

Veritas Storage Foundation™ Release Notes

Linux

5.1



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: 5.1

Document version: 5.1.2

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas, Veritas Storage Foundation 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 maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

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

Contacting Technical Support

Customers with a current maintenance 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 non-technical support Web page at the following URL:

customercare.symantec.com

Customer service

Customer Care information is available at the following URL:

www.symantec.com/customercare

Customer Service is available to assist with 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 maintenance 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

Documentation feedback

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:

sfha_docs@symantec.com

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance 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

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4
Chapter 1 Storage Foundation Release Notes	9
Overview of this release	9
About the Simple Admin utility	10
Veritas Operation Services	11
Changes in Storage Foundation	11
Installation and upgrade	11
Storage Foundation	26
Veritas Volume Manager	29
Veritas Volume Replicator	31
Storage Foundation for Databases supported features	32
System requirements	33
Hardware and software requirements	33
Supported Linux operating systems	33
Storage Foundation memory requirements	34
Disk space requirements	34
Storage Foundation and High Availability Solutions 5.1 patches	34
Database requirements	35
Component product release notes	35
Software limitations	35
Veritas Storage Foundation software limitations	36
Veritas Volume Manager software limitations	36
Veritas File System software limitations	37
Veritas Volume Replicator software limitations	38
Fixed issues	39
Veritas Storage Foundation fixed issues	39
Storage Foundation for Databases (SFDB) tools fixed issues	39
Veritas Volume Manager and Veritas Volume Replicator fixed issues	40
Veritas File System fixed issues	59
Known issues	62
Veritas Storage Foundation known issues	62
Veritas Volume Manager known issues	70
Veritas File System known issues	79

Veritas Volume Replicator known issues	82
No longer supported	85
VVR will use TCP as default protocol from next release	87
Documentation	87
Relevant component product release notes	87
Storage Foundation guides	87
Manual Pages	90
 Appendix A	
Veritas Storage Foundation 5.1 for Xen	91
Supported operating systems for Xen	91
Supported features	91
Unsupported features	92
Installing the Veritas Software in Dom0	92
Verifying Software Versions	93
Uninstalling the Veritas Software from Dom0	94

Storage Foundation Release Notes

This chapter includes the following topics:

- [Overview of this release](#)
- [Changes in Storage Foundation](#)
- [Storage Foundation for Databases supported features](#)
- [System requirements](#)
- [Component product release notes](#)
- [Software limitations](#)
- [Fixed issues](#)
- [Known issues](#)
- [No longer supported](#)
- [Documentation](#)

Overview of this release

Before you continue, make sure that you are using the current version of this guide. It is online at:

http://sfdoccentral.symantec.com/sf/5.1/linux/sf_notes.pdf

This document is Document version: 5.1.2.

This document provides release information about the products in the Veritas Storage Foundation 5.1 product line:

- Veritas Storage Foundation™ (Basic, Standard, Standard HA, Enterprise, and Enterprise HA)
- Veritas™ Volume Manager (VxVM)
- Veritas™ File System (VxFS)
- Veritas Storage Foundation™ Cluster File System (SFCFS)
- Veritas Storage Foundation™ Cluster File System for Oracle RAC (SFCFS for Oracle RAC)

See the *Veritas Storage Foundation Installation Guide*.

For the latest patches available for this release, go to: <http://vos.symantec.com/>.

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/335001>

The hardware compatibility list (HCL) contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:

<http://entsupport.symantec.com/docs/330441>

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.

Review this entire document before installing your Veritas Storage Foundation product.

This document does not contain release information for Veritas Cluster Server.

See the *Veritas Cluster Server Release Notes*.

About the Simple Admin utility

Veritas Storage Foundation has an optional utility, called Simple Admin, that you can use with Veritas File System and Veritas Volume Manager. The Simple Admin utility simplifies storage management by providing a single interface to the administrator and by abstracting the administrator from many of the commands needed to create and manage volumes, disks groups, and file systems.

You can download the Simple Admin utility for Veritas Storage Foundation from the following URL:

http://www.symantec.com/business/products/agents_options.jsp?pcid=2245&pvid=203_1

Veritas Operation Services

Veritas Operations Services (VOS) is a Web-based application that is designed specifically for Veritas Storage Foundation and High Availability products. VOS increases operational efficiency and helps improve application availability.

Among its broad set of features, VOS evaluates the systems in your environment to determine if you are ready to install or upgrade Storage Foundation and High Availability products.

To access VOS, go to:

<http://go.symantec.com/vos>

Changes in Storage Foundation

This section describes the changes in Veritas Storage Foundation 5.1.

Installation and upgrade

Storage Foundation installation and upgrade includes the following changes in 5.1:

Veritas keyless licensing

This release of the Veritas products introduces the option to install without a license key. A license key is used during the installation of a software to identify that the user has the right to use the software. Previous releases of Veritas products required that you obtain a license key prior to installation. The installer required that you enter a valid license key before the Veritas software was installed.

The keyless license strategy does not eliminate the need to obtain a license. A software license is a legal instrument governing the usage or redistribution of copyright protected software. The administrator and company representatives must ensure that a server or cluster is entitled to the license level for the products installed. Symantec reserves the right to ensure entitlement and compliance through auditing.

In this release of the Veritas products, the product installer does not require that you enter a license key. The installer prompts you to select one of the following licensing methods:

- Install a license key for the product and features that you want to install.
- Continue to install without a license key.
The installer prompts for the product modes and options that you want to install, and then sets the required product level.

Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server. If you do not comply with the above terms, continuing to use the Veritas product is a violation of your end user license agreement, and results in warning messages. For more information about keyless licensing, see the following URL:
<http://go.symantec.com/sfhakeyless>

If you upgrade to this release from a prior release of the Veritas software, the product installer does not change the license keys that are already installed. The existing license keys may not activate new features in this release.

If you upgrade with the product installer, or if you install or upgrade with a method other than the product installer, you must do one of the following to license the products:

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.
- Use the `vxlicinst` command to install a valid product license key for the 5.1 products you have purchased.

You can also use the above options to change the product levels to another level that you are authorized to use.

Symantec recommends updating to keyless licensing for the following reasons:

- it enables 5.1 functionality.
- it allows you to change the product level easily.

For information about setting or changing the product level, see the Installation Guide for your Veritas product.

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

Packaging updates

[Table 1-1](#) lists the updates related to packages for this release.

Table 1-1 List of packages

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
N/A	VRTScps	New package.	Veritas Cluster Server Coordination Point Server

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
N/A	VRTSvcsea	New package.	Veritas Cluster Server Enterprise Agents
N/A	VRTSsfmh	New package.	Veritas Storage Foundation Managed Host
N/A	VRTSaslapm	New package.	Volume Manager ASL/APM
SYMClma	N/A	Obsolete in 5.0MP3. Functionality dropped.	Symantec License Inventory Agent
VRTSaa	VRTSsfmh	Consolidated into VRTSsfmh.	Veritas Enterprise Administrator action agent
VRTSaclib	N/A	Obsolete in 5.1. Not available for fresh installation. Only available to upgrade customers.	Veritas Cluster Server ACC Library 5.0 by Symantec
VRTSalloc	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation Intelligent Storage Provisioning
VRTSat	VRTSat	No change.	Symantec Product Authentication Service
VRTSatClient	VRTSatClient	No change.	Symantec Product Authentication Service
VRTSatServer	VRTSatServer	No change.	Symantec Product Authentication Service
VRTScavf	VRTScavf	No change.	Veritas Cluster Server Agents for Storage Foundation Cluster File System
VRTSccg	VRTSsfmh	Consolidated into VRTSsfmh.	Veritas Enterprise Administrator Central Control Grid
VRTScfsdc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	Veritas Cluster File System Documentation

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTScmccc	N/A	Obsolete in 5.1. Delivered with 5.x CMC release.	Veritas Cluster Management Console Cluster Connector
VRTScmcdc	N/A	Obsolete in 5.0MP3. Delivered with 5.x CMC release.	User Documentation for Veritas Cluster Management Console
VRTScmcm	N/A	Obsolete in 5.0MP3. Delivered with 5.x CMC release.	Veritas Cluster Management Console for multicluster environments
VRTScmcs	N/A	Obsolete in 5.1. Delivered with 5.x CMC release.	Veritas Cluster Management Console for single cluster environments
VRTScs	N/A	Obsolete in 5.0MP3. Delivered with SFM release.	Veritas Centralized Management for Storage Foundation Management Server
VRTScscm	N/A	Obsolete in 5.1. Available for download from http://go.symantec.com/vcsmc	Veritas Cluster Server Cluster Manager
VRTScscw	N/A	Obsolete in 5.1.	Veritas Cluster Server Configuration Wizards
VRTScsdoc	N/A	Obsolete in 5.0MP3. Delivered with SFM release.	Veritas Enterprise Administrator Central Server Documentation
VRTScsocw	N/A	Obsolete in 5.1.	Veritas Cluster Server Oracle and RAC Configuration Wizards
VRTScssim	N/A	Obsolete in 5.1. Available for download from http://go.symantec.com/vcsmc .	Veritas Cluster Server Simulator
VRTScutil	VRTScutil	Expanded to include few VCS packages.	Veritas Cluster Utility
VRTScweb	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Enterprise Administrator Central Server Documentation

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSd2gui	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation Graphical User Interface for DB2
VRTSdb2ed	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation for DB2
VRTSdbac	VRTSdbac	No change.	Veritas Oracle Real Application Cluster Support Package
VRTSdbcom	VRTSdbed	Consolidated into VRTSdbed.	Veritas Storage Foundation Common Utilities for Databases
VRTSdbdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	Veritas Storage Foundation Documentation for Databases
VRTSdbed	VRTSdbed	Expanded to include DBED packages.	Veritas Storage Foundation for Oracle
VRTSdbms3	N/A	Obsolete in 5.1. Sybase ASA repository no longer used in 5.1.	Symantec Shared DBMS
VRTSdcli	N/A	Obsolete in 5.1.	Veritas Distributed Command Line Interface
VRTSdcp	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Disk Correlator Provider
VRTSddlpr	N/A	Obsolete in 5.1. Functionality merged into VRTSob.	Veritas Device Discovery Layer Services Provider
VRTSdsa	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Datacenter Storage Agent
VRTSdsm	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Datacenter Storage Manager

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSfas	N/A	Obsolete in 5.0MP3. Functionality dropped.	Veritas FlashSnap Agent for Symmetrix
VRTSfasag	N/A	Obsolete in 5.0MP3. Functionality dropped.	Veritas Cluster Server Agents for Veritas FlashSnap Agent for Symmetrix
VRTSfasdc	N/A	Obsolete in 5.0MP3. Functionality dropped.	Veritas FlashSnap Agent for Symmetrix Documentation
VRTSfsdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	Veritas File System Documentation
VRTSfsman	VRTSvxf	Consolidated into VRTSvxf.	Veritas File System - Manual Pages
VRTSfsmnd	VRTSfssdk	Consolidated into VRTSfssdk.	Veritas File System SDK - Manual Pages
VRTSfspro	VRTSob	Consolidated into VRTSob.	Veritas File System Management Services Provider
VRTSfssdk	VRTSfssdk	No change.	Veritas File System SDK
VRTSfsweb	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas File System Provider Web Client Extension
VRTSgab	VRTSgab	No change.	Veritas Group Membership and Atomic Broadcast
VRTSgapms	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas Generic Array Plug-in for Mapping Services
VRTSgcscha	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas GCS high availability agents
VRTSgcspr	N/A	Obsolete in 5.0MP3. Delieverd with SFM release.	Veritas SAN Global Configuration Server Object Bus Provider

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSglm	VRTSglm	No change.	Veritas Global Lock Manager
VRTSgms	VRTSgms	No change.	Veritas Group Messaging Services
VRTSicsco	N/A	Obsolete in 5.1.	Symantec Infrastructure Core Services Common
VRTSjre	N/A	Obsolete in 5.0MP3.	Veritas Java Runtime Environment Redistribution
VRTSjre15	N/A	Obsolete in 5.1.	Symantec Java Runtime Environment Redistribution
VRTSllt	VRTSllt	No change.	Veritas Low Latency Transport
VRTSsvmconf	VRTSsvmconf	No change.	Veritas Linux LVM to VxVM Converter
VRTSmapro	N/A	Consolidated into VRTSsob.	Veritas Storage Mapping Provider
VRTSmh	VRTSsfmh	Consolidated into VRTSsfmh.	Veritas Storage Foundation Management host
VRTSsob	VRTSsob	No change.	Veritas Enterprise Administrator Service
VRTSsobc33	N/A	Obsolete in 5.1. Functionality Delivered with SFM release	Veritas Enterprise Administrator Core
VRTSsobgui	N/A	Obsolete in 5.1. Functionality Delivered with SFM release.	Veritas Enterprise Administrator
VRTSsobweb	N/A	Obsolete in 5.1. Functionality Delivered with SFM release.	Veritas Enterprise Administrator Web Console
VRTSodm	VRTSodm	No change	Veritas Oracle Disk Manager
VRTSorgui	N/A	Obsolete in 5.1. No longer supported.	Veritas Storage Foundation Graphical User Interface for Oracle

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSspb	N/A	Obsolete in 5.1.	Symantec Private Branch Exchange
VRTSperl	VRTSperl	No change.	Veritas Perl 5.8.8 redistribution
VRTSsmf	N/A	Obsolete in 5.0MP3.	Symantec Service Management Framework
VRTSspt	VRTSspt	No change.	Veritas Software Support Tools
VRTSsybed	N/A	Obsolete in 5.1. Functionality dropped.	Veritas Storage Foundation for Sybase
VRTSvail	N/A	Obsolete in 5.1. Functionality Delivered with SFM release.	Veritas Array Providers
VRTSvcsvcs	VRTSvcsvcs	Expanded to include few VCS packages.	Veritas Cluster Server
VRTSvcsvcsag	VRTSvcsvcsag	Expanded to include agents previously included in VRTSvcsvcsvr.	Veritas Cluster Server Bundled Agents
VRTSvcsvcsdb	VRTSvcsvcsdb	Consolidated into VRTSvcsvcsdb.	Veritas High Availability Agent for DB2
VRTSvcsvcsdc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	User Documentation for Veritas Cluster Server
VRTSvcsvcsdr	VRTSvcsvcsdr	No change.	Veritas Cluster Server Disk Reservation Modules and Utilities by Symantec
VRTSvcsvcsmsg	VRTSvcsvcs	Consolidated into VRTSvcsvcs.	Veritas Cluster Server English Message Catalogs
VRTSvcsvcsmn	VRTSvcsvcs	Consolidated into VRTSvcsvcs.	Manual Pages for Veritas Cluster Server
VRTSvcsvcsor	VRTSvcsvcsor	Consolidated into VRTSvcsvcsor.	Veritas High Availability Agent for Oracle

