option during shutdown of Sybase dataserver used instead of "shutdown with nowait". For Sybase ASE Cluster Edition, the timeout option for shutdown command is supported from versions 15.5 ESD #1 onwards.
- For a Solaris Zone environment, Sybase and SybaseBk agents are updated to check the ContainerInfo attribute for physical to virtual failover support.

Attributes introduced in VCS 6.0

Following are the list of attributes introduced in VCS 6.0:

DNS agent attributes

Changes related to Veritas Cluster Server (VCS)

- **UseGSSAPI:** Use this attribute if the DNS server that is configured is a Windows DNS server. The agent uses “-g” option with `nsupdate` command if this attribute is set to 1.
- **RefreshInterval:** This attribute represents the time interval in seconds after which the DNS agent attempts to refresh the resource records (RRs) on the DNS servers.
- **CleanRRKeys:** Use this attribute to direct the online agent function to clean up all the existing DNS records for the configured keys before adding new records. The default value (0) disables this behavior.

Cluster-level attribute

- **SystemRebootAction:** Use this attribute to determine whether the frozen service groups are ignored on system reboot. If the value of `SystemRebootAction` is `IgnoreFrozenGroup`, product ignores the service groups that are frozen (`TFrozen` and `Frozen`) and takes the remaining service groups offline. If the Value of `SystemRebootAction` is "", product tries to take all service groups offline. See the *Veritas Cluster Server Administrator's Guide* for more information.
- **EnableVMAutoDiscovery:** Enables or disables auto discovery of virtual machines. By default, auto discovery of virtual machines is disabled.
- **SystemRebootAction:** Determines whether frozen service groups are ignored on system reboot.

Service group attributes

- **OnlineClearParent:** When this attribute is enabled for a service group and the service group comes online or is detected online, VCS clears the faults on all online type parent groups, such as online local, online global, and online remote.
- **ProPCV:** Indicates whether the service group is proactively prevented from concurrency violation for ProPCV-enabled resources.
- **TriggerPath:** Enables you to customize the trigger path.
- **TriggerResRestart:** Determines whether or not to invoke the restart trigger if resource restarts.
- **TriggersEnabled:** Determines if a specific trigger is enabled on a node or not.

Resource type attributes:

- **AdvDbg:** Enables activation of advanced debugging.

Zpool agent attributes

- **DeviceDir:** Specifies the directories that the `zpool import` command must search for devices or files.

- **FailMode:** Controls the system behavior in the event of a catastrophic pool failure. The value of this attribute is used as the failmode option while importing the ZFS storage pool. The possible values are **wait**, **continue**, or **panic**.
- **ForceOpt:** Invokes the following commands again with the `-f` option if this attribute is enabled (if the value is set to 1) and if the following commands fail:
 - `zpool export` in the offline entry point
 - `zpool export` in the clean entry point
 - `zpool import` in the online entry point
- **ForceRecoverOpt:** Invokes the `zpool import` command again with `-F` option when this attribute is enabled (if its value is set to 1) and if the `zpool import` command fails.

LDom agent attributes:

- **Memory:** Specifies the amount of memory to be assigned to an LDom. See the *Bundled Agents Reference Guide* for more information on the new attributes.
- **RemoveLDomConfigForMigration:** If enabled, agent removes the LDom configuration from the system during offline and clean operations provided `CfgFile` attribute is configured. You must enable this attribute if domain migration is planned for the logical domain. This is because, domain migration cannot be performed if logical domain configuration is present on target node for migration.
- **IPAddress:** The IP address to be assigned to an exclusive IP zone at this site after a cross-site failover. The agent writes the IP address inside the Zone Root in the `/etc/hostname.Device` file, where `Device` is the value of `Device` key.
- **Netmask:** The netmask to be used in an exclusive IP zone at this site after a cross-site failover. The agent writes the netmask inside the Zone Root in the file `/etc/netmasks`.
- **Gateway (Default Router):** The default gateway used by this Zone at this site.
- **DNS:** The domain name to use within the Zone at this site.
- **ConfigureNetwork:** Specifies if the LDom agent configures the network-bootarguments `PROM` variable of the guest domain.

Sybase agent attributes:

- **Quorum_dev:** The quorum device manages the cluster membership, stores cluster configuration data and contains information shared among server instances and nodes. It must be a disk accessible to all nodes in the cluster.

Specify fully qualified quorum device name. Note: This attribute should be specified only for cluster edition. For example:
`/dev/vx/rdisk/Sybase_install_dg/quorum_vol.`

- `interfaces_File`: Specifies the location of interfaces file for the Sybase instance. If this attribute is configured, `[-I interfaces file]` option is used when connecting to the isql session. If this attribute is not configured, the agent does not use the `-I` option.
- `ShutdownWaitLimit`: Specifies the maximum number of seconds for which the agent waits for the Sybase instance to stop after issuing the `shutdown with wait` command, and before attempting to issue the `kill -15 <data server-pid>` command, if required.
- `DelayAfterOnline`: Specifies the number of seconds that elapse after the Online entry point is complete and before the next monitor cycle is invoked.
- `DelayAfterOffline`: Specifies the number of seconds that elapse after the Offline entry point is complete and before the next monitor cycle is invoked.

