

Veritas Storage Foundation™ Release Notes

Solaris

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.

Storage Foundation Release Notes

This document 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/solaris/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)

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/334829>

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 |
| 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 |
| VRTScmccc | N/A | Obsolete in 5.1. Delivered with 5.x CMC release. | Veritas Cluster Management Console Cluster Connector |
| VRTScmedc | N/A | Obsolete in 5.0MP3. Delivered with 5.x CMC release. | User Documentation for Veritas Cluster Management Console |

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

| 5.0 Package Name | 5.1 Package name | Explanation of changes | Package description |
|------------------|------------------|--|---|
| 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 |
| 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 |

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

| 5.0 Package Name | 5.1 Package name | Explanation of changes | Package description |
|------------------|------------------|---|---|
| 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 |
| 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 |

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

| 5.0 Package Name | 5.1 Package name | Explanation of changes | Package description |
|------------------|------------------|---|---|
| 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 |
| VRTSglm | VRTSglm | No change. | Veritas Global Lock Manager |
| VRTSgms | VRTSgms | No change. | Veritas Group Messaging Services |
| VRTSsisco | N/A | Obsolete in 5.1. | Symantec Infrastructure Core Services Common |

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

| 5.0 Package Name | 5.1 Package name | Explanation of changes | Package description |
|------------------|------------------|--|--|
| 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 |
| VRTSmapro | N/A | Consolidated into VRTSob. | Veritas Storage Mapping Provider |
| VRTSmh | VRTSsfmh | Consolidated into VRTSsfmh. | Veritas Storage Foundation Management host |
| VRTSob | VRTSob | No change. | Veritas Enterprise Administrator Service |
| VRTSobc33 | N/A | Obsolete in 5.1. Functionality Delivered with SFM release | Veritas Enterprise Administrator Core |
| VRTSobgui | N/A | Obsolete in 5.1. Functionality Delivered with SFM release. | Veritas Enterprise Administrator |
| VRTSobweb | 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 |
| VRTSpbx | 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 |

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

| 5.0 Package Name | 5.1 Package name | Explanation of changes | Package description |
|------------------|------------------|---|---|
| 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 VRTSvcsvsvr. | Veritas Cluster Server Bundled Agents |
| VRTSvcsvcsdb | VRTSvcsvcsea | Consolidated into VRTSvcsvcsea. | 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 |
| VRTSvcsvcsmsg | VRTSvcsvcs | Consolidated into VRTSvcsvcs. | Veritas Cluster Server English Message Catalogs |
| VRTSvcsvcsmn | VRTSvcsvcs | Consolidated into VRTSvcsvcs. | Manual Pages for Veritas Cluster Server |
| VRTSvcsvcsor | VRTSvcsvcsea | Consolidated into VRTSvcsvcsea. | Veritas High Availability Agent for Oracle |
| VRTSvcsvcssy | VRTSvcsvcsea | Consolidated into VRTSvcsvcsea. | Veritas High Availability Agent for Sybase |
| VRTSvcsvsvr | VRTSvcsvcsag | Consolidated into VRTSvcsvcsag. | Veritas Cluster Server Agents for Veritas Volume Replicator |
| VRTSvcsvdid | N/A | Obsolete in 5.1. | Veritas Device Identification API |
| VRTSvcsvlic | VRTSvcsvlic | No change. | Symantec License Utilities |

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

| 5.0 Package Name | 5.1 Package name | Explanation of changes | Package description |
|------------------|------------------|--|--|
| 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 |
| 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 |

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

| 5.0 Package Name | 5.1 Package name | Explanation of changes | Package description |
|------------------|------------------|--|---------------------------------|
| 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 packages: installs only the basic functionality for the selected product.
- Recommended packages: installs the full feature set without optional packages.
- All packages: installs all available packages.

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

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

Veritas extension for Oracle Disk Manager package 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 package for ODM by default when Storage Foundation is installed.

For Storage Foundation Cluster File System, the `GMSVRTSgms.pkg` package is also installed.

Change in Storage Foundation packaging standard on Solaris

With this release, the Storage Foundation packages on Solaris use the Solaris single file package standard and not the compressed packages. This change eliminates the need to copy the packages to a temporary directory, unzip and extract these packages, and then install. You can now directly install the packages from the product disc.

Rootpath option to uninstall scripts

The `-rootpath` option is used to specify the path from where the packages must be uninstalled. You must use this option if you did not install the packages on the default location.

On Solaris, `-rootpath` passes `-R <root_path>` to `pkgrm`.

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.

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.

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.

Installer support for alternate boot disk

The installer program supports install, uninstallation, and upgrades on alternate boot disks for Solaris.

Refer to the *Veritas Storage Foundation and High Availability Installation Guide's* section on Live Upgrade.

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 aggregate links during installation

The installer program asks if you want to use an aggregate NIC, if so it configures the `llttab` file for you. Note that the installer program does not detect aggregate 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 5.0, 5.0 MP1, 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 SF 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`
- Quick I/O commands: `qio_getdbfiles`, `qio_recreate`, `qio_convertdbfiles`
- 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: commands** `dbdst_makelbfs`, `vxdbs_fstatsummary`, `dbdst_fiostat_collector`, `vxdbs_get_datafile_stats`
- `dbed_saveconfig`, `dbed_checkconfig`
- `dbed_ckptplan`, `dbed_ckptpolicy`
- `qio_convertdbfiles -f` option which is used to check for file fragmentation
- `dbed_scheduler`

The following features are longer supported in release 5.1:

- Storage Foundation for DB2 tools
- Storage Foundation for Sybase tools

Storage Foundation complies with Solaris 10 Service Management Facility (SMF)

Storage Foundation 5.1 complies with Solaris Service Management Facility (SMF). For more information about SMF, refer to the Sun website.

Oracle Disk Manager driver

Prior to this release, the `VRTSodm` package has the `/etc/init.d/odm` script that is used to manage the ODM driver and mount point. The script is applicable to Solaris 9 and 10 on SPARC and Solaris 10 x86_64.

In the `/etc/init.d/odm` script, it has the `start`, `stop`, `restart`, and `status` options. For example, the `start` option starts the ODM driver by loading the driver into the running Solaris kernel and creates a mount point (`/dev/odm`).