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSvcssy	VRTSvcsea	Consolidated into VRTSvcsea.	Veritas High Availability Agent for Sybase
VRTSvcsvr	VRTSvcstag	Consolidated into VRTSvcstag.	Veritas Cluster Server Agents for Veritas Volume Replicator
VRTSvdid	N/A	Obsolete in 5.1.	Veritas Device Identification API
VRTSvlic	VRTSvlic	No change.	Symantec License Utilities
VRTSvmdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	User Documentation for Veritas Volume Manager
VRTSvmman	VRTSvxvm	Consolidated into VRTSvxvm.	Manual Pages for Veritas Volume Manager
VRTSvmpro	N/A	Consolidated into VRTSob.	Veritas Volume Manager Management Services Provider
VRTSvmweb	N/A	Obsolete in 5.0MP3. Delivered with SFM release.	Veritas Volume Manager Management Services Web Client Extensions
VRTSvrdoc	N/A	Obsolete in 5.0MP3. Documentation available in DVD media as PDFs.	User Documentation for Veritas Volume Replicator
VRTSvrpro	N/A	Consolidated into VRTSob.	Veritas Volume Replicator Management Services Provider
VRTSvrw	N/A	Obsolete in 5.1. Delivered with SFM release.	Veritas Volume Replicator Web Console
VRTSvsvc	N/A	Obsolete in 5.0MP3.	Veritas Volume Server and Client Provider
VRTSvxfen	VRTSvxfen	No change.	Veritas I/O Fencing

Table 1-1 List of packages (*continued*)

5.0 Package Name	5.1 Package name	Explanation of changes	Package description
VRTSvxfs	VRTSvxfs	Expanded to include VRTSfsman (man pages). On Linux: VRTSvxfs-common and VRTSvxfs-platform packages are consolidated into single VRTSvxfs package.	Veritas File System
VRTSvxmsa	N/A	Obsolete in 5.1. Functionality dropped.	Veritas VxMS Mapping Service, Application Libraries
VRTSvxvm	VRTSvxvm	Expanded to include VRTSvxman (man pages). On Linux: VRTSvxvm-common and VRTSvxvm-platform packages are consolidated into single VRTSvxvm package.	Veritas Volume Manager binaries
VRTSweb	N/A	Obsolete in 5.1.	Symantec Web Server

Install options for Storage Foundation High Availability and the Storage Foundation Cluster File System High Availability

The product installation programs now prompt you for whether you want to install the high availability packages when you are installing Storage Foundation or Storage Foundation Cluster File System. This change enables you to explicitly choose which functionality is installed. In previous releases, the installed license key determined what functionality was installed.

The product installer displays Storage Foundation High Availability and Storage Foundation Cluster File System High Availability on the product selection menu.

Improved response file generation

You can now create response files without performing a live installation, using the installation simulator.

Option to install only the minimal packages

The product installer now provides several options for which packages to install. For each product, you can install the minimal packages, the recommended packages or all of the packages.

When you install with the product installer, you can select from one of the following install options:

- Minimal RPMs: installs only the basic functionality for the selected product.
- Recommended RPMs: installs the full feature set without optional RPMs.
- All RPMs: installs all available RPMs.

When you install with a product install script, such as `installsf`, you can use the following options to the install script:

- `-minpkgs`: displays the RPMs and patches required for basic functionality.
- `-recpkgs`: displays the recommended RPMs and patches.
- `-allpkgs`: displays all available RPMs and patches.

Veritas extension for Oracle Disk Manager RPM is installed by default for Storage Foundation and Storage Foundation Cluster File System

The Veritas extension for Oracle Disk Manager package is supported with a Storage Foundation license.

In this release, the product installer now installs the required RPM for ODM by default when Storage Foundation is installed.

For Storage Foundation Cluster File System, the GMS

`VRTSgms-5.1.00.00-A39_<dist>.x86_64.rpm` RPM is also installed.

where `<dist>` is the name of the Linux OS either `rhel5`, `sles10`, or `sles11`.

Option to create response file templates

You can use the `-makeresponsefile` option of the installer to create response file templates.

The installer also generates a response file after each successful installer task, such as installation, configuration, uninstallation, or upgrade. These response files contain the details that you provided to the installer questions in the form of values for the response file variables. The response file also contains descriptions and explanations of the variables and their values.

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

Option to start or stop Storage Foundation HA

After the installation and configuration is complete, the installer starts the processes that the installed products use. You can use the product installer to stop or start the processes and load or unload the drivers, if required.

See the *Veritas Storage Foundation and High Availability Installation Guide* for more details.

Support for installer resilience

If an installation or upgrade of Storage Foundation is interrupted, the next time you re-run it the installer discovers the presence of an installer instance. The installer then gives an option to resume the installation or upgrade.

See the *Veritas Storage Foundation and High Availability Installation Guide* for more details.

Support for installing Storage Foundation HA using Kickstart

The Veritas product installer creates a Veritas Kickstart script configuration file to install Storage Foundation HA using the Kickstart utility for RHEL. The file contains the list of Storage Foundation HA RPMs in the correct installation order. The file contains the RPMs in the format that the Kickstart utility can use for installation.

Installation options for Kickstart

The installer program can generate sample configuration scripts that integrate with native installation tools. These sample scripts include all of the relevant Symantec packages in the correct order for installation. This feature is supported for installations, not upgrades.

Installer does not proceed with installation in RSH/SSH disabled environments

In the previous releases of Storage Foundation, in certain secure enterprise environments where RSH or SSH communication was not enabled, the installer installed and configured Storage Foundation only on the local system and the systems with which it could communicate. The installer also generated a response file that you could copy to the other systems in the cluster to identically install and configure Storage Foundation on other systems.

With this release of Storage Foundation, the installer mandates the availability of either RSH or SSH communication between the systems to perform any installer task.

Support for Web-based installer

This release supports an interactive installation using the Web-based installer. You can use a Web-interface to install and configure Storage Foundation HA.

The Web-installer can perform the following functions:

- Install Storage Foundation HA
- Uninstall Storage Foundation HA
- Configure Storage Foundation HA
- Upgrade Storage Foundation HA
- Start and stop Storage Foundation HA
- Perform an installation precheck

The installer program's default answer is no to configure optional features

The installer's default answer to configure optional features is now no. You must enter **y** if you want to configure certain optional features.

Gathering requirements using the installer program

You can use the `-requirements` option of the installer to gather the installation requirements. Web-based installer also provides you with a similar option.

The following information is displayed:

- Required operating system level
- Required patches
- Required disk space
- Other requirements

Support to continue installation after a successful precheck

The installer program has more sophisticated precheck, installation, and configuration options, which follow in outline:

- When you perform a successful precheck, you have the option to continue with the installation.
- After a successful installation, you have the option to continue with the configuration, or you can return to the configuration later.

Selecting default systems for installation

From the local system, the installer program checks for the `/etc/llhosts` for node names. When found, the installer program presents these as default nodes for installation. If the `llhosts` file is not present, then no default node names are provided.

Communication modes

By default, the installer program uses SSH for communication. The installer program switches to RSH if password-less SSH is not enabled.

For RSH communication, the `-rsh` option is available.

The installer programs supports mixed RSH and SSH modes on nodes in a cluster. The installation program can install on systems which may have heterogeneous (RSH and/or SSH) communication modes enabled.

IPv6 support for the installer programs

You can now use the installer to install and configure Storage Foundation HA on systems with IPv4, IPv6, or mixed stack configurations.

The following limitations apply:

- RSH is not supported on Linux
- The installer program does not support link-local IPv6 addresses on Linux.

Adding a node using the `-addnode` option

The `-addnode` option has been added to the installer to add a node to a running cluster. Based on the existing cluster configuration, the installer also configures the new node to use Symantec Product Authentication service and to use I/O fencing.

The installer also supports adding a node to a single node cluster, but stops the cluster during the addition of the node.

Silent and automated installation enhancements for response files

The installer program supports silent installations using response files.

Operations that you can perform using response files follow:

- Fresh installations
- Configurations

- Uninstallations
- Upgrades from previous supported releases

Using bonded links during installation

The installer program asks if you want to use a bonded NIC, if so it configures the `lfttab` file for you. Note that the installer program does not detect bonded links.

Command options to help troubleshoot installations

You can run the installer with the `-debug` option and the `-trace` option to troubleshoot an installation.

Supported paths for Storage Foundation upgrades that do not require a system reboot

When you perform a typical upgrade using the installer program from Storage Foundation or Storage Foundation High Availability versions 4.1 MP4, 5.0, 5.0 MP1, 5.0 MP2, and 5.0 MP3 to version 5.1, a system reboot is not required.

Upgrades that follow any other upgrade paths require a reboot.

Changes related to the installer for cross-product upgrades

This release includes the following changes related to the cross-product upgrades:

- If you try to perform a cross-product upgrade, the installer now gracefully exits with an error message.
For example, if you choose to upgrade SFHA 5.0 MP3 to SFCFS 5.1, the installer displays the following error message:

```
SFHA 5.0.30.00 is installed.  
Upgrading SFHA 5.0.30.00 directly to SFCFS 5.1 is not supported.
```

The installer does not support a direct upgrade from a previous SF or SFHA version to SFCFS or SFCFS Oracle RAC version 5.1. You must upgrade SF or SFHA to version 5.1, and then install the 5.1 version of the stack product. See the appropriate product Installation Guides for upgrade instructions.

- If a previous version of SFHA is installed, the installer supports partial product upgrade.
You can upgrade only VCS or SF to version 5.1. For example, you can upgrade SFHA 5.0 MP3 to SF 5.1. If you want to upgrade the complete SFHA stack later, you can run the `installvcs` program.

See the *Veritas Storage Foundation High Availability Installation Guide* for supported upgrade paths.

Storage Foundation

Storage Foundation includes the following changes in 5.1:

Veritas Enterprise Administrator Graphical User Interface

The Veritas Enterprise Administrator (VEA) console is no longer packaged with Storage Foundation products. Symantec recommends use of Storage Foundation Manager to manage, monitor and report on Storage Foundation product environments. You can download this utility at no charge at <http://go.symantec.com/vom>. If you wish to continue using VEA, a version is available for download from <http://go.symantec.com/vom>.

Changes to Storage Foundation for Databases (SFDB) functionality

The Storage Foundation for Databases (SFDB) feature for enhanced management of Oracle databases is included with Enterprise licensing for Storage Foundation and Storage Foundation HA. The SFDB tools provide enhanced ease-of-use commands which can be run by a database administrator without root privilege to optimize storage for an Oracle database environment. This extended functionality is supported for configurations with single instance Oracle and includes the following:

- Database Checkpoints
- Database FlashSnap
- Database Cloning
- Database Dynamic Storage Tiering

For information on using SFDB tools, see the *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases* guide.

SFDB new features

New features in the Storage Foundation for Databases tools package for database storage management for release 5.1:

- SQLite repository
- Multiple disk group support for FlashSnap
- Mapped mount points for individual volumes for Database Flashsnap clones

- Oracle Dataguard support
- Oracle Enterprise Manager (OEM) Plugin
- Cached ODM support

SFDB feature changes

If you are upgrading from Storage Foundation for Oracle (HA) 4.x or 5.0 to Storage Foundation 5.1, the following changes in functionality will apply.

Commands which have changed:

- `sfua_db_config` functionality is changed: this command is no longer needed to create a SFDB repository. The functionality of `sfua_db_config` is now used to set user and group access to various SFDB directories.
- Use the `dbed_update` command to create a new SQLite SFDB repository.
- `sfua_rept_adm` was used in release 5.0 to perform repository backup and restore and this command will be obsolete in release 5.1.
- The `sfua_rept_util` command is used to perform SQLite repository backup and restore.
- The `sfua_rept_migrate` command is added for migrating the SFDB repository from the 4.x or 5.0 release to the 5.1 release.

Commands which continue to be supported:

- `dbed_update`
- Database Storage Checkpoint commands: `dbed_ckptcreate`, `dbed_ckptdisplay`, `dbed_ckptmount`, `dbed_ckptquota`, `dbed_ckptremove`, `dbed_ckptrollback`, `dbed_clonedb`, `dbed_ckptumount`
- Database Flashsnap commands: `dbed_vmchecksnap`, `dbed_vmclonedb`, `dbed_vmsnap`
- Database Dynamic Storage Tiering commands: `dbdst_addvol`, `dbdst_admin`, `dbdst_classify`, `dbdst_convert`, `dbdst_file_move`, `dbdst_partition_move`, `dbdst_preset_policy`, `dbdst_rmvol`, `dbdst_show_fs`, `dbdst_tbs_move`, `dbdst_report`

SFDB features which are no longer supported

Commands which are no longer supported in release 5.1:

- ORAMAP (`libvxoramap`)
- Storage mapping commands `dbed_analyzer`, `vxstorage_stats`
- DBED providers (DBEDAgent), Java GUI, and `dbed_dbprocli`.

The SFDB Oracle features can only be accessed through the command line interface. However, Veritas Storage Foundation Manager 2.1 (a separately licensed product) can display Oracle database information such as tablespaces, database to LUN mapping, and tablespace to LUN mapping.

- Storage statistics: `commandsdbdst_makelbfs, vxdbts_fstatsummary, dbdst_fiostat_collector, vxdbts_get_datafile_stats`
- `dbed_saveconfig, dbed_checkconfig`
- `dbed_ckptplan, dbed_ckptpolicy`
- `dbed_scheduler`

SmartMove™ feature

SmartMove reduces the time and I/O required to attach or reattach a plex to an existing VxVM volume, in the specific case where a VxVM volume has a VxFS file system mounted on it. The SmartMove feature uses the VxFS information to detect free extents and avoid copying them.

SmartMove provides the following benefits:

- Less I/O is sent through the host, through the storage network and to the disks/LUNs
- Faster plex creation, resulting in faster array migrations
- Ability to migrate from a traditional LUN to a thinly provisioned LUN, removing unused space in the process

See the *Veritas Storage Foundation Advanced Features Administrator's Guide* for more information.

Thin Storage Reclamation support

Thin Storage is an array vendor solution for allocating storage to applications only when the storage is truly needed, from a pool of free storage. Thin Storage attempts to solve the problem of under utilization of available array capacity.

Thin Storage Reclamation-capable arrays and LUNs allow the administrators to release once-used storage to the pool of free storage. Storage is allocated from the free pool when files are created and written to in the file system. However, this storage is not released to the free pool when files get deleted; the administrator must perform the operation of reclaiming this storage for the free pool.

Veritas File System supports reclamation of the free blocks in the file system on Veritas Volume Manager-backed file systems. The operation of reclamation can

be done on a disk, LUN, full file system, or part of a file system using the `vxdisk` and `fsadm` commands, and the `vxfs_ts_reclaim` API.

See the *Veritas Storage Foundation Advanced Features Administrator's Guide* for more information.

Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 5.1:

Support for SSD

A solid-state disk (SSD) is solid-state storage that emulates a hard disk drive (HDD) interface. Solid-state storage is a nonvolatile storage medium that employs integrated circuits rather than magnetic or optical media. The main advantage of solid-state storage is that it contains no mechanical parts. As a result, data transfer to and from solid-state storage media takes place at a much higher speed than is possible with HDD storage.

SSD devices are supported with Veritas Volume Manager. VxVM automatically discovers SSD devices and those can be displayed using the `vxdisk` and the `vxprint` commands. In this release, we support auto detection of SSD devices for FusionIO, EMC CLARiiON, Hitachi AM/WMS series, and Hitachi USPV/USPVM arrays. In case the devices are not being auto detected, the devices can be manually flagged as SSD or HDD using the `vxdisk` command.

See the `vxdisk(1M)` manual page for details.

SSD or HDD devices can be specified for allocation of VxVM volumes and the `vxassist` command with the `mediatype` attribute can be used to allocate storage on specific devices.

See the `vxassist(1M)` manual page for details.

Since SSD devices have faster read rates when compared to HDD devices, VxVM automatically picks up SSD plexes present in a mirror for reads, unless overridden by setting preferred read policy.

Note: The disk group version should be 150 or later for VxVM support for SSD devices.

See the *Veritas Volume Manager Administrator's Guide* for more information.

Handling of minor number conflicts

The volume device minor numbers in a disk group to be imported may conflict with existing volume devices. In releases of VxVM prior to 5.1, the conflicts resulted in failures; either the disk group imported failed, or the slave node failed to join for a shared disk group. When this happened, you had to run the `vxldg remminor` command manually to resolve the minor conflicts.

In this release, VxVM can automatically resolve minor number conflicts. When there exists a minor conflict when a disk group is imported, the disk group is automatically assigned a new base minor, and the volumes in the disk group are reminored based on the new base minor. You do not need to run the `vxldg remminor` command to resolve the minor conflicts.

See the *Veritas Volume Manager Administrator's Guide* for details.

Enhancements to the vxconfigd daemon

The `vxconfigd` daemon can now process some queries while a disk group import is in progress. This enhancement means that the user and agent scripts are able to check the status of Veritas Volume Manager objects during that time. In previous releases, such operations would block waiting for the import to finish.

In this release, the following commands may proceed in parallel with a disk group import:

- `vxldctl [-c] mode`
- `vxldg list [<dgname>]`
- `vxldisk list [<daname>]`
- `vxprint [-Aqht]`

Use the following command to disable this functionality:

```
# vxldctl request_threads 0
```

Upgrading the array support

The Storage Foundation 5.1 release includes all array support in a single RPM, `VRTSaslapm`. The array support RPM includes the array support previously included in the `VRTSvxvm` RPM. The array support RPM also includes support previously packaged as external array support libraries (ASLs) and array policy modules (APMs).

See the 5.1 Hardware Compatibility List for information about supported arrays.

<http://entsupport.symantec.com/docs/330441>

When you upgrade Storage Foundation products with the product installer, the installer automatically upgrades the array support. If you upgrade Storage Foundation products with manual steps, you should remove any external ASLs or APMs that were installed previously on your system. Installing the VRTSvxvm RPM exits with an error if external ASLs or APMs are detected.

After you have installed Storage Foundation 5.1, Symantec provides support for new disk arrays through updates to the `VRTSaslapm` package.

For more information about array support, see the *Veritas Volume Manager Administrator's Guide*.

The `vxdumpinq` utility renamed to the `vxscsiinq` utility

The diagnostic utility `vxdumpinq` has been renamed to the `vxscsiinq` utility.

Enclosure-based naming is now the default naming scheme

In this release, the enclosure-base naming scheme is now the default method Veritas Volume Manager uses for naming devices. The Veritas product installer no longer prompts you to select a naming scheme. If you perform a fresh installation, the naming scheme is set to enclosure-based naming.

When you upgrade from a previous release of a Storage Foundation product, the naming scheme is set to enclosure-based naming, with the following exception. If you explicitly set the naming scheme for the existing installation to operating system-based naming, the setting is preserved when you upgrade. That is, if you used the `vxddladm set namingscheme=osn` command for the existing installation, the upgraded 5.1 product retains the operating system-based naming.

To change the naming scheme after installation or upgrade, use the following command:

```
# vxddladm set namingscheme=osn|ebn [persistence=yes|no] \  
[lowercase=yes|no] [use_avid=yes|no]
```

For more information about device naming, see the *Veritas Volume Manager Administrator's Guide*.

Veritas Volume Replicator

Veritas Volume Replicator includes the following changes in 5.1:

SmartMove for VVR

The SmartMove for VVR feature enables VVR to leverage information from VxFS knowledge of the file system blocks in use to optimize the time and network bandwidth required for initial resync of replicated volumes.

See the *Veritas Volume Replicator Administrator's Guide* for more information on SmartMove for VVR.

Veritas Volume Replicator supports IPv6

Veritas Volume Replicator supports IPv6 in this release. IPv6 is supported only with disk group version 150 or later.

The Internet Protocol version 6 (IPv6) is the next-generation Internet Layer protocol for packet-switched networks and the Internet. IPv4 is the first version of the protocol to be widely deployed. IPv6 has a much larger address space than IPv4. This results from the use of a 128-bit address, whereas IPv4 uses only 32 bits. This expansion provides flexibility in allocating addresses and routing traffic and eliminates the primary need for network address translation (NAT). IPv6 also implemented new features that simplify aspects of address assignment and network renumbering when changing Internet connectivity providers. Network security is integrated into the design of the IPv6 architecture.

See “[IPv6 software limitations](#)” on page 38.

See the *Veritas Volume Replicator Administrator's Guide* for more information on VVR IP terminology.

See the *Veritas Storage Foundation Installation Guide* for more information on planning and upgrading VVR from a previous version of IPv4 to IPv6.

Storage Foundation for Databases supported features

See “[Changes to Storage Foundation for Databases \(SFDB\) functionality](#)” on page 26.

Note: Storage Foundation supports running SFDB tools with Oracle databases only.

For the most current information on Storage Foundation products and single instance Oracle versions supported, see:

<http://entsupport.symantec.com/docs/331625>

System requirements

This section describes the system requirements for this release.

Hardware and software requirements

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://entsupport.symantec.com/docs/330441>

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

Supported Linux operating systems

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

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/335001>

The Veritas 5.1 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux (OEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)

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

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

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported Red Hat erratas and updates and SUSE service packs is available in the following TechNote. Read this TechNote before you install Symantec products.

<http://entsupport.symantec.com/docs/335001>

Xen platform for Linux

The Veritas 5.1 release is also supported on the Xen platform for Linux, with some restrictions.

See “[Supported operating systems for Xen](#)” on page 91.

VMware Environment

For information about the use of this product in a VMware Environment, refer to <http://entsupport.symantec.com/docs/289033>

Mandatory patch required for Oracle Bug 4130116

If you are running Oracle versions 9.2.0.6 or 9.2.0.7, you must apply the Oracle patch for Oracle Bug 4130116. Contact Oracle to obtain this patch, and for details on how to apply it.

Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Disk space requirements

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

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

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

Storage Foundation and High Availability Solutions 5.1 patches

Symantec strongly recommends that you install Storage Foundation and High Availability Solutions (SFHA) 5.1 Patch 1 immediately after you install SFHA 5.1.

The patch for rhel5_x86_64 is available at the following URL:

<https://vos.symantec.com/patch/detail/2957>

The patch for sles10_x86_64 is available at the following URL:

<https://vos.symantec.com/patch/detail/2958>

The patch for sles11_x86_64 is available at the following URL:

<https://vos.symantec.com/patch/detail/2959>

Database requirements

The following TechNote identifies the most current information on supported database and operating system combinations:

<http://entsupport.symantec.com/docs/331625>

Note: Storage Foundation supports running Oracle, DB2, and Sybase on VxFS and VxVM.

Storage Foundation does not support running SFDB tools with DB2 and Sybase.

Component product release notes

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

The component product release notes for this release are included as PDF files on the software disc:

- *Veritas Storage Foundation Cluster File System Release Notes* (*sfcfs_notes.pdf*)
- *Veritas Storage Foundation Cluster File System for Oracle RAC Release Notes* (*sfcfsrac_notes.pdf*)
- *Veritas Cluster Server Release Notes* (*vcs_notes.pdf*)

Software limitations

The following sections describe Storage Foundation software limitations that exist in this release.

See the *Veritas Cluster Server Release Notes* for VCS software limitations.

Veritas Storage Foundation software limitations

The following are software limitations in the 5.1 release of Veritas Storage Foundation.

RSH does not support IPv6 address

The product installer does not support installation by RSH using an IPv6 address as a system name in the command line.

Veritas Volume Manager software limitations

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

Cluster Volume Manager (CVM) fail back behavior for non-Active/Active arrays (1441769)

This describes the fail back behavior for non-Active/Active arrays in a CVM cluster. This behavior applies to A/P, A/PF, APG, A/A-A, and ALUA arrays.

When all of the Primary paths fail or are disabled in a non-Active/Active array in a CVM cluster, the cluster-wide failover is triggered. All hosts in the cluster start using the Secondary path to the array. When the Primary path is enabled, the hosts fail back to the Primary path. However, suppose that one of the hosts in the cluster is shut down or brought out of the cluster while the Primary path is disabled. If the Primary path is then enabled, it does not trigger failback. The remaining hosts in the cluster continue to use the Secondary path. When the disabled host is rebooted and rejoins the cluster, all of the hosts in the cluster will continue using the Secondary path. This is expected behavior.

For A/P,APG, A/A-A, and ALUA arrays, if the disabled host is rebooted and rejoins the cluster before the Primary path is enabled, enabling the path does trigger the failback. In this case, all of the hosts in the cluster will fail back to the Primary path.

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 DMP restore daemon cycle to 60 seconds. The default value of this tunable is 300 seconds. The change is persistent across reboots.

Issue the following command at the prompt:

```
# vxddmpadm settune dmp_restore_internal=60
```

To verify the new setting, use the following command:

```
# vxdmpadm gettune dmp_restore_internal
```

Veritas File System software limitations

The following are software limitations in this release of Veritas File System

Linux I/O Scheduler for Database Workloads (1446361)

Symantec recommends using the Linux deadline I/O scheduler for database workloads on both Red Hat and SUSE distributions.

To configure a system to use this scheduler, include the `elevator=deadline` parameter in the boot arguments of the GRUB or LILO configuration file.

The location of the appropriate configuration file depends on the system's architecture and Linux distribution:

Configuration File	Architecture and Distribution
<code>/boot/grub/menu.lst</code>	RHEL5 x86_64, SLES10 x86_64, and SLES11 x86_64

For the GRUB configuration files, add the `elevator=deadline` parameter to the kernel command. For example, change:

```
title RHEL5UP3
  root (hd1,1)
  kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2
  initrd /boot/initrd-2.6.18-128.el5.img
```

To:

```
title RHEL5UP3
  root (hd1,1)
  kernel /boot/vmlinuz-2.6.18-128.el5 ro root=/dev/sdb2 \
  elevator=deadline
  initrd /boot/initrd-2.6.18-128.el5.img
```

A setting for the elevator parameter is always included by SUSE in its ELILO and GRUB configuration files. In this case, change the parameter from `elevator=cfq` to `elevator=deadline`.

Reboot the system once the appropriate file has been modified.

See the Linux operating system documentation for more information on I/O schedulers.

umount can hang when inotify watches are used (1590324)

If inotify watches are used, then an unmount can hang in the `vx_softcnt_flush()` call. The hang occurs because inotify watches increment the `i_count` variable and cause the `v_os_hold` value to remain elevated until the inotify watcher releases the hold.

Cached ODM

You can enable Cached ODM only for files on local file systems. It can not be enabled for files on Cluster File System.

Veritas Volume Replicator software limitations

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

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 will not be supported, because the IPv6-only node has no IPv4 address configured on it and so 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 will not be 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 will 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 5.1 and the prior major releases of Storage Foundation (5.0 or 5.0MP3). Replication between versions is

supported for disk group versions 120, 140, and 150 only. Both primary and secondary host must be using a supported disk group version.

Fixed issues

The following sections describe Storage Foundation issues that were fixed in this release.

See the *Veritas Cluster Server Release Notes* for VCS fixed issues.

Veritas Storage Foundation fixed issues

There are no Veritas Storage Foundation fixed issues in the 5.1 release.

Storage Foundation for Databases (SFDB) tools fixed issues

[Table 1-2](#) describes fixed issues in the Storage Foundation 5.1 release.

Table 1-2 Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Incident	Description
1728611	The DBED directories that are present in <code>/var/vx/vxdba</code> are no longer world writable.
1719617	The owners of OS_DIRECTORIES are not changed after installing SFDB package.
1535235	Allows user defined mount points for snapshot clones.
1535234	OEM storage plugins are usable across platforms.
1535232	Standby (dataguard) support for checkpoint flashsnap.
1535231	All DBED commands work from the CVM slave.
1535230	Extend support for multi-dg support with snapshot.
1531216	Fixed the <code>dbed_vmchecksnap</code> script error, if the "dco" volume did not have the string "_dco" in its name.
1525096	The DBED files that are present in <code>/tmp</code> are no longer world writable.
1525094	Changed the DBED repository database engine from Sybase ASA to SQLite.
1512850	Maked SFDB script IPV6 compliant.

Table 1-2 Veritas Storage Foundation for Databases (SFDB) tools fixed issues
(continued)

Incident	Description
1505177	Improved boot performance of DBED components.
1503356	The DBED scripts clean up correctly on failure.
1480379	Simplified the Storage Foundation for Databases packaging.
1361604	Fixed the usage examples in the dbdst_report manual page.
1121091	You do not need to add the _ora suffix to use the dbdst manual pages.
862687, 862092	Database FlashSnap clones supports DB_RECOVERY_FILE_DESTINATION as the sole mandatory archive log destination.

Veritas Volume Manager and Veritas Volume Replicator fixed issues

The Storage Foundation 5.1 includes cumulative fixes for Veritas Volume Manager and Veritas Volume Replicator since the 5.0MP3 release. The following tables describe these fixed issues.

[Table 1-3](#) describes fixed issues in the Veritas Volume Manager 5.1 release.

