

Veritas Storage Foundation™ Release Notes

Linux

5.0 Release Update 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: 5.0 RU1

Document version: 5.0RU1.1

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

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/techsupp/

Contacting Technical Support

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

www.symantec.com/techsupp/

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

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

Licensing and registration

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

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

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

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](#)
- [Installation](#)
- [Required patches](#)
- [System requirements](#)
- [Component product release notes](#)
- [Software limitations](#)
- [Fixed issues](#)
- [Known issues](#)
- [No longer supported](#)
- [Documentation](#)
- [Documentation errata](#)

Overview of this release

This document provides release information about the products in the Veritas Storage Foundation 5.0 Release Update 1 (RU1) product line:

- Veritas Storage Foundation™ (Basic, Standard, Standard HA, Enterprise, and Enterprise HA)

- Veritas™ Volume Manager (VxVM)
- Veritas™ File System (VxFS)
- Veritas™ Volume Replicator (VVR)
- Veritas Storage Foundation™ Cluster File System (SFCFS)

Each of these products is activated by a single license key. You must obtain a license key before installing the product.

See the *Veritas Storage Foundation Installation Guide*.

The Veritas Storage Foundation 5.0 RU1 release is based on 5.0 MP3 RP2 and includes support for SuSE Linux Enterprise Server (SLES) 11. All issues in RU1 are cumulative since the 5.0 release. However, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation Cluster File System RAC are not available in the Veritas Storage Foundation 5.0 RU1 release because Oracle and DB2 have not yet provided support for databases on SLES 11. Additionally, there is no support for Xen kernel on SLES11 in the RU1 release.

For the latest information on updates, patches, and known issues regarding this release, see the following TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/281993>

The hardware compatibility list (HCL) is available at:

<http://entsupport.symantec.com/docs/324247>

The hardware TechNote is available at:

<http://entsupport.symantec.com/docs/283282>

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

Changes in Storage Foundation

This section describes the changes in Veritas Storage Foundation 5.0 Release Update 1 (RU1).

Support for SuSE Linux Enterprise Server 11

The Storage Foundation software stack has been enhanced in 5.0 Release Update 1 to support SuSE Linux Enterprise Server 11.

Storage Foundation

Storage Foundation 5.0 Release Update (RU) 1 is cumulative and based on Veritas Storage Foundation 5.0 MP3 RP2 and includes support for SuSE Linux Enterprise Server (SLES) 11. Review previous versions of the Storage Foundation 5.0 (Linux) Release Notes at the following URL:

<http://entsupport.symantec.com/docs/306947>

Storage Foundation and Storage Foundation Cluster File System documentation changes

There are changes to the Storage Foundation and Storage Foundation Cluster File System documentation. The *Veritas Storage Foundation Cluster File System Installation Guide* ([sf_cfs_install.pdf](#)) and the *Veritas Storage Foundation Cluster File System Release Notes* ([sf_cfs_notes.pdf](#)) no longer exist. The Storage Foundation Cluster File System installation content exists in the *Veritas Storage Foundation Installation Guide* ([sf_install.pdf](#)) and the Storage Foundation Cluster File System Release Notes content exists in the *Veritas Storage Foundation Release Notes* ([sf_notes.pdf](#)).

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

In the 5.0 RU1 release, this feature is turned on automatically for Thin disks and off by default for non-Thin disks. To turn on and use the full functionality the SmartMove feature, refer to the Volume Manager Administration Guide.

To use the SmartMove feature, VxVM and VxFS must be version 5.0 MP3 or later.

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.

Veritas Volume Manager

Veritas Volume Manager includes the following changes in 5.0 Release Update 1:

Support for iSCSI devices

Veritas Volume Manager now provides support for the use of iSCSI devices in DMP. VxVM provides the same functionality available on Fibre Channel devices for iSCSI devices. This feature is available on all the supported iSCSI arrays. In addition, VxVM provides new interfaces to obtain Fibre Channel and iSCSI configuration information visible to the host. Wherever support from the operating system is available, you can also view and set iSCSI session parameters to improve the performance of iSCSI devices.

Boot disk recovery

The process to create a failback disk has been simplified. You can now use the VxVM commands to break off the root mirrors and make the mirror disk bootable.

The resulting mirror of the boot disk can be used as a failback disk in case of upgrade failure.

See the `vxrootadm(1M)` manual page for more information

Enhancements to the Dynamic Multipathing feature

This release provides a number of enhancements to the Dynamic Multipathing (DMP) features of VxVM. These enhancements simplify administration, and improve display of detailed information about the connected storage.

Improved Dynamic Multipathing device naming

The DMP device naming feature has been enhanced to provide a more consistent and user friendly approach for naming the DMP devices.

The following enhancements apply regardless of the specified naming scheme:

- DMP now enables you to assign customized names for DMP devices. You can specify customized names for individual devices, or you can use a file containing user-defined names to assign multiple names.
- You can specify a DMP device name to commands using the name of any of its subpaths. The output displays the DMP device name assigned.
- In a symmetric cluster, the DDL-generated enclosure-based names for DMP devices are now consistent across all the nodes in the cluster.
- Device names can be made persistent. This is the default for the enclosure-based naming (EBN) scheme.

The following enhancements apply to the EBN naming scheme:

- DDL generates the device name in the format *enclosure_index*. If you specify the `use_avid` argument, the name is generated with the Array Volume ID for the index number to provide a more meaningful name.

Default behavior for I/O throttling

By default, DMP is now configured with no I/O throttling. In previous releases, I/O throttling was set to on. Use the `vxdmppadm setattr` command with the `recoveryoption` keyword to configure I/O throttling for DMP.

Specifying a minimum number of active paths

You can now configure a minimum redundancy level, which is the minimum number of paths for the devices under an enclosure. Use the `redundancy` option of the `vxdmppadm getdmpnode` command to display any devices that have fewer than the minimum number of paths. You can also configure DMP to notify you when the number of active paths falls below the configured minimum.

Enhanced subpaths listing

The `vxdmpadm getsubpaths` command now provides the ability to list all subpaths known to DMP, subpaths of an enclosure, or subpaths through an array port or pwwn. To list the paths through an array port, specify either a combination of enclosure name and array port id, or the array port WWN.

The default listing of the `vxdmpadm getsubpaths` command is sorted by enclosure name, then by DMP node and within that by pathname. The new option `-s` enables you to sort the output based on path name, DMP node name, enclosure name, or host controller name.

Enhanced I/O statistics

The following enhancements have been made to I/O statistics:

Queued and Erroneous I/O counts

The `vxdmpadm iostat show` command now provides options to display queued I/O counts (`-q` option) and erroneous I/O counts (`-e` option). These options are applicable for DMP node, path and controller.

Filter zero entries

The `vxdmpadm iostat show` command now provides the `-z` option to filter out entities for which all data entries are zero. This option is especially useful in a cluster environment, when many paths are required for failover capabilities, but the paths are not being used for I/O.

Specifying units for statistics data

You can now specify the units in which the statistics data is displayed. The `-u` option accepts `k`, `m` and `g` arguments to display throughput in kilo-, mega-, and giga- system blocks. The `us` argument displays average read/write time in microseconds. By default, the read/write times are displayed in milliseconds up to 2 decimal places. The throughput data is displayed in terms of 'BLOCKS' and the output is scaled, meaning that the small values are displayed in small units and the larger values are displayed in bigger units, keeping significant digits constant. The `bytes` argument to `-u` option can be used to display throughput in exact number of bytes.

Cumulative I/O statistics

The `vxdmpadm iostat` command now has a `groupby` clause to provide cumulative I/O statistics listing per `dmpnode`, controller, array port id, host-array controller pair and enclosure. If the `groupby` clause is not specified then the statistics are displayed per path.

Miscellaneous improvements to DMP I/O statistics

The following improvements have been made to the DMP I/O statistics output:

The way in which average read/write time is calculated has been corrected.

By default, the average read/write time is displayed in milliseconds up to two places after the decimal point. Use the new option `-u us` to display the average read/write time in microseconds.

The average I/O size is set to 512 bytes.

If the output exceeds 80 columns, then the output is wrapped.

Making DMP restore options persistent

The restore policy, restore interval, and restore period are now persistent across reboot. In addition to being set as options to the `vxddmpadm start restore` command, these attributes can also be set using the `vxddmpadm settune` command. The new tunables are: `dmp_restore_policy`, `dmp_restore_interval`, and `dmp_restore_cycles`.

In addition, there is a new tunable, `dmp_enable_restore`, which enables the path restoration thread to be started.

New log file location for DMP events

The log file location for DMP events is `/var/adm/vx/dmpevents.log`. For backward compatibility, `/etc/vx/dmpevents.log` is a soft link to `/var/adm/vx/dmpevents.log`.

Extended device attributes displayed in vxdisk list

The `vxdisk list` command now displays extended device attributes like hardware mirrors for certain arrays.

Display use_all_paths attribute for an enclosure

Display value of `use_all_paths` attribute for an enclosure.

Viewing information about the ASLs installed on the system

The `/usr/lib/vxvm/diag.d/vxcheckasl` command has been enhanced to provide all the information regarding the ASLs (all those installed in the system), the devices (all seen by OS) and all the possible ways in which these ASLs can interact with these devices.

Displaying the count of LUNs in an enclosure

The `vxddmpadm listenclosure` command now displays the count of LUNs in its default output.

Displaying LUN serial number

The `vxddmpadm getdmpnode` command now includes the option `-v` to display the LUN serial number along with other information.

Displaying HBA details

The `vxddmpadm getctlr` output has been enhanced to display HBA vendor details and the Controller ID. For iSCSI devices, the Controller ID is the IQN or IEEE-format based name. For FC devices, the Controller ID is the WWN. Because the WWN is obtained from Event Source Daemon, this field is blank if the Event Source Daemon is not running.

New exclude and include options for the vxddmpadm command

The `vxddmpadm` command now includes `exclude` and `include` commands to suppress or unsuppress devices from VxVM, respectively.

This provides a command line interface for these operations, which previously required user interaction.

New command for reporting DMP node information

The `vxddmpadm` command now includes the `list` command to display information about a DMP node, including all of the attributes that are set for that DMP node. The `vxddmpadm list` command can be used for a specified `dmpnode`, all `dmpnodes`, all `dmpnodes` on the path name or `dmpnodename`, or all `dmpnodes` in an enclosure.

Setting attributes for all enclosures

The `vxddmpadm setattr` command now has the `all` option for enclosure, array type and arrayname. The `all` option allows you to set the attributes (`iopolicy`, `failover_policy`, `recoveryoption`) on all the enclosures specified. Also, `vxddmpadm setattr arraytype array_type` sets the attribute for all array types derived from the given `array_type`.

Support for ALUA JBOD devices

DDL has now improved the support for JBOD devices to include ALUA JBOD devices. DMP now provides immediate basic support for any ALUA compliant array.