On a system where the `VRTSodm` package is installed, you can execute the `/etc/init.d/odm` script as user root.

```
# /etc/init.d/odm start
```

On Solaris 10 SPARC and x86-64, Sun Microsystems implemented a new facility called the Service Management Facility (SMF). One of the services in the SMF is a new interface to replace the scripts or tasks that are located in the `/etc/init.d` directory. When a driver or application is brought under the SMF management, the SMF user interface must be used. The corresponding scripts in the `/etc/init.d` directory must be removed from the system.

In this release, the ODM driver on Solaris 10 is converting the management of the ODM driver to use the SMF. In SMF, the ODM driver is identified as `vxodm`.

On Solaris 9, where `VRTSodm` package is supported, you should continue to use the `/etc/init.d/odm` script.

On Solaris 10 systems, after you installed the SMF aware `VRTSodm` package and rebooted the system, the `svcs(1)` command is used to list the status of the ODM driver.

To perform operations using the SMF interface on Solaris 10 in non-global zones:

See the *Veritas Storage Foundation High Availability Virtualization Guide* for more information on Veritas extension for Oracle Disk Manager.

To display the status of the ODM driver, type:

```
# /usr/bin/svcs vxodm
STATE      STIME      FMRI
online     15:29:07  svc:/system/vxodm:default
```

The `svcs vxodm` is equivalent to `/etc/init.d/odm status`.

To stop the ODM driver, type:

```
# /usr/sbin/svcadm disable vxodm
# /usr/bin/svcs vxodm
STATE      STIME      FMRI
disabled   14:02:45  svc:/system/vxodm:default
```

The `svcadm disable vxodm` is equivalent to `/etc/init.d/odm stop`.

To start the ODM driver, type:

```
# /usr/sbin/svcadm enable vxodm
# /usr/bin/svcs vxodm
STATE      STIME      FMRI
online     15:29:24  svc:/system/vxodm:default
```

The `svcadm enable vxodm` is equivalent to `/etc/init.d/odm start`.

The `restart` option is used to disable and then enable the ODM driver.

To restart the ODM driver, type:

```
# /usr/sbin/svcadm restart vxodm
# /usr/bin/svcs vxodm
STATE      STIME      FMRI
online     15:30:56  svc:/system/vxodm:default
```

The `svcadm restart vxodm` is equivalent to `/etc/init.d/odm restart`.

In SMF, there is also a `refresh` option that is used to re-read the ODM service configuration file that is found in the `/var/svc/manifest/system/vxodm/odm.xml` file.

To refresh the ODM service configuration file, type:

```
# /usr/sbin/svcadm refresh vxodm
```

On standalone and cluster environment, the same SMF commands are used. There is no change to the ODM's GAB port `d` implementation in the cluster environment.

The ODM driver is dependent on the VxFS driver. The VxFS's SMF service is identified by the name `vxfsldlic`. The `vxfsldlic` SMF service must be online before the `vxodm` SMF service can come online. In a cluster environment, the `vxodm` SMF service is dependent on the GAB's SMF service. The GAB SMF service is identified by the name `gab`. You can use the `svcs(1)` command to list the status for the `vxfsldlic` and the `gab` SMF services.

To display the status for the `vxfsldlic` SMF service, type:

```
# svcs vxfsldlic
STATE          STIME          FMRI
online         Apr_01         svc:/system/vxfs/vxfsldlic:default
```

To display the status for the `gab` SMF service, type:

```
# svcs gab
STATE          STIME          FMRI
online         Apr_01         svc:/system/gab:default
```

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.

Note: The Storage Foundation Thin Reclamation feature is not supported on the Solaris x64 operating environment.

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.

Veritas Volume Manager coexistence with SVM and ZFS

Solaris Volume Manager (SVM) is a logical volume manager software provided by SUN. ZFS is a type of file system presenting a pooled storage model developed by SUN. File systems can directly draw from a common storage pool (zpool). Veritas Volume Manager (VxVM) can be used on the same system as SVM and ZFS disks.

This release of VxVM improves the relationship between VxVM with ZFS and SVM in the following ways:

- VxVM protects devices in use by SVM or ZFS from VxVM operations that may overwrite the disk. These operations include initializing the disk for use by VxVM or encapsulating the disk.
- VxVM classifies and displays the SVM and ZFS disks with the correct labeling. The output of the `vxdisk list` command displays the label SVM or ZFS, respectively.
- VxVM displays an error message when you attempt to perform an operation on a device that is in use by SVM or ZFS.

Note: This release of VxVM can discover ZFS disks for Solaris releases up to Solaris 10 Update 7.

VxVM does not discover ZFS disks for Solaris 10 Update 8 and later versions due to a change in Solaris 10 Update 8. See the Symantec Support web site to check for patch availability at <http://entsupport.symantec.com/docs/334829>.

Before you can manage an SVM or ZFS with VxVM, you must remove it from SVM or ZFS control. Similarly, to begin managing a VxVM disk with SVM or ZFS, you must remove the disk from VxVM control.

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

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 `vxvg reminor` 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 `vxvg reminor` 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:

- `vxctl [-c] mode`
- `vxvg list [<dgname>]`
- `vxdisk list [<daname>]`
- `vxprint [-Aqht]`

Use the following command to disable this functionality:

```
# vxdctl request_threads 0
```

DMP coexistence with native multipathing

This release introduces support for using Veritas Dynamic Multipathing (DMP) with raw devices. A new tunable, `dmp_native_multipathing`, controls the behavior. If the `dmp_native_multipathing` tunable is set to on, DMP intercepts I/O requests, operations such as open and close, and ioctls that are sent on the raw device path. If the `dmp_native_multipathing` tunable is set to off, these requests are sent directly to the raw device.

DMP should not be enabled if EMC PowerPath is installed, or if MPxIO is enabled. In those cases, DMP is not required since the devices use native multipathing.

Upgrading the array support

The Storage Foundation 5.1 release includes all array support in a single package, VRTSaslapm. The array support package includes the array support previously included in the VRTSvxvm package. The array support package 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 package 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 vxdmpinq utility renamed to the vxscsiinq utility

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

Enclosure-based naming is now the default naming scheme

In this release, the enclosure-based 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 39.

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 24.

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 Solaris operating systems

This release of the Veritas products is supported on the following Solaris operating systems:

- Solaris 9 (SPARC Platform 32-bit and 64-bit)

- Solaris 10 (SPARC or x64 Platform 64-bit)

If necessary, upgrade Solaris before you install the Veritas products.