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues

Incident	Description
1822681	memory leak in vxio/voldr1_cleansio_start.
1822200	VRAS:Diff sync failed when remaining sync size is 4TB.
1819777	Panic issue in voldiosio_start() as race window exists while handling duplicate DA records
1810749	CR 6874695 - vxlustart -V deleted existing BEs
1805826	panic in vol_klog_clear_trans on Solaris x86
1804262	VVR:File system I/O of size bigger than 256k fails with error ENXIO after 2TB(>2G blocks)offset.
1797540	VxVM: vxdisk resize intermittently causes vxconfigd to dump core.
1795541	vxddladm disablevscli does not work for this customer

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
1782036	Sun: SC cannot open libvxvmisc.so after upgrade to 5.1
1779257	VVR:Disable Secondary logging through a tunable.
1765779	Man Page change for vxiod
1764972	vxdiskadm option 5 fails with "/usr/lib/vxvm/voladm.d/bin/disk.repl"
1762561	DMP: System panic when perform excludearray operation with powerpath
1762534	vxctl settz and vxconfigd core dump if TZ environment variable is not set.
1755869	tunable addition: gabmaxsend and receiver flowcontrol watermark
1755830	kmsg: sender: the logic for resend of messages needs to be optimized
1755810	kmsg: sender thread is woken up unnecessarily during flowcontrol
1755788	for a broadcast message, sender thread may end up sending the same message multiple times (not resend)
1755735	recovery I/Os get broken down to voliomem_chunk_size
1755707	vxtask list shows the same taskid for parent and child tasks
1755689	During recovery, -o delayrecover option does not work as expected for value of 0
1755628	kmsg layer: with heavy messaging in the cluster the receiver thread slows down processing
1755519	kmsg layer: receiver side flowcontrol is not supported
1755466	vol_find_ilock: searching of ilock is inefficient
1745992	CVR:I/O hang in 4 node CVR cluster
1744672	Oakmont::Primary slave hangs in volcvm_rvgrecovery_send_ioct() TC remote_write_reconfigure_2.tc
1744224	FMR3: multiple vxplex attach cmds running in parallel on a volume lead to clearing DCO map and subsequently lead to corruption
1742702	vxvmconvert fails, probably due to wrong disk capacity calculation

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
1739513	while mirroring 1 tb storage,after 800gb all VM commands hangs
1733811	System panic on voldco_isdirty code path while doing vxsnap make operation after upgrading from DCO version 10
1732200	[DMP][Usability] When NEW dmp_native_multipathing tunable is set to 'on' - unlabelled LUNs vanish from format until turned off
1728587	VVR: Replication started with a checkpoint remains inconsistent/cant_sync after SRL is drained if the replication is interrupted
1728269	Incorrect cur_pri_path updation for A/PG arrays leading to dmp database inconsistency
1725041	VVR: VRAS: vradm addsec fails with "V-5-52-502 Host name or IP XXX.YYY.ZZZ is not configured or available on the Primary."
1722984	Memory leak in vold_dg_get_clone_disks() .
1718008	Unable to initialize EFI LUNs controlled by EMC Powerpath driver, vxprtvtoc "Syntax Error" occurs.
1715889	Unable to encapsulate an unmanaged EMC DMX PP LUN
1711269	System hang showing kmsg receiver thread hogging CPU
1710030	AIX: Issue with LUN size more than TB
1678370	VM_VVR: RLINK disconnected and "vx" commands hung on Secondary while load in progress
1678292	[SxRT sparc/x64] vxdmpadm get tpdnodename error
1677217	DMP does not autofailback to the Primary paths following LCC card restoration.
1676061	System panic'd after 2 out of 4 paths to disk were removed.
1674847	Fixed an issue with vxconfigd not starting after a system reboot.
1673764	vxconfigd loses licensing information
1673002	Need to remove thousands of empty /tmp/vx.* directories.
1653972	VxVM volume device permissions change after running 'scgdevs' (Suncluster) command

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
1638494	VVR:vxnetd stop causing 100% CPU & vx commands hanging
1638174	oakmont:vxconfigd memory leak found
1637514	Issues with tentative evacuation of disks, when aborted in between using vxevac
1594928	Avoid unnecessary retries on error buffers when disk partition is nullified.
1592700	cvmvoldg monitor failed due to awk: Input line cannot be longer than 3,000 bytes - VRTScavf 5.0.31.0
1591146	mirrored volume grow doesn't works well if "mirror=enclosure" option is used, it leads to data corruption issue.
1589018	num_retries field is getting re-initialized to initial value leading to looping and delay in error handling time.
1588978	vxdmppadm getattr arrayname <array> partitionsize misreports DEFAULT partitionsize for subsequent entries
1555461	DMP did not handle the failure in one half of the customer SAN
1545999	Oracle hangs from time to time - VxVM 5.0 MP3 on RHEL5
1545835	vxconfigd core dump during system boot after VxVM4.1RP4 applied.
1543908	While running vxevac command, Oracle process thread stuck into ogetblk() which leads to i/o hang.
1541662	System panicked in DRL code when running flashsnap
1538053	CVM_MSG_REQ_GSLOCK repeatedly resent resulting in hang
1537821	VxVM private buffer's b_fsid field is not set properly, which breaks PRM
1531406	Race condition between Dynamic Reconfiguration thread and DMP error analysis code path lead to panic in gendmpiodone.
1530126	DMP : dmplinux_unplug() panic on linux, for no associated node in dmpnode
1521341	VRTSLSI APM 5.0-2.0 RHEL x86_64 rpm missing largeSMP support.
1517760	VRAS: vradmind core dump if stats collection is enabled

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
1510252	I/Omem chunks are not combined, causing panic at module unload time
1506690	VxVM utility scripts assume DA name is a valid path name.
1503309	VxVM: DMP doesn't log error returned from below driver/SCSI.
1483298	Man page for vxctl does not include all options that are listed in -H.
1480102	dmp disabled all paths of EVA8100 even though just unplugged primay FC cables
1475710	For VxVM patches the postinstall script must not copy /usr/lib/libthread to /etc/vx/slib/
1475703, 1369599	Export API libvxvm_get_disks() and libvxvm_get_subpaths() should return the same path for the same device
1475697	Provide an API for checking multi- ownership and status of a diskgroup
1475691	VxVM should recognize disks in use by ZFS or SVM
1471821	'initdmp' section in vxctl manpage still incorrect
1471784	[5.0MP3RP1 x64] vm can not create stripe-mirror/mirror-stripe/mirror volume with maxsize.
1471771	vxdisksetup should call 'mount' with complete pathname
1471581	vxconfigd may hang if no SCSI timeout value is set while issuing various SCSI commands
1471487	Critical Minimum Queue and Round-robin Improvements
1471263	machine has panicked when added the disk from dg as a foreign device using "vxdmpadm addforeign".
1470732	5.0MP3 : vxconfigd dumps core if all paths are excluded.
1470251	volslabd utilized 100% cpu time
1470102	vxdmpadm getattr for failoverpolicy and tpdmode is giving usage error.
1469351	User confused about snap operations when split to new DG
1468647	vxdmpdebug fails to find ugettxt

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
1463197	no path disable event occurs during I/O error analysis in dmp when pulling a FC cable out with 5.0MPP3
1460101	[VxVM]Proper handling of DMP PGR ioctl's in case of errors/warnings from below driver.
1443748	In a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load
1443046	System panic in vxio:voldrl_trans_copy
1441123	VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1437869	Need to examine package dependencies, especially wrt SUNWscpu...
1437006	DMP: Evaluation of DMP I/O statistics and TC development to verify correctness of the values displayed
1433535	DDL: Data corruption protection activated message should be reported by vxdisk scandisks and vxctl enable CLI
1430001	VM: VVR: Not able to modify any of the volume manager kernel tunable on 4.1MPP3 perf.
1425250	vx commands are hanging in EDC testing
1424194	vxclustadm reinit doesn't allow new nodes to join CVM cluster
1422008	After installed vm patch on AIX, "install-db" would be created if the vxio in "Defined" state.
1421752	High kernel CPU usage when DRL is turned on
1415547	Tool request - Verification tool to detect corrupted DCO.
1409031	VVR: Rlink fail to connect due to failure of memory allocation for incoming message buffer
1402599	VVR: Print warning message in syslog if it hit high water mark for latency protection
1398914	Support cdsdisk format on devices that do not support scsi modesense on page 3 and page 4.
1396566	Permissions set using vxedit is not consistent

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
1391652	No path disable event and I/O hang occurred when pulled out both paths of A/PG array with 4.0MP2RP8
1388324	Unresolved symbols in libvxvmc.so
1386592	tmp files not cleaned up after running vxdisksetup
1369610	Error message from vxdg import misleading/incorrect when disk is write-protected
1369597	vxlufinish uses wrong solaris command for unmounting alternate root environment - should use luumount
1361625	with use_all_paths=yes, get a reservation conflict
1321282	vxdisk scandisk hung in site failure scenario (EDC Testing)
1317186	Security: Remove tmp files in scripts before using them
1299512	incorrect vxdg free output and vxconfigd core dump
1293910	vxdmproot man page needs correction on arrays supported by DMP for SAN booting
1292633	PFTO value is set for only one paths though the DMP node has multiple paths.
1277808	VVR: Write testcase for link-breakoff snapshot with RVG.
1274204	vxbrk_rootmir fails if a volume name = end of another vol name
1235023	V-5-1-4597 vxdg join failed. Transaction aborted waiting for io drain
1230827	vxdisksetup init succeeds but both prtvtoc & vxdiskunsetup fail with error "No Such Device"
1224778	Write disable R2 device(SRDF)not seen after boot
1222625	VCS 5.x CVMCluster Agent doesn't handle non-C locales.
1212256	Panic in uphysdone due to double Iodone on the buffer
1183283	vxconfigstore -p returns with syntax error
1176510	Enhance vxdiskadm to handle custom-made rootdisk and rootmirror

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
1172961	AIX: VxVM should call uphysio() with a proper value for bufcnt argument
1153020	CVM master's behavior with dgfailpolicy=leave in 5.0 differs from the description of manual
1123203	vxconfigd hang in transaction commit while acquiring rwsleep write lock
1114870	System panic in voliorem_seek() due to wrong calculation of no of dco I/O extents.
1082431	During VCS During cluster shutdown of a node, vxconfigd died requiring a reboot
1082149	vxconvert command fails to convert for VPATH devices
1068259	VVR:Primary hang in case of TCP replication, secondary not sending ack
1060336	vxresize should not roll back if fsadm failed but disabled vxfs
1037289	IO Errors during SVC node add/remove test sequence
1020554	Diskgroup are not importing on every second or 3rd reboot. Needs to manually import it and mount the volumes
1016617	vxsnap refresh on SOSS fails with "VxVM vxvol WARNING V-5-1-10070 Failed to synchronize the volume(s)"
1012999	APIs needed for checking whether diskgroup is shared and status of diskgroup
996429	Unable to set no limit for cmd logs using vxcmdlog -n no_limit
990338	FMR Refreshing a snapshot should keep the same name for the snap object
963951	INSTSNAPTMP marked dco log not getting deleted during vxrecover or volume restart
818389	a scsi-2 release is causing a panic
795129	cmhaltcl on one of the cluster nodes causes one of the disk groups to lose shared flag

Table 1-3 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

Incident	Description
786357	Request to make voldrl_volumemax_drtregs a tunable
597517	unable to initialize EFI labeled >1tb PP devices
543638	vxddmpadm/vxdiskunsetup doesn't work well if tpdmode=native
339282	RFE : Failed to create more than 256 config copies in one DG
339187	CVM activation tag in vxprint -m output breaks vxprint
314738	vxdg split fails if the cvm master changes
298072	vxio reports "Illegal vminor encountered" even when there is no rootability
248925	If vxdg import returns error, parse it

Table 1-4 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP3 fixed issues

Incident	Description
1729558	Fixed issue with multiple vxplex attach commands running in parallel on a volume, which led to a clearing of the DCO map and subsequent corruption in FMR2.
1711339	VVR: Restored the ability to modify VVR tunables via kdb and expose tunables using vxtune.
1677416	Fixed CVM join and takeover issues during shared A/P storage config due to "not breaking" more than 64K size kmsgs.
1664891	Fixed issue with encapsulating linux root disk not changing the grub configuration.
1596811	Fixed an issue with slowness when re-enabling individual paths on vx4.1mp2rp1
1590314	Fixed issue with EMC PP/DMP interoperability for vxddmpadm getsubpaths dmpnodename=emcpower## incorrect details.
1589022	Fixed infinite looping issue in DMP error handling code path due to CLARIION APM leading to an IO hang.
1586879	Fixed issue with vxdisk online (on a single device as well as on a list of devices) taking too long on large configurations.

Table 1-4 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP3 fixed issues (*continued*)

Incident	Description
1537027	Fixed a security issue with <code>ddl_change_naming_scheme()</code> not setting mode when creating <code>.newnames</code> .
1534379	Fixed an issue of <code>vx dg</code> split failure with internal config daemon error - VxVM 5.0 MP1_rp5_hf13.
1529157	Fulfilled request to enhance <code>vx disk -e list</code> output.
1528115	Fixed ' <code>vx disk scandisks \!device=sdx</code> ' partial device discovery not working as expected.
1525819	Fixed an issue with <code>vxconfigbackup</code> not working on the <code>dg</code> with 2TB LUNs.9
1507291	Fixed an issue with <code>dmp_monitor_fabric=on</code> triggering unexpected offlining of paths on DMX4 array on 5.0MP3RP1.
1503168	Fixed an issue with the file <code>/etc/vx/darecs</code> not being read in 5.0MP3. <code>dg</code> causing import failure for <code>nopriv</code> disks.
1502120	Fixed an issue with <code>/dev/dmpconfig</code> device causing SELINUX an audit error.
1510204	Fixed an issue with <code>vxconfigd</code> command hanging after cluster nodes split simulation.
1503242	Fixed an issue with <code>vx dmpadm iostat</code> stats not tallying with those of <code>sar</code> , <code>iostat</code> and <code>vxstat</code> .
1487907	Fixed an issue with <code>vxio</code> has a non-unique UDID for failed disks.
1487888	Fixed an issue with the start up process fails after upgrade to 5.0 MP3.
1487584	Fixed an issue with CVR: I/O hang on logowner after logclient crashes.
1485379	Fixed an issue of <code>vxtask</code> showing Progress with "Progress: inf% 349978624 of 0 Blocks" while LINKSYNCing 34TB vol in SuSE 9.
1483643	Fixed an issue with the <code>raid5</code> volume not starting on Thin Provisioning LUNs of 3PAR.
1483164	Fixed an issue with disks with <code>nolabel</code> state being usable.
1479735	Fixed an issue with a CVR I/O hang on the slave if the master (logowner) crashes with DCM active.