Full support still requires an array support library (ASL) for that array. See the Hardware Compatibility List (HCL) for details about supported arrays.

Enhancements to the Cluster Volume Manager

Cluster Volume Manager has been enhanced.

Local detach policy now supported with Veritas Cluster Server clusters and with Dynamic Multipathing Active/Passive arrays

In a CVM environment, it is possible that a failure of site is local to the node doing the I/O. In such cases, a shared disk group can be associated with a “local detach” policy. With this option, the node that encounters the site failure requests another node or a set of nodes to verify if the site is accessible. If the failure is local to the node, the volume is disabled on that node without doing any configuration updates over the cluster. With the local detach policy, the behavior of cluster configuration with sites is the same as with regular cluster configuration.

Prior to this release, the local detach policy was not recommended for certain environments. This release removes the restrictions for using the local detach policy, as follows:

The local detach policy can now be used in CVM or CFS clusters which are managed by VCS. The new behavior notifies the VCS agents of the local failure.

Also, the local detach policy can now be used in a DMP environment with Active/Passive arrays.

Distributed Volume Recovery

In Cluster Volume Manager (CVM), upon a node crash, the mirror recovery is initiated by the CVM master. Prior to this release, the CVM master also performed all of the recovery I/O. In this release, the CVM master can distribute recovery tasks to other nodes in the cluster. Distributing the recovery tasks is desirable in some situations so that the CVM master can avoid an I/O or CPU bottleneck.

When distribution of volume recovery is turned on, the master distributes recovery tasks in a round-robin fashion to other nodes in the cluster. By default, all of the nodes in the cluster can participate in volume recovery. You can also exclude particular nodes from the volume recovery.

To turn on the distribution of volume recovery, add the keyword `distribute` in the file `/etc/default/vxrecover`. You can also specify the `distribute` keyword when running `vxrecover`.

To disable this feature, remove the `distribute` keyword from the file `/etc/default/vxrecover`. If the feature is not enabled, the master performs the resynchronization tasks locally, as in the previous release.

The distributed volume resynchronization functionality does not depend on the disk group version, and hence works with older version disk groups.

Campus Cluster enhancements

The Campus Cluster feature provides the capability of mirroring volumes across sites, with hosts connected to storage at all sites through a Fibre Channel network.

In this release, the following enhancements have been made to the Campus Cluster feature:

Site tagging of disks or enclosures

The following enhancements to `vxdisk` are related to site tagging:

- Site tagging operations on multiple disks or enclosures are now supported.
- New option to rename a site tag on a disk or enclosure.

Automatic site tagging

The `vx dg settag` command now provides an option for automatic tagging of a site. You can specify that an enclosure is automatically tagged with a particular site name. When you add a LUN belonging to that enclosure to a disk group, the LUN is tagged with the site name specified.

Site renaming

The `vx dg` command has a new `renamesite` option. The `renamesite` option renames the existing site record that is configured on the disk group. The `renamesite` option also associates all of the volume objects to the new site.

Veritas File System

Veritas File System includes the following changes in 5.0 Release Update 1:

Dynamic Storage Tiering enhancements

The Dynamic Storage Tiering (DST) has the following enhancements in this release:

- Dynamic Storage Tiering APIs are enhanced to provide a new interface for managing allocation policies of Storage Checkpoints at creation time and later, and for managing named data stream allocation policies.
- The `fspadm` command now supports UID, GID, and TAG elements in the placement policy XML file.
- Improved `fspadm` command scan performance.
- Suppressed the processing of the chosen RULE.
- Parser support for UID, GID, and TAG elements in a DST policy.

- What-if support for `analyze` and `enforce` without requiring the policy to be assigned.
- Storage Checkpoint data placement support in a DST policy.
- Upgrade to SQLite 3.3.9
- Shared DB thread handle support
- CPU and I/O throttling support for DST scans.
- New command, `fstag`, for file tagging.
- New command, `fspmk`, for creating XML policies.

Loopback device support

This release introduces loopback device support. However, loopback device support has the following limitations and restrictions:

- VxFS does not support loopback devices with encryption on Red Hat Enterprise Linux due to a limitation of Red Hat (Bugzilla 56698).
- VxFS does not support running a loopback device backed by a file.
- VxFS does not support running Storage Foundation Cluster File System on top of loopback devices.

Mount `mntlock` and `mntunlock` options

You can specify the `mntlock` option with the `mount` command, which prevents a file system from being unmounted by an application. This option is useful for an application that does not want the file systems that the application is monitoring to be improperly unmounted by other applications or administrators. Clustering applications, such as Veritas Cluster Server (VCS), are particularly expected to benefit.

The `mntunlock` option of the `vxumount` command reverses the `mntlock` option if you previously locked the file system.

Sendfile support

This release supports the use of the `sendfile()` system call with VxFS. You no longer need to set the Apache configuration option `EnableSendfile Off` when using Apache with VxFS.

Write performance improvement

Depending on the features active on a file system, write performance is improved as much as 10% compared to the 5.0 Maintenance Pack 1 release.

Veritas Storage Foundation Cluster File System

Veritas Storage Foundation Cluster File System includes the following changes in 5.0 Release Update 1:

Number of parallel fsck threads to run during recovery is tunable

In prior releases the number of parallel fsck threads that could be active during recovery was set to 4. In this release the default depends on the number of CPUs in the system, but is tunable within given limits.

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

Upgrading the kernel on a root encapsulated disk

Operating system vendors often release maintenance patches to their products to address security issues and other minor product defects. They may require customers to apply these patches regularly to conform with maintenance contracts or to be eligible for vendor support. Prior to this release, you could not install a kernel patch or upgrade on a root encapsulated system; you needed to unencapsulate the system, apply the upgrade, then reencapsulate the root disk. You can now upgrade the operating system kernel on a root encapsulated system.

Installation

Storage Foundation 5.0 Release Update 1 provides support for SuSE Linux Enterprise Server 11 (SLES 11). This release is based on Veritas Storage Foundation 5.0 MP3 RP1 and RP2.

The following sections summarize the changes since the Veritas Storage Foundation 5.0 release.

Review previous installation-specific Veritas Storage Foundation 5.0 MP3 Release Notes information at:

<http://entsupport.symantec.com/docs/306947>

Veritas Installation Assessment Service

The Veritas Installation Assessment Service (VIAS) utility assists you in getting ready for a Veritas Storage Foundation and High Availability Solutions installation. The VIAS utility allows the preinstallation evaluation of a configuration, to validate it prior to starting an installation or upgrade.

<https://vias.symantec.com/>

Simplified installation and configuration

Installation and configuration procedures have been simplified, based on usability testing.

Installation quick reference

The product installer displays a menu that simplifies the selection of installation options. Select a product to install from the menu to invoke that product's installation script.

Table 1-1 provides a quick overview of a stand-alone installation using the product installer.

Table 1-1 Installation overview

Installation task	For more information, refer to the following section:
Obtain product licenses.	See the <i>Veritas Storage Foundation™ and High Availability Solutions Getting Started Guide</i> Product licensing section for more information.
Download the software, or insert the product DVD.	See the <i>Veritas Storage Foundation™ and High Availability Solutions Getting Started Guide</i> Downloading the software and Mounting a software disc sections for more information.
Set environment variables.	See the <i>Veritas Storage Foundation™ and High Availability Solutions Getting Started Guide</i> Set environment variables section for more information.
Configure the secure shell (SSH) on all nodes.	See the <i>Veritas Storage Foundation™ and High Availability Solutions Getting Started Guide</i> Configuring the secure shell (SSH) section for more information.

Table 1-1 Installation overview (*continued*)

Installation task	For more information, refer to the following section:
Verify that hardware, software, and operating system requirements are met.	See the <i>Veritas Storage Foundation™ and High Availability Solutions Getting Started Guide</i> Verifying requirements section for more information.
Check that sufficient disk space is available.	See the <i>Veritas Storage Foundation™ and High Availability Solutions Getting Started Guide</i> Check disk space section for more information.
Use the installer to install the products.	See “Installing and configuring Storage Foundation using the common product installer” on page 20.

Installing and configuring Storage Foundation using the common product installer

The Veritas product installer is the recommended method to license and install Storage Foundation.

The following sample procedure is based on the installation of Storage Foundation on a single system.

The Veritas 5.0 RU1 release operates on the following operating systems and hardware:

SuSE Linux Enterprise Server 11 (SLES 11) with kernel number 2.6.27.19-5 or later on AMD Opteron or Intel Xeon EM64T (x86_64).

Note: This RU1 release supports only fresh installation of Storage Foundation onto SuSE Linux Enterprise Server 11; upgrade paths from SuSE Linux Enterprise Server 9 or SuSE Linux Enterprise Server 10 are not supported. When performing a manual install of Storage Foundation packages that use driver modules (VM, FS, and VCS), an `unsupported-modules` error may occur; refer to `/etc/modprobe.d/unsupported-modules` for the workaround.

To install Storage Foundation

- 1 To install on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.

- 2 Load and mount the software disc.
- 3 Move to the top-level directory on the disc.

```
# cd /mnt/cdrom
```

- 4 From this directory, type the following command to install on the local system only. Also use this command to install on remote systems using the secure shell (ssh) utilities:

```
# ./installer
```

If you use the remote shell utilities to install on remote systems, additionally specify the `-rsh` option:

```
# ./installer -rsh
```

The sample installation assumes that ssh is used.

- 5 Enter `I` to install and press Return.
- 6 When the list of available products is displayed, select Veritas Storage Foundation, enter the corresponding number, and press Return.
- 7 You are prompted to enter the system names (in the following example, "host1") on which the software is to be installed. Enter the system name or names and then press Return.

```
Enter the system names separated by spaces on which to  
install SF: host1
```

- 8 Enter the product license information.

Each system requires a product license before installation. License keys for additional product features should also be added at this time.

```
Enter a SF license key for host1:
```

```
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X
```

```
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered on  
host1
```

```
SF license registered on host1
```

- 9** You are prompted to enter additional license information, until all licenses for all systems have been entered. Then reply that you have no additional licenses to enter.

```
Do you want to enter another license key for host1?  
[y,n,q] (n) n
```

- 10** You can choose to install required or all . Optional include man pages, for example.

Each option displays the disk space that is required for installation. Select which option you want to install and press Return.

For example, you should see output similar to the following:

```
SF can be installed without optional rpms to conserve  
disk space.
```

```
1) Install required Veritas Storage Foundation rpms -  
   491 MB required  
2) Install all Veritas Storage Foundation rpms -  
   625 MB required
```

```
Select the rpms to be installed on all systems?  
[1-2,q,?] (2) 2
```

- 11** Configure Storage Foundation when prompted.

```
Are you ready to configure SF? [y,n,q] (y) y
```

- 12** You have the option of specifying the default name of a disk group. If you specify a name, it is used for Veritas Volume Manager commands when a disk group is not specified.