SybaseBk agent attribute:

- `interfaces_File`: Specifies the location of the interfaces file for the Sybase instance. If this attribute is configured, `[-I interfaces file]` option is used when connecting to the isql session. If this attribute is not configured, the agent does not use the `-I` option.

Oracle agent attributes

- `DBName`: Specifies the database name. This attribute is required only for a policy-managed database. The value of this attribute must be set to the database name.
- `ManagedBy`: Specifies whether the database is administrator-managed or policy-managed. The default value of this attribute is `ADMIN`. You need not explicitly set its value. In a policy managed RAC database, this attribute must be set to **POLICY**.

DiskGroupSnap agent attributes:

- `FDType`: Specifies the configuration to be used for the firedrill. The possible values of this attribute are `Bronze` and `Gold` (default).

Zone agent attribute:

- `DeleteVCSZoneUser`: If enabled on a non-secure cluster, Zone agent deletes the VCS Zone user created for password-less communication between local-zone and global zone during offline and clean entry points. `DeleteVCSZoneUser` is disabled by default.

- **DROpts:** The value of this attribute consists of the following keys that define disaster recovery options for the Zone agent:
 - DNSDomain
 - DNSSearchPath
 - DNSServers
 - Gateway
 - Device
 - IPAddress
 - Netmask
 - Hostname

In a DR configuration, if one or more of these keys are set, the resource is considered to be DR-enabled. If all the keys stay at their default value (""), then the resource is not DR-enabled even if it is in a disaster recovery configuration.

Changes to product clusters in secure mode

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

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

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

Changes related to High Availability and Disaster Recovery (HA/DR)

Campus Cluster firedrill made easy

Campus Cluster firedrill configuration has been made easier by introducing a new attribute `FDType`. You can control the firedrill flavor by setting just this attribute instead of specifying specific values for other attributes.

You need to install the Global Cluster Option (GCO) license to run firedrill in a Campus Cluster configuration.

Support VVR multiple secondaries with GCO

The RVGPrimary agent is now enhanced to start the replication to the other secondaries in the RDS after a secondary becomes Primary. The online process of the agent initiates a differential synchronization to the other secondaries from the new primary and then starts the replication to those using the checkpoint.

Added disaster recovery support for the AIX and Solaris virtualization environments

Disaster recovery (DR) support is added to make it easier to customize AIX WPARs, Oracle VM Server for SPARC, and Solaris Zones during fail over to a DR site. The respective agents (WPAR, LDom, and Zone) are enhanced to accept cluster (subnet) specific network parameters that are applied when the agents brings the respective resource online.

Changes related to virtualization support in VCS

New LPAR agent on AIX

The LPAR agent monitors the logical partitions running on the same physical host and brings them online and offline. LPAR agent communicates with HMC to perform its operations.

You can use this agent to make the LPARs highly available and to monitor them. This agent is added as a part of virtualization support.

New KVMGuest agent on Linux

The KVMGuest agent monitors the Linux Kernel-based Virtual Machines (KVM guest) and brings the KVM guest online and offline. KVMGuest agent uses virsh commands.

You can use this agent to make the KVM guests highly available and to monitor them. This agent is added as a part of virtualization support.

Domain migration support for Oracle VM server (OVM) for SPARC

Added support for logical domain migration (warm migration on OVM 2.0 and live migration on OVM 2.1) in Oracle VM Server for SPARC environment. LDom agent detects the logical domains being migrated and changes the state of VCS resource accordingly.

See the *Veritas Storage Foundation and High Availability Virtualization Guide* for more information on supported configurations for domain migration.

Added support for multiple I/O domains to the guest domains

Added support for logical domains with access storage and network services from multiple I/O domains. Added a new agent AlternateIO to monitor the storage and network services from primary and alternate I/O domains. Current release supports only two I/O domains for guest domain.

See the *Virtualization Guide* for more information on supported configurations. See the *Bundled Agents Reference Guide* for more information on the AlternateIO agent.

Changes related to Veritas Storage Foundation and High Availability (SFHA)

Storage Foundation and High Availability (SFHA) includes the new features and changes introduced in 6.0 of the underlying products.

See [“Changes related to Veritas Storage Foundation \(SF\)”](#) on page 14.

See [“Changes related to Veritas Cluster Server \(VCS\)”](#) on page 23.

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

SFCFSHA includes the new features and changes introduced in 6.0 of the underlying products.

SFCFSHA includes the following changes in 6.0:

Enhancements to master node selection for failover

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

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

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

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

Availability of shared disk group configuration copies

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

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

Enhancements to CVM detach policies

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

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

Changes related to Symantec VirtualStore (SVS)

Symantec VirtualStore includes the new features and changes introduced in 6.0 of the underlying products.

Symantec VirtualStore includes the following changes in 6.0:

Enabled VMware View Integration

This allows you to import FileSnap cloned machines as a new Desktop Pool in VMware View.