Install all the latest required Solaris patches listed in the product *Release Notes*.

See “[Required Solaris patches](#)” on page 35.

For information about the use of this product in a VMware Environment on Solaris x64, refer to <http://entsupport.symantec.com/docs/289033>

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

<http://entsupport.symantec.com/docs/334829>

Required Solaris patches

Before installing Veritas Storage Foundation, ensure that the correct Solaris patches are installed.

See <http://sunsolve.sun.com> for the latest Solaris patch updates.

The following patches are required for Solaris SPARC:

Table 1-2 Solaris SPARC patches

| Operating system | Sun patch number |
|------------------|--|
| Solaris 9 | 114477-04 122300-29 - required for Live Upgrade |
| Solaris 10 | 118833-36 118918-24 119254-70 119578-30 120011-14 120272-25 123839-07 125503-02 125547-02 125731-05 125891-01 126419-02 126540-02 126897-02 127127-11 127755-01 |

The following patches are required for Solaris x64:

Table 1-3 Solaris x64 patches

| Operating system | Sun patch number |
|------------------|------------------|
| Solaris 10 | 118344-14 |
| | 118855-36 |
| | 119043-11 |
| | 119131-33 |
| | 120012-14 |
| | 125732-05 |
| | 127128-11 |

Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x6000. When you install the Veritas File System package, `VRTSvxfs`, the `VRTSvxfs` packaging scripts check the values of these variables in the kernel. If the values are less than the required values, `VRTSvxfs` increases the values and modifies the `/etc/system` file with the required values. If the `VRTSvxfs` scripts increase the values, the installation proceeds as usual except that you must reboot and restart the installation program. A message displays if a reboot is required.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System. Use the following commands to check the values of the variables:

```
# echo "lwp_default_stksize/X" | mdb -k
lwp_default_stksize:
lwp_default_stksize:          6000

# echo "svc_default_stksize/X" | mdb -k
svc_default_stksize:
svc_default_stksize:          6000
```

If the values shown are less than 6000, you can expect a reboot after installation.

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

To avoid a reboot after installation, you can modify the `/etc/system` file with the appropriate values. Reboot the system prior to installing the packages. Appropriate values to the `/etc/system` file are shown in the following examples:

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
```

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 Solaris SPARC is available at the following URL:

<https://vos.symantec.com/patch/detail/2960>

The patch for Solaris x64 is available at the following URL:

<https://vos.symantec.com/patch/detail/2961>

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 for Oracle RAC Release Notes* (`sfrac_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

There are no Veritas Storage Foundation software limitations in the 5.1 release.

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.

Volume Manager object names cannot include non-ASCII characters (Japanese and Chinese locales)

Volume Manager object names cannot include any non-ASCII characters. The limitation applies to both Japanese and Chinese locales.

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:

```
# vxddmpadm gettune dmp_restore_internal
```

Veritas File System software limitations

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

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.

Veritas Storage Foundation and High Availability features not supported on Solaris x64

The following Storage Foundation and High Availability features that are supported on Solaris SPARC and not supported on Solaris x64:

- Application Templates
- Disk layout versions 4 and 5
- Data Management Application Programming Interface (DMAPI)
- Veritas Cluster Server (VCS) gabdisk support
- Storage Foundation Manager (SFM) client only
- Storage Foundation database editions for DB2
- Localization (L10N)

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 1-4](#) describes fixed issues in the Storage Foundation 5.1 release.

Table 1-4 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 complaint. |
| 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 <code>dbdst_report</code> manual page. |
| 1121091 | You do not need to add the <code>_ora</code> suffix to use the <code>dbdst</code> manual pages. |
| 862687, 862092 | Database FlashSnap clones supports <code>DB_RECOVERY_FILE_DESTINATION</code> 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 1-5](#) describes fixed issues in the Veritas Volume Manager 5.1 release.

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues

| Incident | Description |
|----------|--|
| 1822681 | memory leak in vxio/voldr1_cleansio_start. [SUN bug ID: 6846549] |
| 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 [SUN bug ID: 6874695] |
| 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 |
| 1782036 | Sun: SC cannot open libvxvmso after upgrade to 5.1 [SUN bug ID: 6864609] |
| 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 | vxdctl settz and vxconfigd core dump if TZ environment variable is not set. |
| 1755869 | tunable addition: gabmaxsend and receiver flowcontrol watermark |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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_iocnt() 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 |
| 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 |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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 [SUN bug ID: 6850933] |
| 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 [SUN bug ID: 6835347] |
| 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. |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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 | vxndmpadm 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 gendmpidone. |
| 1530126 | DMP : dmplinux_unplug() panic on linux, for no associated node in dmpnode |
| 1517760 | VRAS: vradmind core dump if stats collection is enabled |
| 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. |
| 1505434 | VxVM:LDOM: Volume becomes inaccessible from the guest in the event of primary domain reboot [SUN bug ID: 6795836] |
| 1503309 | VxVM: DMP doesn't log error returned from below driver/SCSI. |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|------------------|--|
| 1483298 | Man page for vxddctl does not include all options that are listed in -H. [SUN bug ID: 6790194] |
| 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/ [SUN bug ID: 6716991] |
| 1475703, 1369599 | Export API libvxvm_get_disks() and libvxvm_get_subpaths() should return the same path for the same device [SUN bug ID: 6399937] |
| 1475697 | Provide an API for checking multi- ownership and status of a diskgroup [SUN bug ID: 6308791] |
| 1475692 | The size of large VxVM volumes must be reported correctly to Solaris utilities [SUN bug ID:6571880] |
| 1475691 | VxVM should recognize disks in use by ZFS or SVM [SUN bug ID: 6672721, 670507, 6715158] |
| 1471821 | 'initdmp' section in vxddctl manpage still incorrect |
| 1471784 | [5.OMP3RP1 x64] vm can not create stripe-mirror/mirror-stripe/mirror volume with maxsize. [SUN bug ID: 6844425] |
| 1471771 | vxdisksetup should call 'mount' with complete pathname |
| 1471606 | Solaris: patch/package scripts needs to sync boot archive for sparc as well x86 for 2.10u6. [SUN bug ID: 1471606] |
| 1471581 | vxconfigd may hang if no SCSI timeout value is set while issuing various SCSI commands |
| 1471487 | Critical Minimum Queue and Round-robin Improvements |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 1471263 | machine has panicked when added the disk from dg as a foreign device using "vxddmpadm addforeign". [SUN bug ID: 6884813] |
| 1470732 | 5.0MP3 : vxconfigd dumps core if all paths are excluded. |
| 1470251 | volslabd utilized 100% cpu time |
| 1470102 | vxddmpadm getattr for failoverpolicy and tpdmode is giving usage error. |
| 1469351 | User confused about snap operations when split to new DG |
| 1468647 | vxddmpdebug fails to find ugettxt |
| 1463197 | no path disable event occurs during I/O error analysis in dmp when pulling a FC cable out with 5.0MP3 |
| 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:voldr_l_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 4.1MP3 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. |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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 |
| 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 [SUN bug ID: 6737054] |
| 1386592 | tmp files not cleaned up after running vxdisksetup |
| 1369610 | Error message from vxdg import misleading/incorrect when disk is write-protected [SUN bug ID: 6264161] |
| 1369597 | vxlufinish uses wrong solaris command for unmounting alternate root environment - should use luumount [SUN bug ID: 6446847] |
| 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 vxconfigf 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. |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 1288413 | Revisit our copy to /etc/vx/slib for Sol 10/11 [SUN bug ID: 6716991] |
| 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 & vxdisksetup 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 |
| 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 |
| 1057239 | Support of raw device in Solaris local zones [SUN bug ID: 6326200] |

Table 1-5 Veritas Volume Manager and Veritas Volume Replicator 5.1 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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 [SUN bug ID: 6308791] |
| 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 [SUN bug ID: 6569995] |
| 795129 | cmhaltcl on one of the cluster nodes causes one of the disk groups to lose shared flag |
| 786357 | Request to make voldrl_volumemax_drtregs a tunable |
| 597517 | unable to initialize EFI labeled >1tb PP devices |
| 543638 | vxdkmpadm/vxdiskunsetup doesn't work well if tpdmode=native |
| 339282 | RFE : Failed to create more than 256 config copies in one DG [SUN bug ID: 6284604] |
| 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 1-6](#) describes fixed issues in the Veritas Volume Manager 5.0MP3 RP1 release.

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0MP3 RP1 fixed issues

| Incident | Description |
|--------------------|--|
| 1444425 | The vxsnap prepare manual page includes support for the mirror= attribute. |
| 1443748 | Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load. |
| 1443706 1443679 | Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion. |
| 1442369 1224659 | Fixed a bug in vxconfigbackupd script leading to 0 byte binconfig file being created. |
| 1441003 | Fixed a secondary panic due to double free of message with TCP protocol and 16 connection. |
| 1435681 | Fixed an issue with vxesd looping using 100% of one CPU. |
| 1435470 | Fixed an issue with cluster nodes panicking after installing 5.0 MP3. |
| 1433120 | Fixed an issue with after a reboot site read policy is not honored. [SUN bug ID: 6765037] |
| 1269468 | Fixed an issue with vxconfigd core dumps. Fixed a system panic in vxio:voldr1_trans_copy. |
| 1424479 | Fixed an issue with vxdmpadm dumped core when executing vxdmpadm list dmpnode command. [SUN bug ID: 6765037] |
| 1425338 | Fixed an issue with CVR fails to connect rlinks followed by vxconfigd hangs on secondary. |
| 1421088 | Fixed a secondary panic due to a corrupted volsioq_start. |
| 1416930 | Fixed an issue with the vxvm daemon that comes online when the system is rebooted. [SUN bug ID: 6747492, 6787315] |

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0MP3 RP1 fixed issues (*continued*)

| Incident | Description |
|--------------------|--|
| 1414451 | The vxsnap manual page includes mirror=enclosure parameter to avoid being mirrored on the same enclosure. |
| 1385126 | Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery. |
| 1413700 | Fixed an issue with the wrong label on a device lead VxVM to calculate the wrong public region size. |