Enter **n** if you do not want to specify the name of the default disk group at this time. You can set the name of the default disk group after installation.

If you specify the name of a default disk group, this step does not create the disk group. After installation, you can use the `vxdiskadm` command to create the disk group.

```
Do you want to set up a default disk group for each  
system? [y,n,q,?] (y) y
```

- 13** If you responded `y`, then enter the information for the default disk group name.

```
Specify a default disk group name for system host1. [?] dg001
```

- 14** You are prompted to confirm the default disk group.

Note: If `nodg` is displayed, then the host will be configured to have no default disk group.

```
Is this correct? [y,n,q] (y) y
```

- 15** Verify the fully qualified hostname of the systems.

```
Is the fully qualified hostname of system  
"host1" = "host1.domain_name"? [y,n,q] (y) y
```

- 16** This product can be configured as a Storage Foundation Manager managed host or a stand-alone host.

Several prerequisites are necessary to configure the system as a Storage Foundation Manager managed host.

See the *Veritas Storage Foundation Manager Installation Guide* for more information.

```
Enable Veritas Storage Foundation Management Server Management?  
[y,n,q] (y) n
```

- 17** Enabling Veritas Storage Foundation Management Server management simplifies and improves management of the complex data center resources, reducing planned and unplanned down time.

This product can be configured in one of the following two ways:

- Veritas Storage Foundation Management Server managed host (Recommended)
- Standalone host

Pre-requisites for Veritas Storage Foundation Management Server managed host:

- Veritas Storage Foundation Management Server (SFM) is configured
- SFM is currently running and fully operational
- The Authentication Broker used by SFM is running

There are no pre-requisites for deploying as a standalone host.

To enable centralized management using Storage Foundation Manager, download Veritas Storage Foundation Management Server from:

[http:// www.symantec.com/sfm](http://www.symantec.com/sfm)

See the *Veritas Storage Foundation Manager Installation Guide* for more information.

- 18** The installation and configuration complete automatically, and the processes are started.

Check the log file, if needed, to confirm the installation and configuration.

Installation log files, summary file, and response file are saved at:

```
/opt/VRTS/install/logs/installer-****
```

Note: If you choose not to use the Veritas product installer, you will need to edit `allow_unsupported_modules` in `/etc/modprobe.d/unsupported-modules`. Please refer to Novell support document 7002793.

Support for SuSE Linux Enterprise Server 11

The Storage Foundation software stack has been enhanced in 5.0 Release Update 1 to support SuSE Linux Enterprise Server 11.

Required patches

There are no required patches for this release.

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems for Storage Foundation 5.0 Release Update 1

This section describes the Storage Foundation supported Linux operating system for this release.

Storage Foundation operates on the following Linux operating system and kernels distributed by SuSE:

- SuSE Linux Enterprise Server 11 (SLES 11) with kernel 2.6.27.19-5 or later on AMD Opteron or Intel Xeon EM64T (x86_64).

This 5.0 RU1 release supports only SuSE Linux Enterprise Server 11. You must perform a fresh installation of Veritas Storage Foundation software onto SLES 11.

VMware Environment support

The Storage Foundation 5.0 RU1 release is tested with the VMware ESX Server 3.5. For information about the use of the 5.0 RU1 product in a VMware Environment, refer to the following URL:

<http://entsupport.symantec.com/docs/289033>

Storage Foundation Cluster File System node requirements

All nodes in a Cluster File System must have the same operating system version SuSE Linux Enterprise Server 11 GMC installed.

Component product release notes

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

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

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

Software limitations

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

Veritas Storage Foundation software limitations

Software limitations in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

<http://entsupport.symantec.com/docs/283859>

5.0 Release Update 1 Veritas Storage Foundation software limitations

The following are additional software limitations in the 5.0 Release Update 1 release of Veritas Storage Foundation.

Xen platform, Security-enhanced Linux, and database support limitations

The Veritas Storage Foundation 5.0 RU1 release on SuSE Linux Enterprise Server (SLES) 11 is not supported on the Xen platform for Linux.

There is no support for the Xen kernel on SLES11 in this release.

Security-enhanced Linux (SE Linux) is not supported in the Veritas Storage Foundation 5.0 RU1 release.

Because of the database and Xen support limitations on SuSE Linux Enterprise Server (SLES) 11, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation Cluster File System for Oracle RAC is not available on SuSE Linux Enterprise Server (SLES) 11 at the time of this release.

For the latest information on updates, patches, and known issues regarding this release, see the following TechNote on the Symantec Technical Support Web site :

<http://entsupport.symantec.com/docs/281993>

VEA may fail to function (1192238)

VEA processes fail to function on SLES9 (SP3 or later) and SLES10 (SP1 and SP2) with the existing `libgcc` rpm version shipped by SuSE.

This issue is fixed in SLES 11 by default.

Workaround

For the VEA processes to function, you have to download the new `libgcc42-32bit-4.2.1_20070724-17.x86_64.rpm` and apply it.

To download the new rpm and apply it

- 1 Refer to the TechNote for the new rpm download link::

<http://entsupport.symantec.com/docs/281993>

- 2 Apply the new rpm::

```
# rpm -Uvh libgcc42-32bit-4.2.1_20070724-17.x86_64.rpm --force
```

odmmkfile command must be run as a privileged user

The `odmmkfile` command must be run as a privileged user. This behavior has changed from the previous releases where this command is run by any user.

5.0 MP2 Veritas Storage Foundation software limitations

There are no additional Veritas Storage Foundation software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas Storage Foundation software limitations

There are no additional Veritas Storage Foundation software limitations in the 5.0 MP1 release.

Veritas Volume Manager software limitations

Software limitations for the 5.0 RU1 release are cumulative and listed by product in this section. Software limitations in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

<http://entsupport.symantec.com/docs/283859>

5.0 RU1 Veritas Volume Manager software limitations

There are no new additional software limitations in this 5.0 RU1 release of Veritas Volume Manager.

5.0 MP3 Veritas Volume Manager software limitations

The following incident is the new software limitation for the 5.0 MP3 release of Veritas Volume Manager.

Devices and some paths are not discovered properly with IBM's DS4700 disk array after a reboot (1205369)

On a system with IBM's DS4700 disk array, use fewer than 30 LUNs to ensure that the disk array discovers all of the devices and paths after a reboot.

5.0 MP2 Veritas Volume Manager software limitations

There are no additional Veritas Volume Manager software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas Volume Manager software limitations

There are no additional Veritas Volume Manager software limitations in the 5.0 MP1 release.

Veritas File System software limitations

Software limitations in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

<http://entsupport.symantec.com/docs/283859>

5.0 Release Update 1 Veritas File system software limitations

The following are new additional software limitations in this 5.0 Release Update 1 release of Veritas File System.

Support of 32 terabyte file systems

Only Veritas Storage Foundation Enterprise and Veritas Storage Foundation Enterprise HA support file systems that are greater than 32 TB.

5.0 MP3 Veritas File system software limitations

There are no additional Veritas File System software limitations in the 5.0 MP3 release.

5.0 MP2 Veritas File system software limitations

There are no additional Veritas File System software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas File system software limitations

There are no additional Veritas File System software limitations in the 5.0 MP1 release.

Veritas Storage Foundation Cluster File System software limitations

Software limitations in the Veritas Storage Foundation Cluster File System are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/283857>

5.0 Release Update 1 Veritas Storage Foundation Cluster File System software limitations

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

Compatibility with previous versions of Veritas File System

A disk layout Version 7 file system created with VxFS 5.0 software will not be accessible if the VxFS 5.0 file system software is removed and the system is reverted to VxFS 4.1.

Quick I/O, ODM, mount -o cio, and the VX_CONCURRENT advisory are mutually exclusive

The `VX_CONCURRENT` advisory cannot be set on a file that is actively open by Quick I/O or ODM. A file that has the `VX_CONCURRENT` advisory set may not be concurrently opened by Quick I/O or ODM. Quick I/O and ODM access are not allowed for any files on a file system that is mounted with the `-o cio` mount option.

Consistent distribution and kernel version for Storage Foundation Cluster File System

All the nodes in a SFCFS cluster must be at the same OS version and patch level. In addition, mixing nodes running 32-bit kernel with nodes running 64-bit kernel is not supported with SFCFS.

5.0 MP2 Veritas Storage Foundation Cluster File System software limitations

There are no additional software limitations in the 5.0 MP2 release of Veritas Storage Foundation Cluster File System.

5.0 MP1 Veritas Storage Foundation Cluster File System software limitations

There are no additional Veritas Storage Foundation Cluster File System software limitations in the 5.0 MP1 release.

Veritas Volume Replicator software limitations

Software limitations in the 5.0 release are listed in the *Veritas Volume Replicator 5.0 Release Notes*, which is available at the following URL.

<http://entsupport.symantec.com/docs/290244>

5.0 MP3 Veritas Volume Replicator software limitations

There are no new additional Veritas Volume Replicator software limitations in this release.

5.0 MP2 Veritas Volume Replicator software limitations

There are no additional Veritas Volume Replicator software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas Volume Replicator software limitations

There are no Veritas Volume Replicator software limitations in the 5.0 MP1 release.

5.0 Veritas Volume Replicator software limitations

The following are additional software limitations in the 5.0 release of Veritas Volume Replicator.

RAID-5 volume

VVR does not support Volume Manager RAID-5 volumes as part of RVGs. Hardware RAID-5 is supported.

Disk Group Split and Join

VVR does not support Disk Group Split and Join. This means that you cannot use Disk Group Split and Join on data volumes in an RVG. However, you can take snapshots of data volumes and use DGSJ on the snapshots.

Importing a disk group with VVR objects

If a private disk group containing VVR objects is imported on two nodes, as well as the problems documented for VxVM objects, the SRL is corrupted and a full resynchronization of the Secondary is required.

Volumes in boot disk group

Symantec does not recommend having replicated volumes in the boot disk group (`bootdg`).

Selecting records using search expressions

Selecting RVG and RLINK records using search expressions with the `vxprint -e pattern` command is not supported.

Adding a Secondary

When adding a Secondary to the RDS, the `vradmin addsec` command cannot be entered from the host being added. If the RDS contains the Primary only, the command must be entered on the Primary.

Issues related to replication in a shared environment

The following issues relate to replication in a shared environment:

Creating a Primary RVG when the data volumes and SRL are on a shared disk

When creating a Primary whose data volumes and SRL are located on a shared disk group, the `vradmin createpri` command must be issued on the master node of the cluster.

Creating a Secondary RVG when the data volumes and SRL are on a shared disk

When adding a Secondary whose data volumes and SRL are located on a shared disk group to a RDS, the `vradmin addsec` command requires the Secondary host name must be resolvable and up on the master node of the Secondary cluster.