See the *Symantec VirtualStore Installation and Configuration Guide* for more information on creating virtual machine clones using Symantec FileSnap.

Direct NFS support in Clustered NFS environments for Oracle

Added support for NFS coming from Clustered NFS. Symantec VirtualStore (SVS) can serve as the backend storage for Oracle databases.

Changes related to Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)

Direct NFS (dNFS) is an optimized NFS client that provides faster and more scalable access to NFS storage located on NAS storage accessible over TCP/IP.

See the *Symantec VirtualStore Administrator's Guide* for more information on deploying Oracle with Clustered NFS and VirtualStore utilities for the Oracle database.

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

Administering iSCSI with VirtualStore

The `svsiscsiadm` command provides a mechanism to simplify the administration of exporting iSCSI LUNs backed by files residing on the VirtualStore file system. This utilizes the iSCSI target driver implementation that is shipped with the operating system.

See the *Symantec VirtualStore Administrator's Guide*.

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

Changes related to Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)

SF Oracle RAC includes the new features and changes introduced in 6.0 of the underlying products.

Support for Web-based installation and configuration

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

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

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

- Installation and configuration of product
- Uninstallation of product

Support for policy-managed database environments

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

Changes related to Veritas Storage Foundation for Sybase ASE CE (SF Sybase CE)

SF Sybase CE includes the new features and changes introduced in 6.0 of the underlying products.

About Veritas Storage Foundation for Sybase ASE CE

Veritas Storage Foundation™ for Sybase® Adaptive Server Enterprise Cluster Edition (product) by Symantec leverages proprietary storage management and high availability technologies to enable robust, manageable, and scalable deployment of application on UNIX platforms. The solution uses cluster file system technology that provides the dual advantage of easy file system management as well as the use of familiar operating system tools and utilities in managing databases.

product integrates existing Symantec storage management and clustering technologies into a flexible solution which administrators can use to:

- Create a standard toward application and database management in data centers. product provides flexible support for many types of applications and databases.
- Set up an infrastructure for application that simplifies database management while fully integrating with application clustering solution.
- Apply existing expertise of Symantec technologies toward this product.

The solution stack comprises the Veritas Cluster Server (VCS), Veritas Cluster Volume Manager (CVM), Veritas Cluster File System (CFS), and Veritas Storage Foundation, which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

Benefits of SF Sybase CE

SF Sybase CE provides the following benefits:

- Use of a generic clustered file system (CFS) technology or a local file system (VxFS) technology for storing and managing application installation binaries.
- Support for file system-based management. product provides a generic clustered file system technology for storing and managing application data files as well as other application data.
- Use of Cluster File System (CFS) for the application quorum device.
- Support for a standardized approach toward application and database management. A single-vendor solution for the complete product software stack lets you devise a standardized approach toward application and database management. Further, administrators can apply existing expertise of Veritas technologies toward product.
- Easy administration and monitoring of product clusters from a single web console.
- Enhanced scalability and availability with access to multiple application instances per database in a cluster.
- Prevention of data corruption in split-brain scenarios with robust SCSI-3 Persistent Reservation (PR) based I/O fencing.
- Support for sharing all types of files, in addition to application database files, across nodes.
- Increased availability and performance using Veritas Dynamic Multi-Pathing (DMP). DMP provides wide storage array support for protection from failures and performance bottlenecks in the Host Bus Adapters (HBAs) and Storage Area Network (SAN) switches.
- Fast disaster recovery with minimal downtime and interruption to users. Users can transition from a local high availability site to a wide-area disaster recovery environment with primary and secondary sites. If a node fails, clients that are attached to the failed node can reconnect to a surviving node and resume access to the shared database. Recovery after failure in the product environment is far quicker than recovery for a failover database.
- Support for block-level replication using VVR.

Licensing changes in the SFHA Solutions 6.0 release

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

- The Cluster File System license is deprecated. CFS customers are entitled to the Storage Foundation Cluster File System High Availability (SFCFS HA) functionality.

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

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

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

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

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

Table 1-2 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.
<i>Veritas Storage Foundation and High Availability Release Notes</i>	Release-specific information for Veritas Storage Foundation and High Availability users.
<i>Veritas Storage Foundation and High Availability Solutions Solutions Guide</i>	Solutions and use cases for Veritas Storage Foundation and High Availability Solutions.
<i>Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide</i>	Troubleshooting information for Veritas Storage Foundation and High Availability Solutions.
<i>Veritas Storage Foundation and High Availability Solutions Virtualization Guide</i>	Virtualization-related information for Veritas Storage Foundation and High Availability Solutions.
<i>Symantec VirtualStore Release Notes</i>	Release-specific information Symantec VirtualStore.

Table 1-2 New documents (*continued*)

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

[Table 1-3](#) lists the documents that are deprecated in this release.

Table 1-3 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 1-3 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 1-4](#) lists documents that are no longer bundled with the binaries. These documents are now available online.

Table 1-4 Online documents

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

Changes related to product documentation