| 1412784 | Fixed an issue with the system hanging while creating volumes in a guest Ldom. [SUN bug ID: 6744348] |
| 1410216 | Fixed a secondary log error causing rlink disconnect after IBC unfreeze. |
| 1409991 | Fixed an issue with vxclust configuration caused the cluster to panic. [SUN bug ID: 6745612] |
| 1409986 | Fixed a segmentation fault on x64 system when running the vxdumpadm list dmpnode all command. [SUN bug ID: 6744469] |
| 1403370 1401188 | Fixed a system panic after running the vxctl enable or vxconfig -k commands. |
| 1288468 | Fixed an issue with vxconfig sleeping and no vx commands were responding. |
| 1402144 | Fixed a system panic due to invalid pointer being passed to bcopy() by volkio_to_kio_copy. [SUN bug ID: 6742479] |
| 1397877 | Enhanced the vxresize manual page to run from non-CVM master. |
| 1397540 | Fixed an issue with the vxsnap restore manual page is unable to properly freeze or thaw file systems in a CVM environment. |
| 1393570 | Fixed a FC-Switch port failure resulting in the loss one of four paths. |
| 1393030 | Fixed an issue with the vxdiskunsetup manual page failing when the dmpnode is not the primary path. |

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0MP3 RP1 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 1389511 | Fixed issue that was unable not to force import diskgroup version 80 in VxVM 5.0. |
| 1373432 | Fixed a system panic in bcopy() due to null passed in from voliectl_copyin() |
| 1386980 | Fixed a system panic in vol_putdisk() code. |
| 1385996 | Fixed a rootdisk with B0 subdisk rendering unbootable after its removed and replaced with itself. |
| 1382977 | Fixed a system panic due to memory allocation. |
| 1375354 | Fixed an issue with vxcached never deletes old snaps when cache hits HWM. |
| 1374927 | Fixed an issue with vxvm-startup2 does not set VISSWAP flag if swap device is encapsulated and mirrored. |
| 1368752 | Fixed an issue when there are no mirrors to read, VOL_READ_MIRRORS ioctl returns -1 instead of 1. |
| | Fixed an issue with VSCSI: A/P LBI/O policy not working with enabled DMP support on boot devices. |
| 1281274 | Fixed an issue with vxplex core dumps during vxassist addlog due to DRL log length being less than 33 blocks. [SUN bug ID: 6510589, 1-19025230] |
| 1114699 | Fixed the vxtask command to display the resync progress subtask for shared volumes with DRL. |
| 1230351 | Fixed a system panic in vol_klog_start() due to accessing freed mv read_sio. [SUN bug ID: 6775184] |
| | Fixed the vxdg -n [newdg] deport [origdg] command causing a memory leak. |
| 1135462 | Fixed issue that was unable not to import disk group. |
| 1058665 | Fixed the vxdiskunsetup command failing when disk access name does not match the physical path name. |

Table 1-6 Veritas Volume Manager and Veritas Volume Replicator 5.0MP3 RP1 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 853207 | Fixed an issue with 4.1 vxclust reconfig step 2 timed out on joining; node, reconfiguration looping. [SUN bug ID: 6439209] |
| 424397 | 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 1-7](#) describes fixed issues in the Veritas Volume Manager 5.0MP3 RP2 release.

Table 1-7 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues

| Incident | Description |
|----------|--|
| 850816 | You can now delete snap objects from a mounted volume. |
| 1097258 | The vxconfigd daemon no longer hangs when an array is disconnected. |
| 1108839 | Turning on the dmp_cache_open tunable no longer slows down the vxconfigd daemon when run with 2048 dual path LUNs. |
| 1184280 | Added additional debug messages around the VE_BADPROTOV error message to improve debugging. |
| 1189199 | Fixed the cause of a system panic that occurred when you unloaded the vxdmp driver. [SUN bug ID: 6633627] |
| 1195591 | Fixed the cause of a panic when a cluster had an empty RVG. |
| 1224659 | Fixed an issue in which the vxconfigbackup -p script sometimes created a zero-length .binconfig file. |
| 1259467 | Fixed an issue in which the accept() call entered an infinite loop. [SUN bug ID: 6529218] |
| 1286298 | Fixed an issue in which proper locks were not taken in all necessary places while modifying last_sent_seqno. |
| 1287975 | The vxclustadm command has a segmentation fault when the main.cf file contains lines that are greater than 512 characters. |

Table 1-7 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 1302064 | Fixed an issue in which EFI disks could not be initialized or set up after formatting the disks. |
| 1321272 | Fixed the an issue in which some VxVM commands hung after disconnecting, then reconnecting to the FC site link. |
| 1321298 | Fixed the cause of a vxconfigd daemon core dump that occurred after reconnecting the FC site link and heartbeat link. |
| 1370927 | Fixed an issue in which the VTOC of disks in a cluster became corrupted. |
| 1374603 | Fixed a cause of data corruption in the dmp_bypass_iodone() call. |
| 1380386 | The appropriate number of I/O threads are now created for systems with more than 8 CPUs. |
| 1388883 | Fixed an issue in which rebooting a controller caused the diskgroups to be disabled. |
| 1402443 | Fixed the cause of a system panic in the kmsg_udp_payload() call. [SUN bug ID: 6730498] |
| 1408367 | Fixed the cause of a system panic when mutex_panic() was called from vol_rwsleep_wrlock(). |
| 1414336 | Fixed an issue in which some disk devices did not appear in the vxdisk list command output. |
| 1414469 | Fixed an issue in which the vxddladm listsupport all did not display up-to-date information. |
| 1416080 | Fixed the cause of a system panic in the vol_change_disk() routine that was due to NULL pointer dereference. |
| 1418659 | Fixed an issue in which a Jumpstart installation of the 4.1 MP2 and 4.1 MP2 RP3 patches created duplicate entries in the /var/svc/profile/upgrade file. |
| 1421353 | Fixed an issue in which I/O got stuck in the drl_logbusy queue due to corruption of the age node LRU list. |
| 1425338 | Fixed an issue in which connect rlinks failed to be connected, followed by vxconfigd hanging on a secondary node. |

Table 1-7 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 1437281 | Fixed the cause of an error with the <code>vxddmpadm -v getdmpnode enclosure=<name></code> command when a LUN was removed incorrectly. |
| 1446208 | Changed message V-5-1-2140 from an error message to an informational message. |
| 1450348 | Fixed a potential hang/panic that was due to a race condition between an RU thread and a volume read completing during DCM replay. |
| 1452957 | Fixed a panic in the <code>bcopy()</code> call from <code>dmp_recv_scsipkt()</code> . |
| 1457132 | Fixed the cause of data corruption when running the <code>vxddmpadm disable path</code> and <code>vxddmpadm disable ctlr</code> commands. |
| 1457758 | Fixed an issue in which the <code>vxdiskadm</code> command failed to replace a disk that was removed. |
| 1458792 | Fixed in issue in which the <code>*unit_io</code> and <code>*pref_io</code> tunables became set to 32 MB after upgrading from the Storage Foundation 5.0 MP1 release to the 5.0 MP3 release. |
| 1459831 | Fixed an issue in which replication hung due to a deadlock on a secondary that had a TCP multiconnection and was managed by <code>nmcom</code> . |
| 1461314 | DMP no longer uses the SCSI bypass on single path disks for path-suppressing TPD. [SUN bug ID: 6887215] |
| 1461717 | Fixed an issue in which the <code>vxsnap make</code> command caused the <code>vxconfigd</code> daemon to hang. |
| 1463547 | Fixed the cause of a <code>vxconfigd</code> core dump that occurred when dynamically reconfiguring a LUN. |
| 1469487 | The I/O buffer start time is no longer modified as part of error analysis. [SUN bug ID: 6778439] |
| 1471658 | Fixed the cause of a <code>vxconfigd</code> daemon core dump that occurred in the <code>priv_get_all_udid_entry()</code> call. [SUN bug ID: 6827895] |

Table 1-7 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 1471763 | Fixed the cause of the following error: build_devlink_list: readlink failed for /dev/vx/rdisk/ludg: Invalid argument [SUN bug ID: 6741762] |
| 1472736 | Fixed the cause of a system panic in the vxdmp module that was due to a NULL pointer dereference. [SUN bug ID: 6761745, 1-24676247] |
| 1473638 | Fixed the cause of a failover in the IOCTL context for coordinator disks. |
| 1475707 | Added an error message for attempting to import unwritable disks. [SUN bug ID: 6264161] |
| 1477143 | The cluster volume manager failback protocol is now triggered when cur_pri is null and at least one DMP node of the same LUN group is DMPNODE_SHARED. |
| 1479729 | Fixed the cause of an I/O hang on the primary node after a secondary node crashed. |
| 1479735 | Fixed the cause of an I/O hang on a slave if the master (logowner) crashed with a data change map active. |