VRW limitations

The following VRW limitations exist.

Increasing font size

Increasing the font size may cause problems in viewing and navigating the wizards.

IBC messaging support

This release of VRW does not support In-Band Control (IBC) messaging.

VCS Agents for VVR limitations

The following VCS Agents for VVR limitations exist.

Onlining and offlining the RVG resources

Currently a problem with Veritas Volume Manager affects the RVG resources. When many RVGs are defined, the number of requests to the Volume Manager to online or offline can overload the `vxconfigd` process of VxVM, preventing some RVGs from going online or offline.

Workaround

The RVG resources have been set to allow only a single instance of an RVG resource to be brought online or offline. This slows down the onlining and offlining of service groups containing RVGs.

Limitations in the RVGSnapshot agent

Fire drill setup (`fdsetup`) supports applications using one disk group only. The RVGSnapshot agent does not support volume sets.

Fixed issues

The following sections describe Veritas Storage Foundation issues that were fixed in this 5.0 Release Update 1 release.

Veritas Storage Foundation fixed issues

The fixed issues for the 5.0 Release Update 1 release are listed by product. Fixed issues in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

<http://entsupport.symantec.com/docs/283859>

5.0 RU1 Veritas Storage Foundation fixed issues

Storage Foundation 5.0 RU1 is based on Veritas Storage Foundation 5.0 MP3 RP2 and includes support for SLES 11. To review the Veritas Storage Foundation 5.0 MP3 Release Notes, use the following URL:

<http://entsupport.symantec.com/docs/306947>

Table 1-2 5.0 RU1 Veritas Storage Foundation fixed issues

Incident	Description
1192238	VEA processes fail to function on SLES9 (SP3 or later) and SLES10 (SP1 and SP2) with the existing <code>libgcc</code> rpm version shipped by SuSE. This issue is fixed in SLES 11 by default.

5.0 MP3 Veritas Storage Foundation fixed issues

Table 1-3 Storage Foundation 5.0 MP3 RP2 fixed issues

Incident	Description
1366109	Fixed an SE Linux-related issue for mapping provider.

5.0 MP2 Veritas Storage Foundation fixed issues

There are no additional Veritas Storage Foundation fixed issues in the 5.0 MP2 release.

5.0 MP1 Veritas Storage Foundation fixed issues

There are no additional Veritas Storage Foundation fixed issues in the 5.0 MP1 release.

Veritas Volume Manager fixed issues

The following sections list Veritas Volume Manager fixed issues for Storage Foundation 5.0 Release Update 1 release. Fixed issues in the 5.0 release are listed in the Veritas Storage Foundation 5.0 Release Notes, which is available at the following URL:

<http://entsupport.symantec.com/docs/283859>

Veritas Volume Manager 5.0 MP3 fixed issues

Table 1-4 describes fixed issues in the Veritas Volume Manager 5.0 MP3, RP1, and RP2 releases.

Table 1-4 Veritas Volume Manager fixed issues

Incident	Description
425273	Fixed an issue with VVR RU thread not starting nio after it is created from than waiting for all replicas to have NIO's created.
536853	If the vxdiskadm command is used to replace a disk in site-consistent disk group, the new disk was expected to be tagged with the same site name as the disk that is being replaced.
592685	Fixed a vxdumpadm getdmpnode help usage error.
602953	In the VEA client for Microsoft Windows systems, existing volume tags were not displayed when adding a new volume tag.

Table 1-4 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
607096	Incorrect error message when importing a disk group.
779060	Enhanced VxVM to perform faster failover/plex detach for mirrored volumes.
797829	Encapsulated root disk could not be mirrored.
840452	Search did not return any objects for non-Administrator users.
934142	Fixed a failure that occurred when importing a version 60 diskgroup.
971002	Fixed a failure that occurred when using the <code>vxdmadm setattr</code> command to update the dynamic multipathing events.log.
973256	VxVM displayed the same dynamic multipathing node for two different LUNs.
990003	<code>vxdtl upgrade</code> would hang during an upgrade from VxVM 4.1 to 5.0.
993551	Growing a volume with a DCO association caused <code>vxconfigd</code> to hang.
994127	The <code>vxvltune</code> command could not set <code>maxiosize</code> to a value greater than 512.
1034591	Changed DMP get path state code to check only for write protection for Symmetrix.
1035027	Fixed a memory leak in the VxVM and DMP plugin of VxMS.
1046926	The same DMP node was used for two different disks when the LUN number of a DMP node was deallocated and reassigned to another physical device.
1050994	Using the <code>vxassist grow</code> command on a striped-mirror, log volume failed with the "Cannot update volume" error message.
1053089	Enhanced the <code>vxrecover</code> command to handle multiple recoveries.
1057146	DCO maps were cleared due to a failed attempt to attach a plex, which led to data corruption.
1065478	Fixed a panic in DMP due to stale DMP nodes.
1067870	A DCO volume moved to another disk when the volume was grown, due to <code>vxassist grow</code> not honoring the allocation specification.

Table 1-4 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1078615	Added support for initialization of LUNs greater than 2 TB.
1084199	The Oracle recovery process would loop when trying to read the next mirror.
1086982	Fixed a false serial split brain detection during DR.
1087183	The <code>vxconfigd</code> daemon core dumped on the primary during <code>flashsnap</code> .
1092390	Fixed a memory leak in <code>vxconfigd</code> .
1097834	Added a failover mode for Clarion arrays.
1104933	Fixed an incorrect calculation of disk capacity done by the <code>vxdisksetup</code> command.
1107078	The <code>vxconfigd</code> daemon core dumped at startup in CBR.
1107156	Added instantaneous device suppression support in DMP.
1120131	Fixed a system panic in DCO detach during a link failover.
1150763	Fixed an issue in which running <code>vxdisk updateudid</code> on an imported diskgroup disk could render the diskgroup unable to be imported.
1152033	The <code>vxconfigd</code> daemon core dumped during a BCV diskgroup import.
1176580	Fixed a CVM node panic that occurred when cables were unplugged from an array.
1182475	Fixed the <code>vxdg split</code> failing if the CVM master changes.
1183331	The <code>vxconfigrestore -p</code> command failed with a syntax error.
1192105	Fixed the <code>vxdg -n [newdg] deport [origdg]</code> command causing a memory leak.
1194738	Fixed a panic in <code>vxio</code> due to an incorrect calculation of number of DCOs.
1195666	After removing SDD, the <code>vxconvert</code> utility failed in the analysis phase with the following error message: <code>vxicap pvptr = NULL before reading VGRA</code>

Table 1-4 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1201755	The <code>vxrelocd</code> daemon failed to work during a simulated disk failure when using a layered volume.
1204200	The <code>vx dg join</code> command failed due to a transaction abort.
1204961	The <code>vxresize</code> command failed to resize with the following error message: "Cannot determine file system current size for volume."
1204981	Fixed an error that occurred with the <code>vx disk scandisk</code> command for <code>/dev/cciss/*</code> devices.
1205065	Fixed an incorrect error message for disabling the root disk or a failed disk.
1210160	Different <code>iotime</code> computation units used by <code>vxio</code> and <code>vx dmp</code> resulted in a false insane DMP node.
1213564	Fixed a panic that occurred during a DCO update.
1214356	The <code>vxconfigd</code> daemon transactions failed after a minor number conflict and temporary remapping of shared volumes.
1224737	Deported diskgroup disks from Clariion A/PF LUNs that were not on the preferred path were not accessible.
1225953	Improved I/O performance with VxFS filesystems on mirror-concat VxVM volumes with DCO and DRL.
1228536	Running <code>vx dg flush</code> on a slave node in a CVM cluster disabled the disk.
1230360	Fixed a system panic when one of the mirrored volume's plexes went down due to LUN access failure.
1232013	The <code>vxconfigd</code> daemon core dumped while importing a diskgroup.
1245037	The "Operations requires transaction" error occurred during snapshot creation.
1248303	Fixed a panic in <code>vx dmp</code> that was caused by an invalid CPU table address.
1253830	Fixed an issue with HP CCISS raid devices failing in <code>vxmconvert</code> .
1264340	The <code>vx dmpadm getattr arrayname name recoveryoption</code> command core dumped when more than one enclosure of a given type was attached.

Table 1-4 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1265972	After a crash or reboot, a node failed to re-join the CVM cluster.
1266730	Fixed the <code>vxtask</code> command to display the resync progress subtask for shared volumes with DRL.
1267134	The <code>vxconfigd</code> daemon core dumped due to stack corruption.
1280387	Fixed a panic that occurred when using the <code>vxsnap admir</code> command to take a link-breakoff snapshot.
1288418	The current primary path was not getting set on a DMP node in a CVM environment when all of the ports on a switch were disabled from the host side, and then the ports were enabled.
1293922	Enhanced VxVM to include an online plex comparison utility.
1299156	Fixed by increasing the time out value for the STPG command.
1299252	VxVM commands hung because <code>vxconfigd</code> was waiting for a configuration update to finish.
1321195	Fixed a panic on CVM nodes.
1317361	Fixed an issue with the <code>vxreattach</code> and <code>vxunreloc</code> commands were leaving behind tmp files.
1372340	Fixed an issue with <code>vxplex</code> core dumps during <code>vxassist addlog</code> due to DRL log length being less than 33 blocks.
1376656	Fixed an issue with <code>vxcached</code> never deletes old snaps when cache hits HWM.
1381783	Improved the performance of snapshot backups.
1382705	Fixed an issue with the <code>vxdmpadm listexclude</code> command dumping core when you run the VRTSexplorer tool.
1385922	Fixed a system panic due to memory allocation.
1387033	Fixed a system panic in <code>bcopy()</code> due to null passed in from <code>volioctl_copyin()</code>
1389512	Able to force import diskgroup version 80 in VxVM 5.0.
1389584	Fixed a system panic in <code>vol_putdisk()</code> code.