Table 1-4 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP3 fixed issues (*continued*)

Incident	Description
1477143	Fixed an issue with vxconfigd sometimes hanging after pulling and re-attaching the fibre channel cable of the master node in clustered file system environments.
1475707	Added error message for unwritable disks.
1475580	Fixed an issue with VTOC getting corrupted by our header "PRIVHEAD".
1470734	Fixed an issue with vxconfigd dumps core if all paths are excluded.
1471658	Fixed an issue with AxRT5.0MP3: vxconfigd core in priv_get_all_udid_entry via vold_get_udid_entry via da_udid_is_any_same_disk.
1469487	Fixed issue with (bp)->start time being modified as part of error processing.
1461717	Fixed an issue with the vxsnap make command that caused vxconfigd and IO to sleep for too long a time.
1461314	Fixed disabling of SCSI bypass for single path disks (so that for path suppressing TPD, DMP does not use SCSI bypass).
1459831	VVR: Fixed replication hangs due to deadlock on secondary nmcom receive with TCP multiconnection.
1459000	Fixed an issue with the failover command on a bad LUN sometimes causing an infinite loop in dmpCLARiiON_issue_failover.
1457132	Fixed an issue with vxdmpadm enable path/ctrl CLI taking more time than they used to take prior to 5.0.
1450348	VVR: Fixed potential hang/panic due to the race between the RU thread and the volume read completion during a DCM replay.
1443752	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
1443706	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1442369	Fixed an issue with a bug in vxconfigbackupd script leading to 0 byte binconfig file being created.

Table 1-4 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP3 fixed issues (*continued*)

Incident	Description
1441131	Fixed an issue with VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1441072	Fixed an issue with siteread policy is not honoured.
1441020	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1435882	Fixed a panic in volkio_to_kio_copy due to passing null ptr to bcopy.
1435471	Fixed an issue with the cluster nodes panicking in voldco_or_pvmbuf_to_pvmbuf code after installing 5.0 MP3.
1431279	Fixed an issue with vxconfigf core dumps.
1428106	Fixed a system panic in vxio:voldr1_trans_copy.
1427284	Fixed an issue with vxdlmpadm dumped core when executing vxdlmpadm list dmpnode command.
1427267	Fixed a CVR panic in VOLSIOQ_MORE due to corrupted volsioq_start queue.
1425919	Fixed an issue with vxesd looping using 100% of one CPU.
1425338	Fixed an issue with a CVR failure to connect rlinks followed by vxconfigf hangs on secondary.
1421353	Fixed an issue with an I/O getting stuck in the drl_logbus queue due to corruption of the agenode LRU list.
1416080	Fixed an issue with system panic in the vol_change_disk() routine due to a NULL dereference.
1414469	Fixed an issue with DMP : vxddladm listsupport all not displaying up-to-date information.
1414441	The vxsnap manual page includes mirror=enclosure parameter to avoid being mirrored on the same enclosure.
1414381	Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery.
1414380	Fixed an issue with VVR Rlink failing to connect due to failure of memory allocation for incoming message buffer.

Table 1-4 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP3 fixed issues (*continued*)

Incident	Description
1414342	Fixed an issue with disk devices not appearing in vxdisk list.
1412785	Fixed an issue with the system hanging while creating volumes in a guest Ldom.
1411636	Fixed a secondary log error causing rlink disconnect after IBC unfreeze.
1408367	Fixed an issue with system panic in mutex_panic() being called from vol_rwsleep_wrlock().
1406185	Fixed a segmentation fault on x64 system when running the vxdmpadm list dmpnode all command.
1403370	Fixed a system panic after running the vxctl enable or vxconfigd -k commands.
1403123	Fixed an issue with vxconfigd sleeping and no vx commands were responding.
1402443	Fixed an issue with server panics in kmsg_udp_payload() 2- Node CVM cluster with SUN Cluster 3.2.0.
1397879	Enhanced the vxresize manual page to run from non-CVM master.
1397712	Fixed an issue with the vxsnap restore manual page is unable to properly freeze or thaw filesystems in a CVM environment.
1396427	Enhanced DMP to handle failing IO when it is not able to interpret sense data.
1389584	Fixed a system panic in vol_putdisk() code.
1389512	Fixed a system panic in bcopy() due to null passed in from voliocl_copyin().
1385922	Fixed a system panic due to memory allocation.
1382705	Fixed an issue with the vxdmpadm listexclude command dumping core when you run the VRTSexplorer tool.
1381783	Improved the performance of snapshot backups.
1376656	Fixed an issue with vxcached never deletes old snaps when cache hits HWM.

Table 1-4 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP3 fixed issues (*continued*)

Incident	Description
1372340	Fixed an issue with vxplex core dumps during vxassist addlog due to DRL log length being less than 33 blocks.
1321298	Fixed an issue with heartbeat after vxconfigd dumped the core following reconnection of fibre channel cable link.
1321272	Fixed an issue with vxcommands hanging after re-connecting the FC-site link.
1317361	Fixed an issue with the vxreattach and vxunreloc commands were leaving behind tmp files.
1287975	Fixed an issue with the vxclustadm command SIGSEGV's when main.cf contains lines with > 512 characters.
1266730	Fixed the vxtask command to display the resync progress subtask for shared volumes with DRL.
1253830	Fixed an issue with HP CCISS raid devices failing in vxvmconvert.
1230360	Fixed a system panic in vol_klog_start() due to accessing freed mv read_sio.
1225953	Improved I/O performance with VxFS filesystems on mirror-concat VxVM volumes with DCO and DRL.
1195591	Fixed an issue of three nodes rebooting due to customer's empty RVG.
1192105	Fixed the vxdg -n [newdg] deport [origdg] command causing a memory leak.
1184280	Fixed an issue with vxconfigd CONNECTION FAILED event due to VE_BADPROTOV error not providing sufficient information.
1182475	Fixed the vxdg split failing if the CVM master changes.
1108839	Turning on the dmp_cache_open tunable slows down vxconfigd when it is run with 2048 dual path LUNs.
1097258	Fixed an issue with vxconfigd hanging when an array is disconnected.
850816	Fixed an issue with parallel vxsnap reattach operations that can cause data corruption and also lead to orphan snap objects.
425273	Fixed an issue with VVR RU thread not starting nio after it is created from than waiting for all replicas to have NIO's created.

[Table 1-5](#) describes fixed issues in the Veritas Volume Manager 5.0MP3 RP1 release.

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.0MP3 RP1 fixed issues

Incident	Description
1470734	Fixed an issue with vxconfig dumps core if all paths are excluded.
1443752	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1441131	Fixed an issue with VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1441072	Fixed an issue with siteread policy is not honored.
1441020	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1435882	Fixed a panic in volkio_to_kio_copy due to passing null ptr to bcopy.
1435471	Fixed an issue with the cluster nodes panicking in voldco_or_pvmbuf_to_pvmbufcode after installing 5.0 MP3.
1431279	Fixed an issue with vxconfigd core dumps.
1428106	Fixed a system panic in vxio:voldrl_trans_copy.
1427284	Fixed an issue with vxdmpadm dumped core when executing vxdmpadm list dmpnode command.
1427267	Fixed a CVR panic in VOLSIOQ_MORE due to corrupted volsioq_start queue.
1425919	Fixed an issue with vxesdlooping using 100% of one CPU.
1425434	Fixed an issue with CVR fails to connect rlinks followed by vxconfigd hangs on secondary.
1414441	The vxsnap manual page includes mirror=enclosure parameter to avoid being mirrored on the same enclosure.
1414381	Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery.

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.0MP3 RP1 fixed issues (*continued*)

Incident	Description
1414380	Fixed an issue with VVR Rlink failing to connect due to failure of memory allocation for incoming message buffer.
1412785	Fixed an issue with the system hanging while creating volumes in a guest Ldom.
1411636	Fixed a secondary log error causing rlink disconnect after IBC unfreeze.
1406185	Fixed a segmentation fault on x64 system when running the vxdmpadm list dmpnode all command.
1403370	Fixed a system panic after running the vxctl enable or vxconfigd -k commands.
1403123	Fixed an issue with vxconfigd sleeping and no vx commands were responding.
1397879	Enhanced the vxresize manual page to run from non-CVM master.
1397712	Fixed an issue with the vxsnap restore manual page is unable to properly freeze or thaw file systems in a CVM environment.
1396427	Enhanced DMP to handle failing IO when it is not able to interpret sense data.
1389584	Fixed a system panic in vol_putdisk() code.
1389512	Able to force import diskgroup version 80 in VxVM 5.0.
1387033	Fixed a system panic in bcopy() due to null passed in from volioctl_copyin()
1385922	Fixed a system panic due to memory allocation.
1382705	Fixed an issue with the vxdmpadm listexclude command dumping core when you run the VRTSexplorer tool.
1381783	Improved the performance of snapshot backups.
1376656	Fixed an issue with vxcached never deletes old snaps when cache hits HWM.
1368752	Fixed an issue when there are no mirrors to read, VOL_READ_MIRRORS ioctl returns -1 instead of 1.

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.0MP3 RP1 fixed issues (*continued*)

Incident	Description
	Fixed an issue with VSCSI: A/P LBI/O policy not working with enabled DMP support on boot devices.
1317361	Fixed an issue with the vxreattach and vxunreloc commands were leaving behind tmp files.
1314301	Fixed an issue with vxlustart.
1266730	Fixed the vxtask command to display the resync progress subtask for shared volumes with DRL.
1253380	Fixed an issue with HP CCISS raid devices failing in vxvmconvert.
1230360	Fixed a system panic in vol_klog_start() due to accessing freed mv read_sio.
1225953	Improved I/O performance with VxFS filesystems on mirror-concat VxVM volumes with DCO and DRL.
1192105	Fixed the vxdg -n [newdg] deport [origdg] command causing a memory leak.
1182475	Fixed the vxdg split failing if the CVM master changes.
425273	Fixed an issue with VVR RU thread not starting nio after it is created from than waiting for all replicas to have NIO's created.

[Table 1-6](#) describes fixed issues in the Veritas Volume Manager 5.0MP3 RP2 release.

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues

Incident	Description
1502120	Fixed an issue with /dev/dmpconfig device causing SELINUX an audit error.
1510204	Fixed an issue with vxconfigd command hanging after cluster nodes split simulation.
1503242	Fixed an issue with vxdmpadmiostatstats not tallying with those of sar, iostatand vxstat.
1487907	Fixed an issue with vxiohas a non-unique UDID for failed disks.

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1487888	Fixed an issue with the start up process fails after upgrade to 5.0 MP3.
1487584	Fixed an issue with CVR: I/O hang on logowner after logclient crashes.
1475580	Fixed an issue with VTOC getting corrupted by our header "PRIVHEAD".
1470734	Fixed an issue with vxconfig dumps core if all paths are excluded.
1443752	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
1443706	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1442369	Fixed an issue with a bug in vxconfigbackupscript leading to 0 byte binconfig file being created.
1441131	Fixed an issue with VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1441072	Fixed an issue with siteread policy is not honored.
1441020	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1435882	Fixed a panic in volkio_to_kio_copydue to passing null ptr to bcopy.
1435471	Fixed an issue with the cluster nodes panicking in voldco_or_pvmbuf_to_pvmbufcode after installing 5.0 MP3.
1431279	Fixed an issue with vxconfig core dumps.
1428106	Fixed a system panic in vxio:voldrl_trans_copy.
1427284	Fixed an issue with vxdumpadm dumped core when executing vxdumpadm list dmpnode command.
1427267	Fixed a CVR panic in VOLSIOQ_MORE due to corrupted volsioq_start queue.
1425919	Fixed an issue with vxesdlooping using 100% of one CPU.
1425434	Fixed an issue with CVR fails to connect rlinks followed by vxconfig hangs on secondary.

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1414441	The vxsnap manual page includes mirror=enclosure parameter to avoid being mirrored on the same enclosure.
1414381	Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery.
1414380	Fixed an issue with VVR Rlink failing to connect due to failure of memory allocation for incoming message buffer.
1414342	Fixed an issue with disk devices not appearing in vxdisklist.
1412785	Fixed an issue with the system hanging while creating volumes in a guest Ldom.
1411636	Fixed a secondary log error causing rlink disconnect after IBC unfreeze.
1406185	Fixed a segmentation fault on x64 system when running the vxdmpadm list dmpnode all command.
1403370	Fixed a system panic after running the vxctl enableor vxconfigd -k commands.
1403123	Fixed an issue with vxconfigd sleeping and no vx commands were responding.
1397879	Enhanced the vxresize manual page to run from non-CVM master.
1397712	Fixed an issue with the vxsnap restoremanual page is unable to properly freeze or thaw filesystems in a CVM environment.
1396427	Enhanced DMP to handle failing IO when it is not able to interpret sense data.
1389584	Fixed a system panic in vol_putdisk() code.
1389512	Able to force import diskgroup version 80 in VxVM 5.0.
1387033	Fixed a system panic in bcopy() due to null passed in from voliectl_copyin()
1385922	Fixed a system panic due to memory allocation.
1382705	Fixed an issue with the vxdmpadm listexclude command dumping core when you run the VRTSexplorer tool.
1381783	Improved the performance of snapshot backups.

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1376656	Fixed an issue with vxcached never deletes old snaps when cache hits HWM.
1372340	Fixed an issue with vxplex core dumps during vxassist addlog due to DRL log length being less than 33 blocks.
1317361	Fixed an issue with the vxreattachand vxunreloc commands were leaving behind tmp files.
1266730	Fixed the vxtaskcommand to display the resync progress subtask for shared volumes with DRL.
1253830	Fixed an issue with HP CCISS raid devices failing in vxvmconvert.
1230360	Fixed a system panic in vol_klog_start() due to accessing freed mv read_sio.
1225953	Improved I/O performance with VxFS filesystems on mirror-concat VxVM volumes with DCO and DRL.
1192105	Fixed the vx dg -n [newdg] deport [origdg] command causing a memory leak.
1182475	Fixed the vx dg split failing if the CVM master changes.
425273	Fixed an issue with VVR RU thread not starting nio after it is created from than waiting for all replicas to have NIO's created.

Veritas File System fixed issues

The Storage Foundation 5.1 includes cumulative fixes for Veritas File System since the 5.0MP3 release. The following tables describe these fixed issues.

[Table 1-7](#) describes fixed issues in the Veritas File System 5.1 release.

Table 1-7 Veritas File System 5.1 fixed issues

Incident	Description
1074101	The SLES 10 boot sequence no longer displays error messages due to Veritas File System entries in <code>/etc/fstab</code> .
1411871	NFS file locks are no longer broken with the RHEL 5 Update 2 <code>nfsd</code> or <code>lockd</code> daemons for exports of VxFS.

Table 1-7 Veritas File System 5.1 fixed issues (*continued*)

Incident	Description
1477763	The <code>qiostat -l</code> command now shows accurate hit percentages.
1504816	Optimized the <code>vx_convnodata_files()</code> call.
1603264	Fixed a panic that occurred while removing Storage Checkpoints, which requires a very long full <code>fsck</code> .
1664173	Fixed an issue in which the file system is getting full faster than files could be deleted, which caused the following error to display: <code>vxfs: msgcnt 5415 mesg 001: V-2-1: vx_nospace</code>
1672389	The <code>dk -k</code> command now shows the correct space used via Veritas Enterprise Administrator.
1726287	Fixed the cause of a hang in the <code>vx_ireuse()</code> call that was due to <code>inotify</code> watches.
1744587	Fixed a mount failure that occurred while mounting file systems that are greater than 10 TB.
1745875	Fixed a hang issue with the <code>fsckptadm create</code> command.
1795939	Fixed the "Invalid argument" error displayed by the <code>hagetcf</code> command during configuration due to GLM.
1804512	Fixed a <code>umount</code> failure that returned EBUSY.
1805951	Added 64-bit libraries in the publicly available <code>VRTSfssdk</code> package to resolve issues with the VxFS libraries in NetBackup 7.0.
1806756, 1820496	Fixed the cause of the <code>sysctl</code> command displaying the following error: <code>"Invalid argument" reading key "vximc.vxgms"</code>

[Table 1-8](#) describes fixed issues in the Veritas File System 5.0 MP3 RP2 release.