| 1480315 | Fixed an issue in which VxVM performed a full re-sync of a volume that was created in the background when the volume's diskgroup was imported. |
| 1483164 | Fixed an issue in which disks with the NOLABEL state were usable via the CLI. |
| 1483201 | Fixed an issue in which the Device Discovery Layer (DDL) sometimes set the unique disk identifier (UDID) value to INVALID. Multiple disks set to INVALID resulted in the following error: VxVM vxio V-5-0-1056 new disk disk_id has a non-unique UDID |
| 1483643 | Fixed an issue in which a raid 5 volume would not start on 3PAR Thin Provisioning LUNs. |

Table 1-7 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 1484919 | Fixed an issue in which a system that was upgraded to the 5.0 MP3 release could not be booted. 1485379 Fixed an issue in which the vxtask -l list command displayed incorrect progress of the vxsnap addmir command, which was used to link a snapshot volume to the source volume. |
| 1488084 | Fixed an issue in which the vxdmpadm iostat command reported different amounts of read/write blocks than the vxstat, iostat, and sar -d commands. |
| 1500389 | The vxrootadm command now automatically enables the use-nvramrc? variable. [SUN bug ID: 6792686] |
| 1501165 | Changed the V-5-1-2140 message from an error to a warning. [SUN bug ID: 6761748] |
| 1502842 | Fixed an issue in which the dmpolicy.info file did not get updated after upgrading the packages from Storage Foundation (SF) 5.0 MP3 RP1 to SF 5.1. |
| 1503168 | Fixed an issue in which the diskgroup for disks without a private region (nopriv disks) could not be imported. |
| 1507291 | Fixed an issue in which setting the dmp_monitor_fabric value to ON triggered unexpected offlining of paths on a DMX4 array. |
| 1508462 | Fixed the cause of a vxconfigd hang that occurred due to a split brain condition on a cluster. |
| 1512352 | Fixed an issue in which the vxconfigrestore command failed with the following error: VxVM vxconfigrestore ERROR V-5-2-3706 Diskgroup configuration [SUN bug ID: 6822234] |
| 1515581 | Fixed an issue in which recreating a shared diskgroup put CVMVolDg in an empty KSTATE and offlined clustered file systems. |
| 1525121 | Fixed an issue in which EFI disks were in an error state after installing the Storage Foundation 5.0 MP3 RP1 patches. |
| 1525819 | Fixed an issue in which the vxconfigbackup command failed to work on a diskgroup that had 2 TB LUNs. |

Table 1-7 Veritas Volume Manager and Veritas Volume Replicator 5.0 MP3 RP2 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 1527247 | Fixed an issue in which the vxstat command showed twice the I/O activity on a mirror volume compared to the source volume. |
| 1528368 | Fixed the cause of an I/O hang during the data change map transition after performing vxresize operations on the primary node. |
| 1534038 | Fixed an issue in which DMP stats sometimes used invalid I/O stats entries, which led to a panic on the host. |
| 1534379 | Fixed an issue in which the vxdg split command failed with the following error: Internal configuration daemon error |
| 1544051 | Fixed an issue in which the incorrect bit was being checked for an EMC Symmetrix thin device. |
| 1586879 | Improved performance of the vxdisk online command when used on large configurations. |
| 1589022 | Fixed the cause of an infinite loop in the DMP error handling code path with a CLARIION array, which led to an I/O hang. |
| 1589172 | Fixed an issue in which the vxdisksetup and vxdiskunsetup commands sometimes failed for EFI disks. |
| 1589881 | Fixed an issue in which the dump device was changed to none (dumps disabled) after encapsulating a boot disk. |
| 1590314 | The vxdmpadm getsubpaths dmpnodename command now validates the dmpnodename value before getting the subpath information. |
| 1597868 | Fixed an issue in which, on a secondary node, rlink paused and generated the “Incorrect magic number or unexpected upid” error message, and the secondary_log_err flag got set. |
| 1598706 | Fixed the cause of a system crash that occurred while mirroring the rootdisk. |

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 1-8](#) describes fixed issues in the Veritas File System 5.1 release.

Table 1-8 Veritas File System 5.1 fixed issues

| Incident | Description |
|----------|--|
| 1477763 | The <code>qiostat -l</code> command now shows accurate hit percentages. |
| 1518001 | Fixed a panic in during a <code>vx_do_putpage()</code> call. |
| 1556692 | Fixed a <code>pagezero()</code> panic that occurred when the <code>vmodsort</code> parameter was enabled. [SUN bug ID: 6815195] |
| 1603264 | Fixed a panic that occurred while removing Storage Checkpoints, which requires a very long full <code>fsck</code> . |
| 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. |
| 1804512 | Fixed a <code>umount</code> failure that returned EBUSY. |
| 1834048 | Fixed a panic that occurred in the <code>page_unlock()</code> call. |

[Table 1-9](#) describes fixed issues in the Veritas File System 5.0 MP3 RP2 release.

Table 1-9 Veritas File System 5.0 MP3 RP2 fixed issues

| Incident | Description |
|----------|--|
| 1370823 | Fixed an issue in which running a full <code>fsck</code> did not fix a file system. |
| 1401516 | Fixed the cause of a hang that occurred after locking a file system, disconnecting the storage cable, then using <code>fsadm</code> to unlock the file system. |
| 1412465 | Fixed an issue in which the <code>vxresize</code> command failed to resize the volume, even though the file system was successfully resized. |
| 1426951 | Fixed some badly formed <code>printf()</code> statements in <code>vxm_getinfo()</code> that caused a system panic. |
| 1441487 | Changed GMS to use the standard <code>gab_api_init()</code> call to avoid a possible GAB panic. |
| 1445511 | The <code>vx_cds_control()</code> call now releases active level 1 on an error path. |

Table 1-9 Veritas File System 5.0 MP3 RP2 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 1468377 | You can now shrink a file system regardless of where the structural files reside on that file system. |
| 1484888 | Fixed an issue in which the cache hit percentage shown by <code>qiostat -l</code> command was inaccurate. |
| 1517415 | Fixed the cause of a core dump when running the <code>ncheck</code> command. |
| 1526581 | <code>vx_tflush_map()</code> no longer disables the file system if a map is marked as bad, but there is no I/O error. |
| 1588199 | Fixed an issue in which <code>dm_get_allocinfo()</code> failed with the EIO error for ext4 inodes with indirect pointers. |
| 1601187 | Reverted default <code>max_seqio_extent_size</code> to 2048, from 104857. |
| 1634788 | Fixed an issue in which the <code>fsadm</code> command dumped core intermittently when trying to defragment a file system. |

[Table 1-10](#) describes fixed issues in the Veritas File System 5.0 MP3 RP1 release.

Table 1-10 Veritas File System 5.0 MP3 RP1 fixed issues

| Incident | Description |
|----------|---|
| 1413494 | Fixed a failure of the <code>umount -f</code> command to unmount a VxFS file system. |
| 1414175 | Improved VxFS performance. |
| 1414178 | Fixed an issue with VxFS using too much CPU while looking for odd-sized extents (<code>vxi_alloc_fail</code>). |
| 1415188 | Fixed a full <code>fsck</code> core dump that was caused by running out of swap space, which resulted in a malloc failure. |
| 1417973 | Eliminated a benign error that occurred on globally-mounted VxFS file systems in a SunCluster environment when using the <code>scswitch</code> command or <code>mount</code> command. |
| 1423867 | Optimized <code>vx_convndata_files()</code> . |
| 1428661 | Improved the performance of <code>fsadm resize</code> on SFCFS. |
| 1433066 | Fixed a case of looping in <code>vx_do_putpage()</code> due to having a page beyond <code>i_wsize</code> . |

Table 1-10 Veritas File System 5.0 MP3 RP1 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 1434438 | Fixed a panic in <code>vx_unlockmap()</code> due to a null <code>ml_trap</code> pointer. |
| 1437490 | The <code>fsclustadm</code> command's <code>lltadb.c</code> is now mult-threaded safe for CFMountAgent. |

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.

Removing the VRTSodm 5.1 package may leave /dev/odm mounted in non-global zones preventing the odm module from unloading (1857357)

If you remove the VRTSodm 5.1 package, `/dev/odm` may be left mounted in non-global zones that prevents the odm module from being unloaded. This occurs if you uninstall 5.1 or if you use the `pkgrm` command to remove the VRTSodm 5.1 package manually with non-global zones configured.

Workaround

Ensure `/dev/odm` is unmounted in each non-global zone before you uninstall 5.1 or if remove the VRTSodm 5.1 package manually using the `pkgrm` command.

To unmount /dev/odm

- ◆ Unmount `/dev/odm` in each non-global zone:

```
global# zlogin myzone
myzone# umount /dev/odm
```

Clone command fails on an Oracle RAC database (1399393)

The commands `dbed_vmclondb` and `dbed_clondb` 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_clondb -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.

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 diskgroups. 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:

```
# vxdbg 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 names 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.

Veritas Live Upgrade scripts support upgrade only on UFS root file systems (1634451)

The Veritas Live Upgrade scripts, `vxlustart` and `vxlufinish`, support upgrading only on UFS root file systems. The `vxlustart` and `vxlufinish` scripts do not support upgrades on ZFS or SVM root disks.

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.

Boot fails after installing or removing Storage Foundation packages from a Solaris 9 system to a remote Solaris 10 system (1747640)

The following issue occurs if you install or remove a Storage Foundation package or patch from a Sparc Solaris 9 system to a remote Solaris 10 system, using the `-R rootpath` option of the `pkgadd`, `patchadd`, `pkgrm` or `patchrm` commands.

Generally, when you install or remove a Storage Foundation package on a Solaris 10 system, the package scripts update the boot archive. However if the local system is Solaris 9 and the remote system is Solaris 10, the scripts fail to update the boot archive on the Solaris 10 system.

Note: The boot archive is synchronized correctly when you upgrade Storage Foundation using Solaris Live Upgrade.

Workaround

The workaround is to manually clear the boot archive when you boot the alternate.

The SUN boot process detects that the boot archive is out sync and displays instructions for how to correct the situation.

For example:

WARNING: The following files in / differ from the boot archive:

```
stale //kernel/drv/sparcv9/vxportal
stale //kernel/drv/vxportal.conf
stale //kernel/fs/sparcv9/vxfs
```

```
...
new      /kernel/drv/vxlo.SunOS_5.10
new      /kernel/drv/vxlo.conf
changed /kernel/drv/vxspec.SunOS_5.9
changed /kernel/drv/vxspec.conf
```

The recommended action is to reboot to the failsafe archive to correct the above inconsistency. To accomplish this, on a GRUB-based platform, reboot and select the "Solaris failsafe" option from the boot menu. On an OBP-based platform, reboot then type "boot -F failsafe". Then follow the prompts to update the boot archive. Alternately, to continue booting at your own risk, you may clear the service by running: "svcadm clear system/boot-archive"

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.

Sybase ASE version 15.0.3 causes segmentation fault on some Solaris version (1819595)

Sybase ASE 15.0.3 produces segmentation fault on Solaris SPARC 10 Update 6 in a pure IPv6 environment. However, Sybase ASE 15.0.3 works on Solaris SPARC 10 Update 5.

When running Sybase ASE 15.0.3 GA on a pure IPv6 environment on Solaris SPARC 10 Update 6, you may receive a segmentation fault message. For example:

```
Building Adaptive Server 'CDGV240AIPV6':
Writing entry into directory services...
Directory services entry complete.
Building master device...
Segmentation Fault - core dumped
Task failed
Server 'CDGV240AIPV6' was not created.
```

This is a Sybase known issue. You should use Sybase Adaptive Server Enterprise Suite version 15.0.3 ESD 1 that supports Solaris 10 Update 6 or later. For details, refer to the Sybase Product Download Center regarding ESD 1.

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 Deported under the Disgroup tab in the SFM GUI.

Workaround

To resolve this known issue

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

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

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_rac11dg1: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dg1.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dg1 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.