Table 1-4 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1396427	Enhanced DMP to handle failing IO when it is not able to interpret sense data.
1397712	Fixed an issue with the <code>vxsnap</code> restore manual page is unable to properly freeze or thaw filesystems in a CVM environment.
1397879	Enhanced the <code>vxresize</code> manual page to run from non-CVM master.
1403123	Fixed an issue with <code>vxconfigd</code> sleeping and no <code>vx</code> commands were responding.
1403370	Fixed a system panic after running the <code>vxctl enable</code> or <code>vxconfigd -k</code> commands.
1406185	Fixed a segmentation fault on x64 system when running the <code>vxmpadm list dmpnode all</code> command.
1411636	Fixed a secondary log error causing rlink disconnect after IBC unfreeze.
1412785	Fixed an issue with the system hanging while creating volumes in a guest Ldom.
1414342	Fixed an issue with disk devices not appearing in <code>vxdisk</code> list.
1414380	Fixed an issue with VVR Rlink failing to connect due to failure of memory allocation for incoming message buffer.
1414381	Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery.
1414441	The <code>vxsnap</code> manual page includes <code>mirror=enclosure</code> parameter to avoid being mirrored on the same enclosure.
1425434	Fixed an issue with CVR fails to connect rlinks followed by <code>vxconfigd</code> hangs on secondary.
1425919	Fixed an issue with <code>vxesd</code> looping using 100% of one CPU.
1427267	Fixed a CVR panic in <code>VOLSIOQ_MORE</code> ue to corrupted <code>volsioq_start</code> queue.
1427284	Fixed an issue with <code>vxmpadm</code> dumped core when executing <code>vxmpadm list dmpnode</code> command.
1428106	Fixed a system panic in <code>vxio:voldrl_trans_copy</code> .

Table 1-4 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1431279	Fixed an issue with <code>vxconfigd</code> core dumps.
1435471	Fixed an issue with the cluster nodes panicking in <code>voldco_or_pvmbuf_to_pvmbuf</code> code after installing 5.0 MP3.
1435882	Fixed a panic in <code>volkio_to_kio_copy</code> due to passing null ptr to <code>bcopy</code> .
1441020	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1441072	Fixed an issue with <code>siteread</code> policy is not honoured.
1441131	Fixed an issue with VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1442369	Fixed an issue with a bug in <code>vxconfigbackupd</code> script leading to 0 byte <code>binconfig</code> file being created.
1443706	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1443752	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
1470734	Fixed an issue with <code>vxconfigd</code> dumps core if all paths are excluded.
1475580	Fixed an issue with VTOC getting corrupted by our header "PRIVHEAD".
1487888	Fixed an issue with the start up process fails after upgrade to 5.0 MP3. 1487584 Fixed an issue with CVR: I/O hang on logowner after logclient crashes.
1487907	Fixed an issue with <code>vxio</code> has a non-unique UDID for failed disks.
1503242	Fixed an issue with <code>vxddmpadm iostat</code> stats not tallying with those of <code>sar</code> , <code>iostat</code> and <code>vxstat</code> .
1510204	Fixed an issue with <code>vxconfigd</code> command hanging after cluster nodes split simulation.
1502120	Fixed an issue with <code>/dev/dmptconfig</code> device causing SELINUX an audit error.

Veritas Volume Manager 5.0 MP1 and 5.0 MP2 fixed issues

Table 1-5 describes fixed issues in the Veritas Volume Manager 5.0 MP1 and 5.0 MP2 releases.

Table 1-5 Veritas Veritas Volume Manager fixed issues

Incident	Description
528677	Volume relayout is now supported for site-confined volumes and for site-consistent volumes.
540351	Reattaching a site when the disks were in the serial-split brain condition gave an error message.
540523	Under some circumstances, DMP nodes could be incorrectly enabled.
563524	Split, join and move operations failed on a source disk group that had any site-confined volumes.
584200	The vxmake command could not be used to recreate site records. This is now supported if the -d option is used to read from a description file.
601274	In a CVM cluster, DMP did not fail over to a secondary path when the primary paths were disconnected.
605743	If a disk group were split from a source disk group, volumes in the split-off disk group did not retain their volume tags.
609199	When the vxdmpadm disable command was applied to a primary path on one node in a CVM cluster, the other nodes did not fail over to the secondary path.
611333, 622508	DMP could not obtain the correct serial number for a device if its LUN serial number contained a comma. This problem was seen on EMC Symmetrix arrays with more than 8096 LUNs.
614061, 614787	Adding cache volumes (used by space-optimized instant snapshots) to volume sets could cause data corruption and system panics.
617331, 631334	I/O was not restored on a path that was re-enabled after a failback or a non-disruptive upgrade (NDU) operation.
618068	A system panic could occur when EMC PowerPath was configured to coexist with DMP as a third-party driver.
618317	A system crash could occur while bringing up cluster if I/O were performed on a unopened path.

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

Incident	Description
619958	<p>After encapsulating a SAN-connected root disk on an EMC CLARiiON CX700 array, a system panic occurred during booting when multiple paths were enabled.</p> <p>Systems with an encapsulated root disk on an Active/Passive array in explicit failover mode (A/PF array) failed to boot if more than one path was enabled.</p> <p>Systems with an encapsulated root disk on an Active/Passive array (A/P array) took a longer time than usual to boot.</p>
621832	Immediately after installation, the vxesd daemon had the DVD mount point as its current working directory, which prevented the DVD from being unmounted.
625877	The error "/etc/vx/vxvm-startup: line 241: /usr/sbin/vxdddadm: No such file or directory" was seen at boot time.
643089	Relayout from mirror-stripe to concat-mirror did not work for site-consistent volumes.
645749	Growing a volume by a specified amount did not work for a site-consistent volume with more than 2 disks per site.
771691	The vxrootmir command would not copy the Master Boot Record (MBR) to a root mirror disk. This behavior was generic to all arrays.
778352	Mirroring the root disk to a SAN boot disk could cause data corruption. This behavior was generic to all arrays.
793159	Automatic reattachment of a remote site did not work correctly.
801445	The DMP feature to detect and respond to intermittently failing paths was turned off by default in the 5.0 release, and the values of the dmp_health_time and dmp_path_age tunables were both set to 0. This feature is now enabled by default in 5.0 MP1. The default values of dmp_health_time and dmp_path_age are 60 and 300 seconds respectively.

Veritas Enterprise Administrator 5.0 MP3 fixed issues

There are no additional Veritas Enterprise Administrator fixed issues in the Veritas Storage Foundation 5.0 MP3 release.

Veritas Enterprise Administrator 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-6](#) describes fixed issues in the Veritas Enterprise Administrator 5.0 MP1 and 5.0 MP2 releases.

Table 1-6 Veritas Enterprise Administrator fixed issues

Incident	Description
578688	The maximum size of the Alert and Task logs has been documented as 2MB.
596284	An Action pull-down menu item did not exist for the Layout View, the Disk View or the Volume View.
599060	Controller states were reported as "Not Healthy" when they are actually healthy, and "Healthy" when they were actually not healthy.
614761	The volume set creation wizard showed cache volumes in the "Available Volumes" list.
616661	When connecting to the central host, an "OutOfBoundException" error could occur.
618146	A Java exception error occurred in the Statistics View.

Veritas Web GUI 5.0 MP3 fixed issues

There are no additional Veritas Web GUI fixed issues in the 5.0 MP3 release.

Veritas Web GUI 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-7](#) describes fixed issues in the Veritas Web GUI 5.0 MP1 and 5.0 MP2 releases.

Table 1-7 Veritas Web GUI fixed issues

Incident	Description
564455	Removing a volume from a volume set returned a Java exception.
565072	Creating a file system on a disabled volume returned both success and failure messages.
566619	The Scan Disks By Controller View did not list the available controllers.
574410	Attempting to create a volume without an existing disk group produced a misleading error.

Table 1-7 Veritas Web GUI fixed issues (*continued*)

Incident	Description
575262	Disabling a path to a SENA storage array produced an erroneous message.
576794	Ghost entries for disconnected disks in the All Disks View could not be removed by using the GUI.
596648	Messages about failures to import disk groups were not displayed by the Web GUI.
601157	The wizard could report that an ISP volume was created successfully when the command log showed that it was not.
605468	Forcibly removing a volume from a volume set displayed an erroneous message.
607026	At least one object had to be selected in the GUI before a disk could be initialized.
608573	Deleting a volume that had just been deleted produced a Java exception.
611894	Removing a disk from a disk group displayed an erroneous message.
615395	Attempting to delete an active cache volume failed with an error message that was incomplete.
619039	Messages about exceeding the Storage Foundation Basic soft limitations were not displayed by the Web GUI.
639751	Help for the Scan Disks by Controller page was missing.

Veritas File System fixed issues

The following sections list Veritas File System fixed issues for Storage Foundation 5.0 Release Update 1 release. Fixed issues in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/283859>

Veritas File System 5.0 RU1 fixed issues

Table 1-8 describes fixed issues in the Veritas File System 5.0 RU1 release.

Table 1-8 Veritas File System fixed issues

Incident	Description
1468375	<code>fsadm</code> is unable to shrink a file system, sometimes even when the file system is empty.
1653052	During a <code>vxfsconvert</code> , a memory fault causes a core dump.

Veritas File System 5.0 MP3 fixed issues

[Table 1-9](#) describes fixed issues in the Veritas File System 5.0 MP3, RP1, and RP2 releases.

Table 1-9 Veritas File System fixed issues

Incident	Description
841017	Fixed an erroneous EFAULT error that was returned for Linux's asynchronous I/O against VxFS files.
995917	Fixed an inode table overflow that occurred when upgrading disk layout Version 5 to Version 6 with the <code>vxupgrade</code> command.
1048438	Fixed a performance degradation in <code>vx_logbuf_clean()</code> .
1051524	Fixed an issue in which direct I/O to a disabled data volume in a volume set would cause a panic.
1087692	Fixed an issue in which the <code>vx_quota</code> command could fail with "permission denied" error.
1093636	Backed out changes to <code>vx_extfree()</code> introduced in incident 544066.
1099215	Fixed a possible infinite loop in <code>vx_write_alloc3()</code> due to a fragmented file system.
1143812	Removed the check for a Concurrent I/O license.
1167380	Fixed an incorrect use of the <code>fse_funmounted</code> flag that resulted in EIO errors.
1207239	Fixed a panic in <code>vx_getblk_cmn()</code> that could be caused when a Storage Checkpoint was mounted after the <code>ls -@</code> command was run against the primary file system.
1213253	The <code>fsadm -b</code> command sometimes failed to shrink a file system.

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

Incident	Description
1254405	Fixed a panic in <code>vx_dnlc_purge_ip()</code> that was caused by a dereference to a NULL <code>d_flhead</code> pointer.
1255517	The <code>df</code> command had different output depending on if a file system was local or part of a cluster.
1283708	Fixed an issue in which records specified with the <code>vxdump -B</code> command were incorrectly calculated if the value was larger than 2 gigabytes.

Veritas File System 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-10](#) describes fixed issues in the Veritas File System 5.0 MP1 and 5.0 MP2 releases.

Table 1-10 Veritas File System fixed issues

Incident	Description
616323	For WebGUI online help, the following issues have been fixed: For the Remount Storage Checkpoint operation, the More info link on the second wizard page does not function properly for cluster file systems. For the Unmount Storage Checkpoint operation, the More info link on the second wizard page does not function properly for cluster file systems.
770917	Inode ownership issues detected in large directory related code paths have been fixed.
770935	Prevented the system from panicking when setting access time (<code>atime</code>) or modification time (<code>mtime</code>) of named data streams by calling <code>vxfs_nattr_utimes()</code> API on 32-bit kernel.
770953	<code>fsck</code> used to create the lost+found directory with the <code>rwrxrwx</code> permissions if it decided to create one. Now, it creates the directory with the <code>rwxr-xr-x</code> permissions, which is consistent with the behavior of <code>mkfs</code> .