Table 1-8 Veritas File System 5.0 MP3 RP2 fixed issues

Incident	Description
1411872	NFS locking fix on RHEL 5.
1435656	Made multi-threading improvements to the <code>fsclustadm</code> library

Table 1-8 Veritas File System 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1456149	Improvements when allocating odd sized extents.
1456156	Potential for CFS hang during file read ahead.
1466362	Changed the minimum exclusion zone size for better performance with Thin Reclamation.
1488246	Fixed up directory entry bounds checking which could cause files not to be removed.

[Table 1-9](#) describes fixed issues in the Veritas File System 5.0 MP3 RP1 release.

Table 1-9 Veritas File System 5.0 MP3 RP1 fixed issues

Incident	Description
1323920	Fixed the <code>odmmkfile</code> and <code>qiomkfile</code> to only run as root user.
1385559	Fixed a policy enforcement when converting a “data” checkpoint to a “nodata” checkpoint.
1385560	Fixed an issue when reusing a special-device inodes can cause memory corruption.
1404722	Corrected the <code>vxumount</code> command’s usage text.
1407592	Added support for a fake mount of a file system.
1412163	Fixed the cause of a core dump by a VxFS function call while setting Dynamic Storage Tiering attributes.
1412583	Fixed the cause of a core dump if the <code>chkptname</code> parameter was omitted when running the <code>fsapadm enforceckpt</code> command.
1413464	Changed the maximum direct I/O size to 1 MB.
1414719	The <code>du</code> and <code>bdf</code> commands now report the same amount of free space without a delay after files are deleted.
1424181	Improved performance when converting a data Storage Checkpoint to a nodata Storage Checkpoint.
1428641	Fixed an issue with shrinking a CFS which is over 69% full can cause temporary performance issues.
1459277	Fixed an issue in which Thin Storage could not be reclaimed.

Known issues

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

See the *Veritas Cluster Server Release Notes* for VCS known issues.

Veritas Storage Foundation known issues

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

5.1 installer does not allow direct upgrade from older version of AT to 5.1 SFHA product stacks

The 5.1 installer does not allow upgrading between products and upgrading versions in a single upgrade operation. If an older version of AT is already installed on the systems by some other Veritas product, installing any Storage Foundation High Availability product fails with the following message:

```
The following errors were discovered on the systems:  
AT 4.x.x.x is installed. Upgrading AT 4.x.x.x directly to SF 5.1  
is not supported.  
Product Installed - Product Version - System Name  
AT                - 4.x.x.x          - d08b  
First run the install<AT> script to upgrade product to AT 5.1  
and then run install<SF>
```

Workaround

You must first upgrade AT to the latest version of AT, then you can install the desired SFHA 5.1 product. For example, to upgrade from AT 4.x to SFHA 5.1, first upgrade to AT 5.0, and then install SFHA 5.1.

To upgrade AT and SFHA products

- 1 Run the `installat` script, located under the `authentication_service` product directory in the media.
- 2 After the successful upgrade of AT, use the installer script to install the desired SFHA 5.1 product.

Veritas Storage Foundation for Databases (SFDB) tools known issues

The following are known issues with Storage Foundation for Databases (SFDB) tools.

Clone command fails on an Oracle RAC database (1399393)

The commands `dbed_vmclondb` and `dbed_clonedb` will fail on an Oracle RAC database when the clone SID name contains the primary SID name in the beginning.

For example, the following commands, which have “Prod” as the primary SID and “Prod1” as the clone SID, produce the following error message:

```
# dbed_vmclondb -S Prod -o recoverdb \  
new_sid=Prod1,server_name=srv_name -f snapplan -r relocate_path  
# dbed_vmclondb -S Prod -o mountdb \  
new_sid=Prod1,server_name=srv_name -f snapplan -r relocate_path  
# dbed_clonedb -S Prod1 -m mount_point -c ckpt_name  
ERROR V-81-4882 An error occurred while reconfiguring Oracle instance  
'clone_SID'
```

Workaround

Do not use a clone SID name that contains primary SID name in the beginning.

Database fails over during Flashsnap operations (1469310)

In an SFHA environment, if the database fails over during Flashsnap operations such as the `dbed_vmsnap -o resync` command and various error messages appear. This issue occurs because Flashsnap commands do not create a VCS resource for the SNAP disk group. As such, when the database fails over, only the primary disk group is moved to another node.

Workaround

There is no workaround for this issue.

The error messages depend on the timing of the database failover. To fix the problem, you need to bring the FlashSnap state to `SNAP_READY`. Depending on the failure, you may have to use base VxVM commands to reattach mirrors. After mirrors are attached, you need to wait until the mirrors are in `SNAPDONE` state. Re-validate the snapplan again.

Clone command fails for instant checkpoint on Logical Standby database (1736516)

The `dbed_clonedb` command does not work on Logical standby database when using instant checkpoint.

Workaround

Clone the database using an alternate kind of checkpoint, such as the online checkpoint, which is more commonly used.

Clone command fails if archive entry is spread on 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.

dbed_vmclonedb -o recoverdb for offhost fails for Oracle 10gr2 and prior versions (1789290)

When performing an offhost clone for Oracle 10gr2 and prior versions, the clone operations fails. The failure occurs because the `user_dump_dest`, `background_dump_dest`, and `core_dump_dest` parameters are defined in the `pfile` or `spfile` file of the primary database, but the paths specified for these parameters do not exist on the offhost machine.

Workaround

Before performing the offhost clone operation, create the directory paths on offhost machine that are specified for the `user_dump_dest`, `background_dump_dest`, and `core_dump_dest` parameters in the `pfile` or `spfile` file of the primary database.

Storage Foundation for Databases (SFDB) tools support for using Oracle Data Guard with Oracle RAC (1801265)

Storage Foundation for Databases (SFDB) tools do not support use of Data Guard with Oracle RAC in this release.

Workaround

There is no workaround for this issue.

Flashsnap reverse resync command fails on offhost flashsnap cloning (1810711)

Performing `dbed_vmsnap -o reverse_resync_begin` fails if an off-host clone has been created on the snapshot before.

Workaround

Use database checkpoints for this use case.

dbed_vmclonedb displays errors when attempting to take a clone on a STANDBY database in a dataguard environment when you are using the MEMORY_TARGET feature for Oracle 11g (1824713)

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

```
dbed_vmclonedb started at 2009-08-26 11:32:16
Editing remote_login_passwordfile in initclone2.ora.
Altering instance_name parameter in initclone2.ora.
Altering instance_number parameter in initclone2.ora.
Altering thread parameter in initclone2.ora.
SFORA dbed_vmclonedb ERROR V-81-4918 Database clone2 has not been
correctly recovered.
SFORA dbed_vmclonedb ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.20569/recover.log.
```

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

Workaround

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

To resolve this known issue

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

```
# umount /dev/shm
```
- 3 Mount the `/dev/shm` file system with the following options:

```
# mount -t tmpfs shmfs -o size=4096m /dev/shm
```
- 4 Start the database.

Reattach command fails in a multiple disk group environment if the snapshot operation fails (1840672)

In a multiple disk group environment, if the snapshot operation fails then `dbed_vmsnap` fails to reattach all the volumes. This operation must be performed as root user.

Workaround

In case the reattach operation fails, use the following steps to reattach the volumes.

To reattach volumes in a multiple disk group environment if the snapshot operation fails

- 1 Join the snapshot disk groups to primary disk groups. The snapshot disk group name is a concatenation of “SNAPSHOT_DG_PREFIX” parameter value in `snapplan` and primary disk group name. Use the following command to join the disk groups:

```
# vxvg join snapshot_disk_group_name primary_disk_group_name
```

- 2 Start all the volumes in primary disk group.

```
# vxvol -g primary_disk_group_name startall
```

- 3 Reattach the snapshot volumes with primary volumes. The snapshot volume name is a concatenation of “SNAPSHOT_VOL_PREFIX” parameter value in `snapplan` and primary volume name. Use the following command to reattach the volumes.

```
# vxsnap -g primary_disk_group_name reattach snapshot_volume_name  
source=primary_volume_name
```

Repeat this step for all the volumes.

Database FlashSnap `snapplan` creation and validation failed for standby database if standby redo log is mislocated (1873738)

Mislocated standby redo log files may cause errors during Database FlashSnap `snapplan` creation and validation for standby database in a Oracle Data Guard environment.

The `dbed_vmchecksnap` command may fail with the following error messages:

```
$ dbed_vmchecksnap -S stand \  
-H $ORACLE_HOME -f snp -o setdefaults -t dblxxeon02  
SFORA vxsnapadm ERROR V-81-5907 open() /snap_data11r1/FLAS11r1/redo01.log  
failed (No such file or directory).
```

```
SFORA vxsnapadm ERROR V-81-5526 Empty or open file
      /snap_data11r1/FLAS11r1/redo01.log failed.
Snapplan snp for stand.
```

Workaround

Create empty redo log files by using the touch command so that `dbed_vmchecksnap` can succeed without any errors.

For example:

```
$ touch /snap_data11r1/FLAS11r1/redo01.log
$ touch /snap_data11r1/FLAS11r1/redo02.log
$ touch /snap_data11r1/FLAS11r1/redo03.log
```

db2exp may frequently dump core (1854459)

If a host is configured to an SFM central server with DB2 version 9.x, then the command-line interface `db2exp` may frequently dump core.

Workaround

There is a hotfix patch available for this issue. Contact Symantec Technical Support for the hotfix patch.

In an IPv6 environment, db2icrt and db2idrop commands return a segmentation fault error during instance creation and instance removal (1602444)

When using IBM DB2 `db2icrt` command to create a DB2 database instance on a pure IPv6 environment, the `db2icrt` command returns segmentation fault error message. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2icrt -a server -u db2fen1 db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 4700: 26182 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -addfcm -i ${INSTNAME?}
```

The `db2idrop` command also returns segmentation fault, but the instance is removed successfully after the `db2idrop` command is issued. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2idrop db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 3599: 7350 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -remove -s DB2_${INSTNAME?} 2> /dev/null
```

```
DBI1070I Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

This issue has been identified as an IBM issue. Once IBM has fixed this issue, then IBM will provide a hotfix for this segmentation problem.

At this time, you can communicate in a dual-stack to avoid the segmentation fault error message until IBM provides a hotfix.

To communicate in a dual-stack environment

- ◆ Add an IPv6 hostname as an IPv4 loopback address to the `/etc/hosts` file. For example:

```
127.0.0.1 swlx20-v6
```

Or

```
127.0.0.1 swlx20-v6.punipv6.com
```

where `127.0.0.1` is the IPv4 loopback address.

where `swlx20-v6` and `swlx20-v6.punipv6.com` is the IPv6 hostname.

Process start-up may hang during configuration using the installer (1678116)

After you have installed a Storage Foundation product, some VM process may hang during the configuration phase.

Workaround

Kill the installation program, and rerun the configuration again.

Oracle 11gR1 may not work on pure IPv6 environment (1819585)

There is problem running Oracle 11gR1 on a pure IPv6 environment.

Tools like `dbca` may hang during database creation.

There is no workaround for this as Oracle 11gR1 does not fully support pure IPv6 environment. Oracle 11gR2 release may work on a pure IPv6 environment, but it has not been tested or released yet.

Not all the objects are visible in the SFM 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 Deployed 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
```

An error message when you perform off-host clone for RAC and the off-host node is not part of the CVM cluster (1834860)

There is a known issue when you try to perform an off-host clone for RAC and the off-host node is not part of the CVM cluster. You may receive a similar error message:

```
Cannot open file /etc/vx/vxdba/rac11g1/.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'rac11g1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdg ERROR V-5-1-582 Disk group SNAP_rac11dgl: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dgl.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dgl failed.
```

Currently there is no workaound for this known issue. However, if the off-host node is part of the CVM cluster, then off-host clone for RAC works fine.

Also the `dbed_vmclonedb` command does not support `LOCAL_LISTENER` and `REMOTE_LISTENER` in the `init.ora` parameter file of the primary database.

DB2 databases are not visible from the SFM Web console (1850100)

If you upgraded to SF 5.1, DB2 databases will be not visible from the SFM web console.

This will be fixed in the SF 5.1 Patch 1 release.

Workaround

Reinstall is required for SFM DB2-Hotfix (HF020008500-06.sfa), if the host is upgraded to SF 5.1 Use the deployment framework and reinstall the hotfix for DB2 (HF020008500-06.sfa) on the managed host.

To resolve this issue

- 1 In the Web GUI, go to **Settings > Deployment**.
- 2 Select **HF020008500-06 hotfix**.
- 3 Click **Install**.
- 4 Check the **force** option while reinstalling the hotfix.

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 dynamic storage tiering (DST) 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.

On SLES11, when you upgrade from 5.0RU1 to 5.1, the installation script may hang (1865666)

When you upgrade from Storage Foundation 5.0RU1 to 5.1, the `sfinstall` script may hang. This occurs while `sfinstall` is trying to stop the `vxspec` driver.

Workaround

Before you upgrade to Storage Foundation 5.1, deport your disk groups. After the upgrade, import the disk groups back to the system.

Veritas Volume Manager known issues

The following are known issues in this release of Veritas Volume Manager (VxVM).

Required attributes for Veritas Volume Manager (VxVM) devices to avoid boot failures (1411526)

To support iSCSI devices, Veritas Volume Manager (VxVM) does not start non-root devices until runlevel2. The boot process expects all local (non-NFS) mount points in the `/etc/fstab` file to be present at boot time. To avoid boot failures, all VxVM entries in the `/etc/fstab` file must have the `_netdev` attribute, and must not have

the `fsck` required flag set. These attributes enable VxVM to defer mounting of VxVM devices until after VxVM has started.

Performing Thin Reclamation on sliced format disk causes data corruption (1834848)

This issue occurs only for the Thin Reclamation feature on a sliced format disk. When reclaiming the thin disks within the disk group, this issue occurs if the thin disks are initialized using sliced format. The `vxdisk reclaim` command or `fsadm -R` command option corrupts the data on the file system.

Workaround

Avoid performing any reclaim operation on the disk group with thin reclamation disk on it.

If performing the reclaim is required, initialize the thin reclaim disk as `cdsdisk` format.