Live Upgrade issue on Storage Foundation releases 4.1 and 5.0 (1849558)

When performing a Live Upgrade from Storage Foundation 4.1 or 5.0 to Storage Foundation 5.1, the `vxlufinish` script may fail on an encapsulated system. If you are upgrading from Storage Foundation 5.0MP3 to Storage Foundation 5.1, the issue does not occur because of changes in the way DMP handles alias naming.

If you encounter this issue, the `vxlufinish` script fails with the following error:

```
VxVM vxencap ERROR V-5-2-310 The c0t0d0 disk does not appear to
be prepared for this system.
RROR:vxlufinish Failed: /altroot.5.10/usr/lib/vxvm/bin/vxencap -d
-C 1 0176 -c -p 5555 -g altrootdg rootdisk=c0t0d0s2
```

In this example, `c0t0d0` is a device path of the alternate root disk.

Workaround

To work around this issue

- 1 Remove the following file to disable encapsulation of the alternate root:

```
# rm /altroot.5.10/vx_lu.5.10/.encapdisk
```

- 2 Reissue the `vxlufinish` command.
- 3 After the alternate boot environment is active, manually encapsulate the root disk using the `vxdiskadm` utility.

Rare console hangs during Live Upgrade from Solaris 9 to Solaris 10 (1728047)

When you perform a Live Upgrade from Solaris 9 to Solaris 10, the first reboot into the alternate root environment is a configuring reboot. In some very rare cases, the console may hang. The cause is usually a Solaris configuring process that is waiting for user input.

Workaround:

If there is no configuration change occurring for the Live Upgrade, select F2 (or esc-2) to respond to the configuring process.

Veritas Volume Manager known issues

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

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.

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.

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>
```

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 `vxdisk scandisks` command to bring newly discovered devices into the VxVM configuration.

CDS disks fail due to Solaris x86 bug regarding writes on the last sector of a LUN (1846165)

Solaris x86 has a bug [SUN bug ID: 6342431] where last sector of a LUN cannot be written. This causes failure with CDS disks because actual geometry cannot be set to use full capacity. It can lead to small data corruption because backup labels get written to wrong location. This bug does not exist on SPARC.

Symantec provides a workaround for SUN bug IDs 6847000 and 6844425.

Workaround

Use sliced disks for any LUN exhibiting above behavior.

<http://entsupport.symantec.com/docs/331654>

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

[SUN bug ID: 6889095]

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.

VEA starts in the default locale of the system (1630137)

On Solaris 10 and above, the VEA server starts in the default locale of the system, even if the shell invoking the command has another locale. This issue occurs if you start the VEA server with either of the following commands:

```
/opt/VRTSob/bin/vxsvcctl start
```

or

```
/usr/sbin/svcadm enable -s svc:/system/vxsvc:default
```

Work around

Use this procedure to start VEA in a locale different from the default locale of the system.

To start VEA in a locale different from the default locale of the system

- 1 Run the following commands:

```
svccfg -s system/vxsvc setenv LANG required_locale_id  
vccfg -s system/vxsvc setenv LC_ALL required_locale_id
```

For example, to start the VEA in the locale zh_CN.GB18030:

```
svccfg -s system/vxsvc setenv LANG zh_CN.GB18030  
vccfg -s system/vxsvc setenv LC_ALL zh_CN.GB18030
```

- 2 Run the following command:

```
svcadm refresh system/vxsvc
```

- 3 start the VEA server with either of the following commands:

```
/opt/VRTSob/bin/vxsvcctrl start
```

or

```
/usr/sbin/svcadm enable -s svc:/system/vxsvc:default
```

Veritas File System known issues

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

Possible error during an upgrade and when there is a local zone located on a VxFS file system(1675714)

During an upgrade and when there is local zone located on VxFS, you may receive an error message similar to the following:

```
Storage Foundation Uninstall did not complete successfully  
VRTSvxvm package failed to uninstall on pilotv240-1
```

Workaround

You must reboot after the upgrade completes.

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 -F vxfs /dev/vx/dsk/<dg>/<snapshot_volume>
```

For issue 2:

Run the `fsck` command on the replicated data volumes on the new primary site, to restore the consistency of the file system residing on the data volume.

For example:

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

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 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

vradmind 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 vradmind 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 `vradmind repstatus` and `vradmind printrvg` commands:

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

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

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

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

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

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

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:

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

On the master node.

```
# vxrvlg -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`
 - `gio_convertdbfiles -f` option that is used to check for file fragmentation
 - `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 QIO and CIO fast I/O access with DB2 or Sybase databases.

- FlashSnap Agent for Symmetrix (VxFAS), which was previously named the TimeFinder ToolKit (TFTK)

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 for Oracle RAC Release Notes* (`sfrac_notes.pdf`)

Storage Foundation guides

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

[Table 1-11](#) describes the guides in the Veritas Storage Foundation documentation set.

Table 1-11 Guides in Veritas Storage Foundation documentation set

| Guide Title | Filename |
|---|-------------------------|
| <i>Veritas Storage Foundation and High Availability Getting Started Guide</i> | getting_started.pdf |
| <i>Veritas Storage Foundation JumpStart read me</i> | jumpstart_readme.txt |
| <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 and High Availability Solutions Virtualization Guide</i> | sfha_virtualization.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 1-12](#) describes the Veritas Storage Foundation Cluster File System (CFS) documentation set.

Table 1-12 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 |
| <i>Veritas Storage Foundation Advanced Features Administrator's Guide</i> | sf_advanced_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 |
| <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 |

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

| Guide Title | Filename |
|---|--------------------------|
| <i>Dynamic Reconfiguration for Sun Enterprise Servers</i> | vcs_dynamic_reconfig.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 |

Veritas Storage Foundation for Oracle RAC documentation

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

[Table 1-15](#) describes the Storage Foundation for Oracle RAC documentation set.

Table 1-15 Guides in Storage Foundation for Oracle RAC documentation set

| Guide Title | Filename |
|---|-------------------|
| <i>Veritas Storage Foundation™ for Oracle RAC Installation and Configuration Guide</i> | sfrac_install.pdf |
| <i>Veritas Storage Foundation™ for Oracle RAC Release Notes</i> | sfrac_notes.pdf |
| <i>Veritas Storage Foundation™ for Oracle RAC Administrator's Guide</i> | sfrac_admin.pdf |
| <i>Veritas Storage Foundation: Storage and Availability Management for Oracle Databases</i> | sf_adv_ora.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
```