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

Incident	Description
771086	Fixed an <code>fsck</code> problem in which users could end up creating multiple lost+found directories when running the <code>fsck -o full</code> command and answering <code>fsck</code> questions interactively. Now, <code>fsck</code> creates only one. It also checks for multiple lost+found entries and removes duplicate directory entries.
771996	Enhanced VxFS to use less CPU when doing administrative tasks on the devices of multi-volume file systems.
772013	Enhanced the <code>fsck</code> command to enforce the lost+found file name in the root directory of the file system to be a directory file type.
777012	If the system crashed or there was a metadata I/O error, after the <code>fsadm</code> command reorganized the lost+found directory, running the <code>fsck -o full</code> command may not have been able to clean the file system with regard to names that needed to be added to the lost+found directory. The problem happened on single-volume and multi-volume file systems.
785649	A situation where <code>vxfsconvert</code> of a dusty file system loops forever in user-level code when an inode with pending truncation operation is encountered has been fixed.
793022	<code>vxfs_nattr_open()</code> API interface has been fixed to shrink files, as appropriate, when invoked with <code>O_TRUNC</code> flag.
793030	<code>vxfsutil.h</code> uses <code>struct timeval</code> in one of the function declarations, but does not include <code>time.h</code> . This causes user applications to report warnings during compilation. This issue has been fixed by including <code>time.h</code> in <code>vxfsutil.h</code> .
795073	The increased CPU utilization when writing to a file system that is almost full due to more background processing threads than are actually required being enqueued has been fixed.

Veritas Storage Foundation Cluster File System fixed issues

The following sections list fixed issues for Storage Foundation Cluster File System 5.0 Release Update 1 release. Fixed issues in the Veritas Storage Foundation Cluster File System 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/283857>

Veritas Storage Foundation Cluster File System 5.0 MP3 fixed issues

[Table 1-11](#) describes fixed issues in the Veritas Storage Foundation Cluster File System 5.0 MP3, RP1, and RP2 releases.

Table 1-11 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
1155353	Fixed an issue in which upgrading from disk layout Version 6 to Version 7 caused <code>fsck</code> to fail on the upgraded file system.
1167284	Fixed a live lock issue in <code>vx_iupdat_msg()</code> that resulted in the ENOTOWNER error.
1178621	Fixed an SFCFS multi-transaction server inode leak.
1314995	Fixed a leak of attribute inodes that caused the attribute ilist to become very large.

Veritas Storage Foundation Cluster File System 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-12](#) describes fixed issues in the Veritas Storage Foundation Cluster File System 5.0 MP1 and 5.0 MP2 releases.

Table 1-12 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
612406	Push the entire extent for an ilist hole. This will prevent multi-entry holes.
645227	The file system was hanging due to a rare race between the merge and activation of File Change Log (FCL) which lead to a deadlock. The problem has been resolved by ensuring that the merge threads try for the FCL global lock and proceed only if the lock is available.
770917	In some code paths related to large directory support, the cluster member (node) was not taking and holding the ownership of the inode for which the extents are to be re-organized. So, an assertion was being triggered when another member node of the cluster tried to revoke ownership of the inode. The problem has been resolved by correcting the ownership issues in large directory related code paths.

Table 1-12 Veritas Storage Foundation Cluster File System fixed issues
(continued)

Incident	Description
771892	A deadlock was occurring during CFS recovery. Log replay, which is performed as the first step during CFS recovery when a node crashes, was found to block in certain scenarios on an inode that the replay is trying to process. The problem has been resolved by freeing all inodes in the chunk one at a time without marking other inodes as busy.
771970	When a file is truncated, the truncate request is shipped to the cluster member that has the current ownership of the inode. There were certain inconsistencies in pre-splitting a buffer in the presence of clones, when the truncate request was shipped to another node of the cluster leading to an assertion. The problem has been resolved by changing the truncation code paths to handle the pre-splitting of buffers at a later stage during truncation, for both locally mounted and cluster mounted file systems.
771999	A node in a cluster was hanging under heavy CFS activity in the cluster. When all the worker threads servicing callbacks from a queue were blocked, the callbacks that were still lying in the same queue faced a deadlock. The problem has been resolved by queuing the callbacks that need to be processed to unblock other callbacks into a different queue, and processing them at a higher priority.
775611	Upon encountering an I/O error, the file system is correctly disabled but an internal inconsistency check fails. The check failure in this instance is quite benign and cannot cause customer data corruption since the file system is safely disabled upon encountering an I/O error of this nature. To further diagnose this condition, a new consistency check has been added.
782024	Under heavy inode cache stress, file system recovery on a cluster file system was hanging after an I/O error. The problem has been resolved by ensuring that the code paths that attempt to reuse an inode do not wait for a dirty/bad/inactive inode to become free if recovery initiated by a disabled file system is currently in progress. If there are no inodes available for re-use, a new inode will be allocated to ensure that the recovery proceeds to completion.
793030	The vxfsutil.h was using 'struct timeval' in one of the function declarations without including time.h resulting in a warning message during compilation. The problem has been resolved by including time.h in vxfsutil.h.

Veritas Volume Replicator fixed issues

The following sections list Veritas Volume Replicator fixed issues for Storage Foundation 5.0 Release Update 1 release. Fixed issues in the 5.0 release are listed in the *Veritas Volume Replicator 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/290244>

Veritas Volume Replicator 5.0 MP3 fixed issues

[Table 1-13](#) describes fixed issues in the Veritas Volume Replicator 5.0 MP3, RP1, and RP2 releases.

Table 1-13 Veritas Volume Replicator Fixed Issues

Incident	Description
1083609	VVR would have a hang on the primary with a heavy I/O load.
1085178	Fixed a replication hang due to outbound packets stalling.
1085219	Fixed a panic that occurred during the difference-based synchronization of an RVG.
1210105	"Operation requires transaction" errors occurred in VVR environments with a medium to large configuration size.
1281421	Fixed a panic that occurred during DCM protection due to a duplicate free.
1281987	Fixed an I/O hang that occurred after rlink was disconnected.
1315725	vxio had uncorrectable read/write errors with Oracle SmartSync.
1091734	Enhanced VVR to Allow <code>vol_rp_increment</code> and <code>vol_rp_decrement</code> to be tunables that can be modified from the <code>vxtune</code> interface.
840217	The <code>vradmin syncrvg</code> and the <code>vradmin syncvol</code> commands did not work correctly for volumes larger than 1TB.

Veritas Volume Replicator vradmin 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-14](#) describes fixed issues in the Veritas Volume Replicator vradmin 5.0 MP1 and 5.0 MP2 releases.

Table 1-14 Veritas Volume Replicator vradmin fixed issues

Incident	Description
641439	A security issue was discovered that could have resulted in a Low or Medium Severity attack against the VVR Administration service port, TCP/8199. The attacker would have needed to gain access to the network or gotten the user to visit a malicious site from which the attacker could initiate the attack. An attack could crash the vradmind service (which auto restarts in 60 seconds). Potentially, an attack could degrade system performance if the attack was sustained.
776831	Migration of the Primary could not be done after one node of Primary cluster panicked.
786185	Replicating from the Bunker to a Secondary generated startrep notification continuously, causing the GUI to hang.

Veritas Volume Replicator Web GUI 5.0 MP3 fixed issues

There are no additional Veritas Volume Replicator Web GUI fixed issues in the 5.0 MP3 release.

Veritas Volume Replicator Web GUI 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-15](#) describes fixed issues in the Veritas Veritas Volume Replicator Web GUI 5.0 MP1 and 5.0 MP2 releases.

Table 1-15 Veritas Volume Replicator Web GUI fixed issues

Incident	Description
576729	When a user without sufficient privileges tried to use the Create Primary wizard, the operation resulted in an incorrect error message.
611792	When a user tried to create a Primary without specifying the RVG name, the operation resulted in an RVG with an invalid configuration being created.
615758	If a CVM master node contains a private disk group, a Primary RVG created in that disk group was not displayed in the VVR Web GUI.
615769	When you created a Primary with the VVR Web GUI while connected to the master node of a CVM cluster, the Create Primary wizard did not display private disk group names in the disk group selection list.

Table 1-15 Veritas Volume Replicator Web GUI fixed issues (*continued*)

Incident	Description
615834	In a shared disk group environment, in some cases, clicking on the link for a Secondary disk group did not display the view of the disk group.
766453	If the list of disk groups is empty, clicking the Next button showed a blank screen.
768497	In the Create Primary wizard, sometimes non-free volumes were displayed in the volume list as free volumes.
770478	In some situations, the Current Logging field showed the value SRL when it should have been DCM.
776618	The Creating a Primary wizard failed when the list of data volumes was too long.
784039	In a VVR setup using a bunker Secondary with the STORAGE protocol, if the bunker disk group had been deported and imported several times, VRW displayed incorrect information about the RDS.
785051	The Deactivate Bunker operation was failing to find the activated bunker.
859597	A volume could be resized incorrectly and data could be lost if the requested size contained a decimal point.

Veritas Volume Replicator VEA 5.0 MP3 fixed issues

There are no additional Veritas Volume Replicator VEA fixed issues in the 5.0 MP3 release.

Veritas Volume Replicator VEA 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-16](#) describes fixed issues in the Veritas Volume Replicator VEA 5.0 MP1 and 5.0 MP2 releases.

Table 1-16 Veritas Volume Replicator VEA fixed issues

Incident	Description
602261	In some situations, the Current Logging field showed the value SRL when it should have been DCM.

Table 1-16 Veritas Volume Replicator VEA fixed issues (*continued*)

Incident	Description
616709	In the Japanese locale, the Add Bunker wizard page showed truncated text.
776622	The Creating a Primary wizard failed when the list of data volumes was too long.
784039	In a VVR setup using a bunker Secondary with the STORAGE protocol, if the bunker disk group had been deported and imported several times, VVR VEA displayed incorrect information about the RDS.

Known issues

The following are additional known issues for Veritas Storage Foundation.