Existing disk groups and file systems will not be affected, because the thin reclamation feature is only supported from Storage Foundation 5.0MP3 release and onwards.

The path failure due to filer panic in the netapp array may occasionally cause longer path failover time or longer I/O delay (1835139)

The path failure due to filer panic in the netapp array may occasionally cause longer path failover time or the longer I/O delay which varies from 100 seconds to 400 seconds. The long failover time is due to the race between the I/O code path and the path failover code path, which is more likely to happen with heavy I/O load.

Workaround

There is no workaround for this issue. The poor path failover time is not avoidable at this time.

Veritas Volume Manager 5.1 supports only RHEL 5 Update 3

If you are using RHEL 5, support is limited to RHEL 5 Update 3. This is due to different kernel levels between Update 2 and Update 3. The VxVM 5.1 modules load only with the Update 3 kernel.

Limitation of base minor numbers (130558)

Previous Linux releases, including RHEL 4 U4 and SUSE 9 SP3, did not support minor numbers greater than 255 in an NFS implementation. This limitation was due to a bug in NFS, Red HAT Bugzilla Bug 143897 or SUSE Bugzilla Bug 64552. Volume devices with large minor number could not be remotely mounted on NFS unless a patch was applied.

In previous versions of VxVM, a work around was implemented, so that the `vxpdg` command changed the base minor of the disk group to a small number (such as 100). The volume was allocated with a minor smaller than 255.

In VxVM 5.1, the base minor numbers of shared disk groups and private disk groups are separate ranges. By default, the base minor numbers of shared disk groups cannot be smaller than 33000. Therefore, the `vxpdg` command only applies the work around in the case of private disk groups.

The recommended solution is to install the NFS patch for Linux so that minor numbers greater than 255 can be used. If the NFS patch cannot be installed, the workaround is to disable the division between minor numbers for shared disk groups and private disk groups. To disable the division between minor numbers, set the tunable value `sharedminorstart=0` in the `/etc/default/vxsf` file. You can then change the shared minor to a number smaller than 255.

Manually upgrading rpms from previous releases of Veritas Volume Manager (1590300)

In this release of Veritas Volume Manager, the `VRTSvxvm-common` and `VRTSvxvm-platform` rpms were consolidated to the `VRTSvxvm` rpm.

If you upgrade with the Veritas installer or installation scripts, the rpms are upgraded correctly. This is the recommended method.

If you upgrade any Storage Foundation products with the `rpm` command, the `rpm` command does not correctly handle the upgrade because the old packages have different names. The `rpm` command passes the wrong parameters to the uninstallation scripts. To avoid this problem, use the following version of the `rpm` command to upgrade the rpm:

```
# rpm -U --nopreun --nopostun VRTSvxvm-5.1.OS*.x86_64.rpm
```

Where * is the platform architecture: rhel5, sles10 or sles11.

Host panics once the other node reboots in a 2-node cluster (1603405)

While running PowerPath on SLES 10 SP2, the host panics once the second node of a 2-node cluster reboots.

Workaround

Upgrade PowerPath to version 5.3.1.00.00-103.sles10sp2.

vxrestored daemon fails to restore disabled paths (1663167)

The `vxrestored` daemon fails to restore disabled paths on RHEL 5 with direct attached disks.

Workaround

Enable the `mpt_disable_hotplug_remove` tunable so that path level failover and failback function properly on RHEL 5 machines with direct attached disks.

To enable the `mpt_disable_hotplug_remove` tunable

- 1 Edit the `/etc/modprobe.conf` file and add the following line to the end of the file:

```
options mptsas mpt_disable_hotplug_remove=0
```

- 2 Rebuild the `initrd` image:

```
# mkinitrd -f /boot/initrd-`uname -r`.img `uname -r`
```

- 3 Reboot the system.

DMP cannot detect the re-enabled OS device status on SLES 11 after restoring the daemon (1718573)

Because the `remove_on_dev_loss` parameter of the `scsi_transport_fc` module is removed in SLES 11, the OS device files are removed after a device loss with the `dev_loss_tmo` parameter. When the device comes back online, the port names may have changed, in which case DMP cannot recognize the device's status with the restored daemon.

Workaround

Set the `dev_loss_tmo` parameter to 8000000. This workaround only works with QLogic drivers, as Emulex drivers are not supported on SLES 11.

System hangs or panics after disabling 3 of 4 arrayside ports (1724260)

The system hangs or panics after you disable 3 of 4 arrayside ports.

Workaround

This issue is fixed with a Novell patch for SLES 11 as indicated in Bugzilla ID 524347:

https://bugzilla.novell.com/show_bug.cgi?id=524347

DMP cannot detect LUNs on a node that has an Emulex HBA after rebooting the system (1740205)

External LUNs sometimes are not visible to DMP after rebooting if the Emulex HBA driver loads after VxVM starts up.

Workaround

Ensure that the host adapter driver is loaded in the `initrd` image, which allows VxVM to discover the attached storage at boot time.

SCSI 0x00070000 errors fill the system messages and lead to slow system response (1802485)

SCSI 0x0007 errors fill the system messages and lead to slow system response while working with an HDS AMS500/AMS1000 or HDS USPV array that has multiple active paths.

Workaround

While working with an HDS AMS/USPV array, set the `queue_depth` parameter to a value no higher than 8.

Machine fails to boot after root disk encapsulation on servers with UEFI firmware (1842096)

Certain new servers in the market such as IBM x3650 M2, Dell PowerEdge T610, come with support for the UEFI firmware. UEFI supports booting from legacy MBR type disks with certain restrictions on the disk partitions. One of the restrictions is that each partition must not overlap with other partitions. During root disk encapsulation, it creates an overlapping partition that spans the public region of the root disk. If the check for overlapping partitions is not disabled from the UEFI firmware, then the machine fails to come up following the reboot initiated after running the commands to encapsulate the root disk.

Workaround

- For the IBM x3650 series servers, the UEFI firmware settings should be set to boot with the "Legacy Only" option.
- For the Dell PowerEdge T610 system, set "Boot Mode" to "BIOS" from the "Boot Settings" menu.

The above workarounds have been tested and are recommended in a single-node environment.

Boot time is unusually long on RHEL 5 Update 3 with DS4K arrays (1535260)

The boot time is unusually long on RHEL 5 Update 3 when using DS4K arrays.

Workaround

On RHEL 5 Update 3, you can reduce the boot time by enabling the RDAC module.

To reduce the boot time

- 1 Disconnect all paths connected to the array.
- 2 Run the `mkinitrd` command on an RHEL5 U3 system to include the `scsi_dh_rdac` module:

```
# mkinitrd target_initrd_image kernel_version \  
--preload=scsi_dh_rdac
```

For example:

```
# mkinitrd /boot/my_image `uname -r` --preload=scsi_dh_rdac
```

- 3 Redirect your boot loader to use the new `initrd` image. Replace the `initrd` with the new `initrd` image in the GRUB configuration `/boot/grub/menu.lst` file.

For example, change:

```
initrd /boot/initrd-2.6.16.img
```

To:

```
initrd /boot/my_image
```

where `my_image` is what you specified in step 2

- 4 Connect all paths to the array.
- 5 Reboot the system.

The requestleave policy is not supported if disk cache objects (DCOs) or dirty region logs (DRLs) are configured (1796617)

When disk cache objects (DCOs) or dirty region logs (DRLs) are configured, the `dgfailpolicy` should not be set to `requestleave`. This release 5.1 does not support this scenario because it could result in data corruption or compromise application availability.

Error messages display while running the VRTSexplorer tool (1445897)

While running the VRTSexplorer tool on 5.1, the following error messages may display on the console:

```
VRTSexplorer: Collecting VRAS configuration information.  
VOL_KERN_DUMP ioctl failed: Bad address  
VOL_KERN_DUMP ioctl failed: Bad address  
VOL_KERN_DUMP ioctl failed: Bad address
```

These messages are harmless and may be ignored.

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

VxVM might detect and report a false serial split brain when all 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, run the following command:

```
/etc/vx/diag.d/vxprivutil set <device-path> ssbid=<dm id>
```

The `<dm id>`, which is also the `ssbid`, can be retrieved from the config copy by running the following command:

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

Root disk encapsulation issue (1603309)

Encapsulation of root disk will fail if it has been assigned a customized name with `vxddmpadm(1M)` command. If you wish to encapsulate the root disk, make sure that you have not assigned a customized name to its corresponding DMP node.

See `vxddmpadm(1M)` and the section "Setting customized names for DMP nodes" on page 173 for details.

Remove iSCSI dependency in `vxvm_boot` (1854642)

In some instances, VxVM is unable to start up at boot time if iSCSI service has not started or is not present on the system.

Workaround

Edit the `/etc/init.d/vxvm-boot` file to remove the `iscsi` entry against "Required-Start" keyword. Then, reboot the system.

VxVM starts before OS device scan is done (1635274)

While working with some arrays, VxVM may start before all devices are scanned by the OS. This slow OS device discovery may result in malfunctioning of VM, fencing and VCS due to partial disks seen by VxVM.

Workaround

After the fabric discovery is finished, issue the `vxddisk scandisks` command to bring newly discovered devices into the VxVM configuration.

A physical LUN with large number of cylinders cannot be used as a CDS disk (1846165)

A physical LUN that has number of cylinders greater than 65536 can lead to small data corruption if used as a CDS disk. CDS disk uses a label such that the number of cylinders gets truncated. Backup labels are written to truncated offsets which leads to a single sector corruption in several places.

Workaround

Use sliced disks for any LUN exhibiting above behavior.

<http://entsupport.symantec.com/docs/336650>

Dynamic Lun Expansion of a sliced disk can fail to resize (1850166)

Dynamic Lun Expansion of a sliced disk can fail to resize. This issue occurs because cylinder size is not kept constant for the resize. This issue occurs only with sliced disks

Workaround

Use CDS disks.

After installing Volume Manager, you may be prompted to reinstall it (1704161)

If you remove pre-5.1 Volume Manager packages and then install 5.1 Volume Manager without using the product installer, the following message is displayed:

```
The Volume Manager appears to be installed already. You should use vxdiskadm to add more disks to the system. Installation with vxinstall will attempt to reinstall the Volume Manager from the beginning. Depending upon how your system is currently configured, a reinstallation may fail and could leave your system unusable.
```

```
Are you sure you want to reinstall [y,n,q,?] (default: n)
```

Workaround

When you are prompted to reinstall, enter **y**.

Note: This message is not displayed if you install Volume Manager with the product installer.

To display localized characters on the VEA interface, you need to change your language preferences (1666997)

If you change the locale in your operating system, the Veritas Enterprise Administrator (VEA) interface displays messages half in the new language and half in the previous language.

Workaround

To correct this issue, change the language on your VEA Preferences screen.

The vxsvc process may take more than 90% of the CPU for all the nodes (1822841)

The vxsvc process may take more than 90% of CPU resources for the nodes. During this time, there is little CPU available for I/O, Volume Manager, VCS, or the Veritas File System.

Workaround

To correct this issue, set `POLLSPEED` to `NEVER`.

Disks not discovered by the HBA driver before VxVM starts (1762451)

In SLES11, the `initrd` does not `modprobe` any device driver that requires access to a binary firmware image from the file system. The result is that some HBA drivers (`qla2xxx`, for example) may get loaded later in the boot process even when specified by the `INITRD_MODULES` option in `/etc/sysconfig/kernel`.

During testing of SFHA and SFCFS with fencing enabled, the HBA driver (`qllogic_qla2xxx`) did not complete discovery of the disks prior to VxVM startup. As a result, Volume Manager shows no disks present although they are visible in `ls SCSI`.

Workaround

Run `vxdisk scandisk` and the disks will appear in the Volume Manager configuration. If the system is part of an SFHA or SFCFS configuration, further steps may be required to rejoin the node to the cluster, or to clear faults in the VCS configuration.

On SLES11, converting LVM volumes to VxVM volumes fails when multipathed storage devices are present (1471781)

The `vxvmconvert` utility cannot convert LVM volumes to VxVM volumes when multipathed storage devices are present. This issue occurs only with SLES 11, due to changes in the LVM utilities. If multipathed devices are detected, the `vxvmconvert` utility exits with the following error:

```
vxvmconvert cannot convert multipathed devices on SLES11 systems.  
... Exiting.
```

Veritas File System known issues

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

VxFS module loading fails when freevxfms module is loaded (1736305)

The following module loading error can occur during RPM installation if the `freevxfms` module is loaded:

```
Error in loading module "vxfms". See documentation.
```

```
ERROR: No appropriate VxFS drivers found that can be loaded.  
See VxFS documentation for the list of supported platforms.
```

Workaround

There is no workaround for this issue. You can ignore the error message since the message does not affect the functionality of the `dbed_vmchecksnap` command.

A mount can become busy after being used for NFS advisory locking

If you export a VxFS file system using NFS and you perform file locking from the NFS client, the file system can become unable to be unmounted. In this case, the `umount` command fails with the `EBUSY` error.

Workaround

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

To resolve this known issue

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

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

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

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

- 4 Start the database.

Possible write performance degradation with VxFS local mounts

Some applications that allocate large files without explicit preallocation may exhibit reduced performance with the VxFS 5.1 release compared to the VxFS 5.0

MP3 release due to a change in the default setting for the tunable `max_seqio_extent_size`. One such application is DB2. Hosting DB2 data on a single file system extent maximizes the potential for sequential pre-fetch processing. When DB2 detects an application performing sequential reads against database data, DB2 begins to read ahead and pre-stage data in cache using efficient sequential physical I/Os. If a file contains many extents, then pre-fetch processing is continually interrupted, nullifying the benefits. A larger `max_seqio_extent_size` value reduces the number of extents for DB2 data when adding a data file into a tablespace without explicit preallocation.

The `max_seqio_extent_size` tunable controls the amount of space that VxFS automatically preallocates to files that are allocated by sequential writes. Prior to the 5.0 MP3 release, the default setting for this tunable was 2048 file system blocks. In the 5.0 MP3 release, the default was changed to the number of file system blocks equaling 1 GB. In the 5.1 release, the default value was restored to the original 2048 blocks.

The default value of `max_seqio_extent_size` was increased in 5.0 MP3 to increase the chance that VxFS will allocate the space for large files contiguously, which tends to reduce fragmentation and increase application performance. There are two separate benefits to having a larger `max_seqio_extent_size` value:

- Initial allocation of the file is faster, since VxFS can allocate the file in larger chunks, which is more efficient.
- Later application access to the file is also faster, since accessing less fragmented files is also more efficient.

In the 5.1 release, the default value was changed back to its earlier setting because the larger 5.0 MP3 value can lead to applications experiencing "no space left on device" (ENOSPC) errors if the file system is close to being full and all remaining space is preallocated to files. VxFS attempts to reclaim any unused preallocated space if the space is needed to satisfy other allocation requests, but the current implementation can fail to reclaim such space in some situations.

If your workload has lower performance with the VxFS 5.1 release and you believe that the above change could be the reason, you can use the `vxtunefs` command to increase this tunable to see if performance improves.

To restore the benefits of the higher tunable value

- 1 Increase the tunable back to the 5.0 MP3 value, which is 1 GB divided by the file system block size.

Increasing this tunable also increases the chance that an application may get a spurious ENOSPC error as described above, so change this tunable only for file systems that have plenty of free space.
- 2 Shut down any application that are accessing any large files that were created using the smaller tunable setting.
- 3 Copy those large files to new files, which will be allocated using the higher tunable setting.
- 4 Rename the new files back to the original names.
- 5 Restart any applications were shut down earlier.

Veritas Volume Replicator known issues

The following are known issues in this release of Veritas Volume Replicator.

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 GCO site failover may fail (1558257)

Issue 1:

When `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 GCO 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 GCO site failover and therefore, the file systems on the new primary site may not be fully consistent.

Workaround

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 -t 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 -t vxfs /dev/vx/dsk/<dg>/<data_volume>
```

Running SUSE Linux and using Novell's YaST tool to configure an IPv6 address may result in an error

When Novell's YaST tool is invoked to configure an IPv6 address on a different network interface and if:

- the host name, the DNS server name and domain name are specified to the YaST tool.
- IPv6 address is assigned by the Dynamic Host Configuration Protocol (DHCP).
- the "Write Hostname to /etc/hosts" option is selected (this is selected by default).

This results in the `vradmin` command returning the following error:

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

This happens because the YaST tool can replace the `/etc/hosts` entry containing `127.0.0.2` from the IPv4 host name to the specified new IPv6 host name. For example:

```
127.0.0.2 v6hostname.space.ipv6.com v6hostname
```

Workaround

To resolve this issue

- 1 Edit the `/etc/hosts` file to specify the correct IPv6 address.
- 2 Restart the `vradmin` daemon on all VVR hosts:

```
# /etc/init.d/vras-vradmind.sh restart
```

Storage Foundation 5.0MP3 Rolling Patch 2 required for replication between 5.0 MP3 and 5.1 (1800600)

In order to replicate between primary sites running Storage Foundation 5.0 MP3 and secondary sites running Storage Foundation 5.1, or vice versa, you must install the Storage Foundation 5.0MP3 Rolling Patch 2 on the nodes using 5.0MP3. This patch resolves several outstanding issues for replicating between versions.

The patch is recommended for Storage Foundation, Storage Foundation Cluster File System and Storage Foundation Cluster File System RAC products.

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

Issue: After upgrading VVR to an IPv6-only environment in 5.1 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 will appear 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`) will fail with the error:

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

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

To workaround this issue, restart vradmind on all the cluster nodes using the following commands:

```
# /etc/init.d/vras-vradmind.sh restart
```

Replication hang when VVR logowner is on CVM slave node (1807525)

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:

As a workaround, configure RVGLogOwner agent such that the VVR logowner will always reside on the CVM master node.

On the node where we want to switch the logowner from:

```
# vxrvrg -g dgname set logowner=off rvgname
```

On the master node.

```
# vxrvrg -g dgname set logowner=on rvgname
```

If using VEA to create an Replicated Data Set (RDS) fails, messages display corrupt strings in the Japanese locale (1726499, 1377599)

When using VEA to create an RDS fails, because the volumes do not have a DCM log on all nodes, the message window displays corrupt strings and unlocalized error messages.

No longer supported

This section describes Storage Foundation features that are not supported in this release.

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

- The `seconly` feature in SFCFS
- Storage Expert
- The use of the `vxvoladm` command line utility
- Intelligent Storage Provisioning (ISP)
- When upgrading from 4.x or 5.0 SF for ORACLE (HA)/SFRAC to 5.1 SF (HA)/SFRAC, the following SFDB features will no longer be supported:
 - `Oramap (libvxoramap)` and storage mapping commands (`dbed_analyzer`, `vxstorage_stats`)
 - DBED providers (`DBEDAgent`), Java GUI, and `dbed_dbprocli`. The DBED Oracle GUI features can only be done through command lines. Database cloning, database flashsnap and DBDST scheduling features will not be supported through the GUI. You can use Veritas Storage Foundation Manager 2.1 that is a separate product, to display Oracle database information such as tablespaces, database to lun mapping, and tablespace to lun mapping.
 - `dbdst_makelbfs`, `vxdbts_fstatsummary`, `dbdst_fiostat_collector`, `vxdbts_get_datafile_stats`
 - `dbed_saveconfig`, `dbed_checkconfig`
 - `dbed_ckptplan`, `dbed_ckptpolicy`
 - `dbed_scheduler`
 - The `sfua_db_config` functionality has changed in 5.1. This command is no longer needed to create a SFDB repository. You can issue `dbed_update` to create a new SQLite SFDB repository. The `sfua_db_config` will be used to set user and group access to various SFDB directories instead.
 - The `sfua_rept_adm` command was used in 5.0 to perform repository backup and restore. This command will be obsolete in 5.1. The `sfua_rept_util` command can be used instead to perform SQLite repository backup and restore.

When you upgrade from 4.x or 5.0 SF for DB2 (HA)/SF for Sybase (HA) to 5.1 SF (HA), you will lose all the respective SF for DB2 and SF for Sybase functionalities, but they can use the default Snapshots and Checkpoint features available with Storage Foundation. You can continue to use CIO fast I/O access with DB2 or Sybase databases.

VVR will use TCP as default protocol from next release

In Storage Foundation 5.1 and previous releases, VVR uses UDP as the default transport protocol for communicating between the primary and secondary. In future releases, the default protocol will be switched to TCP. Users will have the option to manually switch back to UDP.

Documentation

Product guides are available on the software disc in PDF format. Symantec's Veritas Storage Foundation documentation is also available on the Symantec website.

Symantec's Veritas Storage Foundation 5.1 documentation set is available at the following URL:

<http://www.symantec.com/business/support/overview.jsp?pid=15107>

Relevant component product release notes

Read the relevant component product release notes before installing any version of Veritas Storage Foundation.

The following documents are relevant component product release notes:

- *Veritas Cluster Server Release Notes* (`vcs_notes.pdf`)
- *Veritas Storage Foundation Cluster File System for Oracle RAC Release Notes* (`sfrac_notes.pdf`)

The following document is the relevant component product release notes:

- *Veritas Cluster Server Release Notes* (`vcs_notes.pdf`)

Storage Foundation guides

The following manuals, along with the online help, comprise the Veritas Storage Foundation documentation set:

Table 1-10 describes the guides in the Veritas Storage Foundation documentation set.

Table 1-10 Guides in Veritas Storage Foundation documentation set

Guide Title	Filename
<i>Veritas Storage Foundation and High Availability Getting Started Guide</i>	<code>getting_started.pdf</code>

Table 1-10 Guides in Veritas Storage Foundation documentation set (*continued*)

Guide Title	Filename
<i>Veritas Storage Foundation read me first</i>	readme_first.txt
<i>Veritas Storage Foundation Release Notes</i>	sf_notes.pdf
<i>Veritas Storage Foundation Installation Guide</i>	sf_install.pdf
<i>Veritas Storage Foundation: Storage and Availability Management for Oracle Databases</i>	sf_adv_ora.pdf
<i>Veritas Storage Foundation Advanced Features Administrator's Guide</i>	sf_advanced_admin.pdf
<i>Veritas File System Administrator's Guide</i>	vxfs_admin.pdf
<i>Veritas File System Programmer's Reference Guide</i>	vxfs_ref.pdf
<i>Veritas Volume Manager Administrator's Guide</i>	vxvm_admin.pdf
<i>Veritas Volume Manager Troubleshooting Guide</i>	vxvm_tshoot.pdf

Veritas Storage Foundation Cluster File System documentation

The following Veritas Storage Foundation Cluster File System documentation is available with all Veritas Storage Foundation HA product suites:

[Table 1-11](#) describes the Veritas Storage Foundation Cluster File System (CFS) documentation set.

Table 1-11 Guides in Veritas Storage Foundation Cluster File System documentation set

Guide Title	Filename
<i>Veritas Storage Foundation Cluster File System Release Notes</i>	sfcfs_notes.pdf
<i>Veritas Storage Foundation Cluster File System Installation Guide</i>	sfcfs_install.pdf
<i>Veritas Storage Foundation Cluster File System Administrator's Guide</i>	sfcfs_admin.pdf
<i>Veritas Storage Foundation: Storage and Availability Management for Oracle Databases</i>	sf_adv_ora.pdf

Table 1-11 Guides in Veritas Storage Foundation Cluster File System documentation set (*continued*)

Guide Title	Filename
<i>Veritas Storage Foundation Advanced Features Administrator's Guide</i>	sf_advanced_admin.pdf

Veritas Storage Foundation Cluster File System for Oracle RAC documentation

The following Veritas Storage Foundation Cluster File System for Oracle RAC documentation is available with all Veritas Storage Foundation HA product suites:

[Table 1-12](#) describes the Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFS RAC) documentation set.

Table 1-12 Guides in Veritas Storage Foundation Cluster File System for Oracle RAC documentation set

Guide Title	Filename
<i>Veritas Storage Foundation Cluster File System for Oracle RAC Release Notes</i>	sfcfsrac_notes.pdf
<i>Veritas Storage Foundation Cluster File System Installation Guide</i>	sfcfsrac_install.pdf
<i>Veritas Storage Foundation Cluster File System Administrator's Guide</i>	sfcfs_admin.pdf

Veritas Cluster Server documentation

The following Veritas Cluster Server documentation is available with all Veritas Storage Foundation HA product suites:

[Table 1-13](#) describes the Veritas Cluster Server documentation set.

Table 1-13 Guides in Veritas Cluster Server documentation set

Guide Title	Filename
<i>Veritas Cluster Server Release Notes</i>	vcs_notes.pdf
<i>Veritas Cluster Server Installation Guide</i>	vcs_install.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents.pdf

Table 1-13 Guides in Veritas Cluster Server documentation set (*continued*)

Guide Title	Filename
<i>Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide</i>	vcs_vvr_agent.pdf
<i>VCS Enterprise Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_install.pdf
<i>VCS Enterprise Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_install.pdf
<i>VCS Enterprise Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_install.pdf

Veritas Volume Replicator documentation

The following Veritas Volume Replicator documentation is available with the Veritas Volume Replicator option:

[Table 1-14](#) describes the Veritas Volume Replicator documentation set.

Table 1-14 Guides in Veritas Volume Replicator documentation set

Guide Title	Filename
<i>Veritas Volume Replicator Administrator's Guide</i>	vvr_admin.pdf
<i>Veritas Volume Replicator Planning and Tuning Guide</i>	vvr_planning.pdf
<i>Veritas Volume Replicator Advisor User's Guide</i>	vvr_advisor_users.pdf

Manual Pages

The Veritas online manual pages are installed in the `/opt/VRTS/man` directory. This directory can be added to the `MANPATH` environment variable.

If the `MANPATH` environment variable does not include `/opt/VRTS/man`, you can view the desired manual page by entering the following command:

```
# man -M /opt/VRTS/man manual_page_name
```

Veritas Storage Foundation 5.1 for Xen

This appendix includes the following topics:

- [Supported operating systems for Xen](#)
- [Supported features](#)
- [Unsupported features](#)
- [Installing the Veritas Software in Dom0](#)
- [Verifying Software Versions](#)
- [Uninstalling the Veritas Software from Dom0](#)

Supported operating systems for Xen

This release provides support for Veritas Storage Foundation 5.1 on the Xen platform for Linux.

Supported operating systems are:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2.
- SUSE Linux Enterprise Server 11 (SLES 11).
- Oracle Enterprise Linux 5 (OEL 5) with Update 3

Supported features

In this release, only Veritas Volume Manager is supported on Xen Dom0.

The following features are provided in this release in the Xen environment:

- Configuration of Veritas Volume Manager volumes in Dom0.
- Export of VxVM volumes to DomU as raw partition devices.
- Support for both SCSI devices and IDE disk devices (`/dev/sdx` and `/dev/hdx`).
- Relayout of a volume in Dom0 that has been exported to DomU is possible. Some reduction in I/O throughput may be experienced during the operation.
- Resizing a volume in Dom0 is not seen in DomU until the domain is rebooted. The file system must be resized independently of the volume. This limitation is imposed by Xen.
- Array Support Libraries (ASLs) and Array Policy Modules (APMs) that are supported in 5.1.
- Root encapsulation.
- SELinux in permissive or enforcing mode.

Unsupported features

The following features of Veritas Storage Foundation Enterprise product are not supported in this release:

- Veritas File System (VxFS).
- Cluster Volume Manager (CVM).
- Cluster File System (CFS).
- Veritas Cluster Server (VCS).
- The VEA server in DomU.
- Multi-volume file systems, volume sets, and Quality of Storage Service (QoS).
- Enhanced I/O support in DomU that would usually be obtained by configuring a VxFS file system on a VxVM volume.
- Resizing a VxVM volume in DomU.

Installing the Veritas Software in Dom0

If you plan to install the Veritas Storage Foundation 5.1 software for the first time and not upgrade an existing system, review the preinstallation instructions.

See the *Veritas Storage Foundation 5.1 Installation Guide* .

For important release information, review the *Veritas Storage Foundation Release Notes*, and all documents in the `release_notes` directory.

To install the Veritas software in Dom0

- 1 Obtain the Veritas Storage Foundation 5.1 software from the software disc or downloaded image.
- 2 Use the Veritas Volume Manager installation script (`installvm`) instead of the generic Storage Foundation `installer` script to install the packages.

Enter the following command from the `platform/volume_manager` directory of the mounted DVD-ROM or downloaded image:

```
# ./installvm -installonly
```

The `-installonly` option is required to perform the installation without configuring the software. Do not create any VxVM volumes before you install the packages that are required to support the Xen environment.

- 3 To configure the software, run the installation script again, this time specifying the `-configure` option.

```
# ./installvm -configure
```

You are prompted to supply a valid Veritas Storage Foundation product license. Presently, there is no separate product license required for enabling Xen support.

See the *Veritas Storage Foundation Installation Guide* for help on how to respond to the installation prompts.

- 4 If the `installvm` script did not reboot the Xen kernel, then manually reboot the Xen kernel now.

Verifying Software Versions

To list the Veritas packages that are installed on your system, enter the following command from the appropriate domain:

```
# rpm -qa | grep VRTS
```

Uninstalling the Veritas Software from Dom0

To uninstall the Veritas software from Dom0

- 1 Log on as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
- 3 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
- 4 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

- 5 To shut down and remove the installed Veritas packages, use the appropriate command in the `/opt/VRTS/install` directory. For example, to uninstall the Veritas Storage Foundation packages, use the following commands:

```
# cd /opt/VRTS/install  
# ./uninstallvm
```