Veritas Storage Foundation known issues

Known issues in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/283859>

Veritas Storage Foundation 5.0 RU1 known issues

Storage Foundation 5.0 RU 1 is based on Veritas Storage Foundation 5.0 MP3 RP2 and includes support for SuSE Linux Enterprise Server 11 (SLES 11). The following known issues are cumulative since the Veritas Storage Foundation 5.0 release. However, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation Cluster File System RAC are not available in Veritas Storage Foundation 5.0 RU1 because Oracle and DB2 have not yet provided support for databases on SuSE Linux Enterprise Server (SLES) 11. Review previous database-specific Veritas Storage Foundation 5.0 MP3 Release Notes known issues at:

<http://entsupport.symantec.com/docs/306947>

Additionally, there is no support for the Xen kernel on SLES11 in the Veritas Storage Foundation 5.0 RU1 release.

vxfsentsthdw needs raw device control interface for its fencing operations (1602658)

In 5.0 RU1, Veritas Storage Foundation High Availability, Storage Foundation Cluster File System, and Veritas Cluster Server `vxfsentsthdw` utility requires use

of a character device interface to write to the disks that it operates on, and this interface is provided by a raw utility. For the raw utility to work, "raw" kernel module should be loaded and configured. If this is not done before running the `vxfcntlsthdx` utility, the following message appears on the console during processing of the `vxfcntlsthdx` run:

```
Cannot open master raw device '/dev/rawctl' (No such file or
directory).
```

On SLES11, the raw driver is not configured automatically during system boot-up.

Workaround

Since the `vxfcntlsthdx` utility depends on the raw module for its operations, it is mandatory to configure it before running the `vxfcntlsthdx` utility. To configure the raw module, load the module after each system reboot using either of the the following commands:

```
# modprobe raw
or
# rcraw start.
```

After configuring the raw module, the `vxfcntlsthdx` script does not display any error message.

HBA discovery and corresponding object discovery failure on SLES 11 for SFM 2.0 managed host package installation (1666808)

When the SFM 2.0 managed host package `VRTSsfmh` is installed for SF SLES 11 (RU1), the LUN's Disks and HBAs tab is empty. HBA driver vendors have revised the model for populating HBA related attributes under `/sys`. This change causes the HBA discovery (and corresponding object discovery for LUNs, Enclosures, and HBAs) to fail. There is no workaround for this discovery issue, but updated components are available in the SFM 2.0 RP1 release.

For the latest information on updates, patches, and known issues regarding this release, see the following TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/281993>

Veritas Storage Foundation 5.0 MP3 known issues

The following are new additional known issues in this 5.0 MP3 release of Veritas Storage Foundation.

Clean up authentication data before starting LxRT5.0 MP3 installation or upgrade(1300779)

Upgrade to Storage Foundation 5.0 MP3, then install or upgrade authentication to version 4.3.34.4. If you have a previous version of authentication data, such as 5.0 authentication version 4.3.23.x already in your system, then install or upgrade does not update the `VRTSatLocal.conf` file. This results in the authentication installation or upgrade failing.

Workaround

Clean up authentication data before upgrading or installing 5.0 MP3. To cleanup the authentication data, enter:

```
# rm -rf /var/VRTSat
```

While configuring Storage Foundation Management Server and the Cluster Management Console through the CPI for authentication passwords, some special characters are not passed correctly through the CPI (1245237)

While configuring Storage Foundation Management Server and the Cluster Management Console through the Common Package Interface (CPI) for authentication passwords, some special characters are not correctly passed through the CPI to the nodes, even though these special characters are accepted by authentication.

The following special characters are not correctly passed through the CPI to the nodes:

- \ (single quote)
- \" (double quote)
- \@ (at sign)
- \\$ (dollar)
- \ (slash)
- * (star)

Workaround

There is no workaround for this issue. When entering authentication passwords, do not use any of the special characters listed above.

PATH setting when using product installer to install on remote nodes

There is a problem with some versions of `rsh` and `ssh` that causes the `PATH` to not get set correctly on remote shell invocations. As a result some of the commands fails to run. Before running the product installer, ensure that the `PATH` is set

correctly for the `root` user. For bash shells, it can be set in the `~/.bashrc` or the `~/.bash_profile` file.

To ensure that the PATH has been set properly

- ◆ Run the following command for each remote node you want to install and check that the command succeeds.

```
# rsh remote_node lsmod
```

Potential problems on systems with a QLogic HBA (935697, 989902, 1010069, 1017379, 1218589, 1594277)

If version 8.01.07 of the QLogic HBA driver is used, there is a potential for a NMI non-recoverable error and system panic if a third host-side path is disabled at the switch side. Other potential problems are for I/O to hang and for re-enabled paths to remain disabled. Issues such as these have been seen on RHEL4, SLES9-SP4 and SLES10 systems.

Workaround

Use version 8.01.06 of the HBA driver, which is available from the QLogic driver downloads site:

http://support.qlogic.com/support/drivers_software.asp

Another known issue (**1594277**) is specific to the 5.0 RU1 release for the SLES 11 operating system. Because the parameter `remove_on_dev_loss` of the `scsi_transport_fc` module is removed in SLES 11, the OS device files are removed after a device loss with `dev_loss_tmo`. When the device comes back online, the port names may have changed and DMP will then not recognize its status with the restored daemon. To prevent this occurrence, use the following workaround.

Workaround

Expand `dev_loss_tmo` to a very large value. This workaround only works with Qlogic driver.

For VxVM, modify the following two files:

Modify the file `/etc/udev/rules.d/40-rport.rules` to set the variable `KERNEL` as:

```
KERNEL=="rport-*",SUBSYSTEM=="fc_remote_ports",ACTION=="add",  
RUN+="/bin/sh -c 'echo 20 > /sys/class/fc_remote_ports/%k/fast_io_fail_tmo'"
```

Modify the file `/etc/modprobe.conf.local` to set the following values:

```
options qla2xxx qlport_down_retry=80000
```

Veritas Storage Foundation 5.0 MP2 known issues

There are no additional known issues in the 5.0 MP2 release of Veritas Storage Foundation.

Veritas Storage Foundation 5.0 MP1 known issues

The following are additional known issues in the 5.0 MP1 release of Veritas Storage Foundation.

LUNs are not displayed by the Storage Foundation Web GUI after upgrade (833516)

Because of a bug in the DDL provider package (`VRTSddlpr`) in the 5.0 release, the entries for the DDL provider are removed from the registry if the `rpm` command is used to upgrade to 5.0 MP1. This results in the Web GUI not being able to display any LUNs that are attached to the system.

Workaround

Specify the `--nopreun` option to the `rpm -U` command when upgrading the `VRTSddlpr` package.

The error does not occur if you use the `installmp` script to upgrade to 5.0 MP1.

Veritas Storage Foundation Basic known issues

Known issues in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/283859>

Veritas Storage Foundation Basic 5.0 MP3 known issues

The following are new additional known issues in this 5.0 MP3 release of Veritas Storage Foundation.

During VEA startup an error message may appear in `/var/log/audit/audit.log` file (1366109)

During VEA startup, with SELinux enabled and with `VRTSmapro` rpm installed. The following error message may appear in `/var/log/audit/audit.log` file.

```
type=AVC msg=audit(1216897889.477:1762): avc: denied { execmod }
for pid=9601 comm="vxpal" path="/opt/VRTSmapro/lib/vxmapping_prov.so"
dev=sdb1 ino=2810212 scontext=root:system_r:unconfined_t:s0-s0:c0.c1023
tcontext=system_u:object_r:lib_t:s0 tclass=file type=SYSCALL
```



```
msg=audit(1216897889.477:1762): arch=40000003 syscall=125
success=yes exit=0 a0=f0400000 a1=13000 a2=5 a3=ffa6c2f0 items=0
ppid=1 pid=9601 auid=0 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0 sgid=0
fsgid=0 tty=pts3 comm="vxpal" exe="/opt/VRTSobc/pal33/bin/vxpal"
subj=root:system_r:unconfined_t:s0-s0:c0.c1023 key=(null)
```

Workaround

Enable text relocation on shared library `vxmapping_prov.so`, enter:

```
# chcon -t textrel_shlib_t /opt/VRTSmapro/lib/vxmapping_prov.so
```

Veritas Volume Manager known issues

Known issues in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/283859>

Veritas Volume Manager 5.0 RU1 known issues

The following are new additional known issues in the 5.0 RU1 release of Veritas Volume Manager.

VxVM devices are not present at system boot time for VxFS fsck and mount (1411526)

On SuSE Linux Enterprise Server 10 and 11, VxVM devices are not present at system boot time for VxFS `fsck` and `mount`. When the system boots, it checks whether the file system needs a `fsck` before mounting, but it fails because the VxVM volume device is not present.

SLES11 does not finish device discovery (1635274)

SuSE Linux Enterprise Server 11 cannot complete device-discovery at boot-up, but Veritas Volume Manager indicates disks are fully discovered at startup.

Device Multipathing is restored after failure but undetectable by SLES11 (1594277)

When Device Multipathing (DMP) is restored after a failure, the SuSE Linux Enterprise Server 11 operating system cannot detect the re-enabled device status.

Volume Manager does not recognize IBM DS4k Series arrays in A/P-F mode (1595248)

Volume Manager does not recognize IBM DS4k Series arrays in A/P-F mode. This release supports DS4K array in A/P-C mode only; efforts to support the array in

A/P-F mode are in progress. For the latest information on updates, patches, and known issues regarding this release, see the following TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/281993>

Veritas Volume Manager 5.0 MP3 known issues

The following are new additional known issues in this 5.0 MP3 release of Veritas Storage Foundation.

vxddladm listsupport all command may not list all of the arrays (1320034)

The `vxddladm listsupport all` command may not list all of the arrays used by DMP.

Workaround

To list all of the supported arrays

- 1 Run the `vxddladm` command:

```
# vxddladm addsupport all
```

The `vxddladm addsupport all` command is informational and harmless to run.

- 2 Run the `vxddladm` command:

```
# vxddladm listsupport all
```

Messages about failback when no failover has occurred (1377822)

On reboot, messages about failback may display on the console even when no failover has occurred. For example:

```
VxVM vxdmp V-5-0-0 failback issued for 226/0x80
```

These messages are harmless and may be ignored.

Limitation of automatic site reattachment feature (1256764)

The site does not reattach automatically in the case where the site storage has disconnected and reconnected to a CVM slave node, but the master node never lost connection to the site storage.

vxvmconvert cannot convert LVM diskgroups to VxVM diskgroups (1362598)

Because of changes to the LVM implementation on the SLES9 SP4 OS platform, `vxvmconvert` cannot convert LVM diskgroups to VxVM diskgroups. The process will fail in the second stage of conversion with the following message:

```
VxVM  ERROR V-5-2-0 Unable to deactivate volume group
LVM_volume_group_name
Some volumes might be still in use.
```

There are no known workarounds for this issue.

Deport operation on a shared disk group fails (1368377)

With all primary paths inaccessible, the deport operation on a shared disk group fails to clear the PGR keys as the DMP database is not up-to-date. The deport operation succeeds but the PGR keys are not cleared as the DMP database is not updated to reflect the inaccessibility of failed primary paths.

Workaround

Running `vxdisk scandisks` before the DG deport operation triggers DMP reconfiguration which updates the DMP database such that a disk is accessible through active paths.

vxvmconvert fails for LVM volumes created on whole CCISS disks (1253830)

`vxvmconvert` on 5.0 MP3 only works for the LVM volumes created on CCISS device disk partitions, not whole CCISS disks. If a Sistine Physical Volume was created on a whole CCISS disk, LVM and LVM2 volumes created on that physical CCISS disk cannot be converted to VxVM volumes by `vxvmconvert`. The operation will fail and can result in the destruction of LVM partitions. For the Physical Volumes created on CCISS device disk partitions, `vxvmconvert` will succeed.

I/O failures result in the disk failing flag (1205380)

In some DMP failover scenarios, I/O retry causes the disk failing flag to be set, although there is nothing wrong with the disks except for the failing flag.

Workaround

Clear the failing flag using the `vxedit` command.

Manually installing the VRTSvxvm patch requires a reboot (1180992)

After you manually install the `VRTSvxvm` patch, you must reboot the system.

Plex disabled in DS4000 disk array (924680)

This issue has been identified in the 5.0 MP1 related to the DS4000 disk array. If there is a high I/O load to the array, a device inquiry may fail. This failure causes the DMP node to be disabled. When the DMP node is disabled, all I/O to the DMP node fails.

Autotagging can cause the reattach of a site to fail in a Campus Cluster (1470548)

Using the autotagging feature for a Campus Cluster can cause the site attach to fail. VxVM displays an error message such as the following:

```
VxVM vxdg ERROR V-5-1-10128 tagid already assigned to disk
```

Workaround:

Tag disks manually; do not use autotagging.

Veritas Volume Manager 5.0 MP2 known issues

There are no additional known issues in the 5.0 MP2 release of Veritas Storage Foundation.

Veritas Volume Manager 5.0 MP1 known issues

There are no additional known issues in the 5.0 MP1 release of Veritas Volume Manager.

Rootability issues of Veritas Volume Manager

There are no additional known issues in the 5.0 MP1 release of Veritas Volume Manager.

Utility issues of Veritas Volume Manager

The following are utility issues in the 5.0 MP1 release of Veritas Volume Manager.

NFS cannot handle minor numbers greater than 255

The NFS implementation in Linux does not support minor numbers greater than 255 unless a patch is applied. Without the patch, volume devices with large minor numbers cannot be remotely mounted via NFS.

This issue does not apply to Volume Manager 5.0 RU1 running SLES 11.

Workaround

Use the `vxldg` command to change the base minor number of the disk group that contains the volumes, as shown here:

```
# vxldg -g diskgroup remminor 2
```

See the description of Red HAT Bugzilla Bug 143897 or SUSE Bugzilla Bug 64552.

DMP issues of Veritas Volume Manager

The following are DMP issues in the 5.0 MP1 release of Veritas Volume Manager.

Identification of ATA and SATA disks (862137)

DMP is unable to identify ATA or SATA disks uniquely. This results in a single DMP virtual device being created for multiple ATA and SATA disks.

Workaround

Disable DMP for ATA and SATA disks.

To disable multipathing for ATA and SATA disks

- 1 Configure the device discovery layer to detect ATA disks as JBOD disks:

```
# vxddladm addjbod vid=ATA pid=*
```

- 2 Run the `vxdiskadm` command and select option 17 (Prevent multipathing/Suppress devices from VxVM's view).
- 3 Select option 7 (Prevent multipathing of disks by specifying a VID:PID combination).
- 4 Enter `ATA:*` as the VID:PID combination.
- 5 Exit from `vxdiskadm`, and reboot the system.

Handling intermittently failing paths

The `dmp_health_time` and `dmp_path_age` tunables control how DMP handles intermittently failing paths. The default values of `dmp_health_time` and `dmp_path_age` are 60 and 300 seconds respectively. The value of `dmp_health_time` represents the minimum time in seconds for which a path must stay healthy. If a path changes state between enabled and disabled on a shorter time scale than this, DMP marks the path as intermittently failing and disables I/O on the path. I/O is not re-enabled on an intermittently failing path until `dmp_path_age` seconds have elapsed without further outage.

The minimum configurable value of `dmp_path_age` is 0, which prevents DMP from detecting intermittently failing paths.

Cluster issues of Veritas Volume Manager

The following are cluster issues in this release of Veritas Volume Manager.

Handling intermittently failing paths in a Campus Cluster

In remote mirror configurations, a site is reattached when its disks come back online. Recovery is then initiated for the plexes of a volume that are configured at that site. Depending on the configuration, recovery of the plexes can take a considerable time and consume considerable resources. To minimize the frequency of having to perform a site reattachment operation, it is recommended that you use the `vxddmpadm settune` command to configure a value smaller than 60 seconds for `dmp_health_time`, and a value larger than 300 seconds for `dmp_path_age`.

Domain controller mode in CVM clusters (603213)

The slave nodes in a CVM cluster only have access to I/O objects. If non-I/O related information (for example, volume tags) are to be made available on a slave node, a command must be shipped to the Storage Agent on the master node for execution. The results are then communicated back to the slave node.

The domain controller mode of VEA allows all nodes of a CVM cluster to be placed in the same domain with a central authentication server. This allows commands to be executed on any node within the domain if the executing process has sufficient rights.

Provided domain controller mode is configured, non-I/O related information is accessible via VEA on any node in a CVM cluster.

However, even if domain controller mode is enabled in a CVM cluster, ISP commands must be run on the master node. ISP commands that are run on a slave node are not redirected to the Storage Agent on the master node. Such commands fail if they require access to non-I/O related information that is unavailable on a slave node.

Veritas Enterprise Administrator known issues

The following are VEA issues in this release of Veritas Volume Manager.

Veritas Web GUI known issues

The following are Veritas Web GUI issues in this release of Veritas Volume Manager.

Error when creating a volume set (615960)

An error such as the following may be seen when attempting to create a volume set that includes a newly created volume:

```
Error: 0xcfff0021 Facility: 0xffff Severity: 0x3 Error number: 0x21  
Object Not Found.
```

Workaround

Refresh the page.

Veritas File System known issues

The Veritas File System 5.0 Release Update 1 known issues are listed by release in this section. Known issues in the Veritas File System 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes* which is available at the following URL:

<http://entsupport.symantec.com/docs/283859>

Veritas File System 5.0 RU1 known issues

The following incidents are new additional Veritas File System known issues in the 5.0 RU1 release.

fsadmin cannot grow a file system that is already full (1468375)

Under some circumstances, the `fsadmin` command cannot grow a file system that is already 100 percent full.

Veritas File System 5.0 MP3 known issues

The following incident is a new additional Veritas File System known issue in the 5.0 RU1 release.

Veritas File System 5.0 MP3 known issues

The following are new additional Veritas File System known issues in this 5.0 MP3 release.

fsck may terminate when applied to unclean file systems from Veritas File System 4.1

Due to an incompatibility in the VxFS `fsck` utility between the 5.0 MP3 and 4.1 releases, `fsck` may terminate during intent log replay. This only affects file systems that were previously running under VxFS 4.1 and were not cleanly unmounted prior to use in VxFS 5.0 MP3. If you encounter this situation, perform a full `fsck`

to bring the file system to a consistent, clean state that is ready to be mounted. See the `fsck_vxfs(1M)` manual page.

Possible error messages using 32-bit commands on raw devices on 64-bit system (1079725)

On 64-bit systems, using 32-bit commands on raw devices sometimes produces error messages in the system log similar to the following example:

```
ioct132(mkfs.vxfs:7471): Unknown cmd fd(3) cmd(564f4c02){16}
arg(fff51138) on /dev/sdc2
```

You can safely ignore these messages.

Unmounting a VxFS file system takes long time after an iSCSI dynamic multipathing failover

Unmounting a VxFS file system can take several minutes after an iSCSI dynamic multipathing failover.

Workaround

The following procedure eliminates the long unmount time.

To eliminate the long unmount time

- 1 After a failover, run the `iostat` command to verify that all I/Os are drained:

```
# iostat
```

- 2 Unmount the file system or stop VCS:

```
# umount /MyFileSystem
```

Or:

```
# hastop -local
```

Fixes for some late-breaking issues

The fixes for some late-breaking issues, including a system panic when using block or character devices on a VxFS file system, and a potential security breach when using the `qiomkfile` and `odmmkfile` utilities, are available in a VxFS Hot Fix accompanying this release.

See the following TechNote for more information:

<http://entsupport.symantec.com/docs/306480>

fsapadm enforceckpt core dumps if you do not specify a Storage Checkpoint

The `fsapadm` command core dumps if you do not specify a Storage Checkpoint when specifying the `enforceckpt` keyword.

Workaround

The proper usage is as follows:

```
fsapadm enforceckpt [-f strict] mountpoint storage_checkpoint
```

Veritas File System 5.0 MP2 known issues

There are no additional known issues in 5.0 MP2 release of Veritas File System.

Veritas File System 5.0 MP1 known issues

The following are additional known issues in the 5.0 MP1 release of Veritas File System.

File Change Log tunable setting for proper functioning of Dynamic Storage Tiering applications

If the active placement policy of a given file system uses I/O or access temperatures, after the policy becomes active by being assigned, you must tune the file system's `fcl_malloc` tunable with the following command:

```
# vxtunefs -o fcl_maxalloc=0 mount_point
```

However, if any applications other than DST use FCL, this setting may conflict with those applications.

Veritas Storage Foundation Cluster File System known issues

Veritas Storage Foundation Cluster File System known issues in the 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*.

The *Veritas Storage Foundation Cluster File System 5.0 Release Notes* can be viewed at the following URL:

<http://entsupport.symantec.com/docs/283857>

Veritas Storage Foundation Cluster File System 5.0 RU1 known issues

The following are new additional Veritas Storage Foundation Cluster File System known issues in this 5.0 RU1 release.

Shared volumes do not come online after disconnecting and reconnecting the storage (1654539)

Shared volumes do not come online after disconnecting and reconnecting the Campus Cluster storage from SuSE Linux Enterprise Server 11. SLES11 removes the old device entry and original paths are not automatically restored by the restore daemon.

Workaround

Set the `dev_loss_tmo` timeout to **8000000**, and create the file `40-rport.rules` in `/etc/udev/rules.d/` so that device references cannot be removed. Prior to SLES 11, `remove_on_dev_loss` was enabled and could be tuned directly. But in SLES 11, `dev_loss_tmo` cannot be tuned directly and gets its value from the `qllogic` driver firmware, which, in turn, recognizes it as `qlport_down_retry`.

Disks not discovered by the HBA driver before VxVM starts (1603696)

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

During testing of SFHA and SFCFS with fencing enabled, the HBA driver (`qllogic qla2xxx`) did not complete discovery of the disks prior to VxVM startup. As a result, Volume Manager shows no disks present although they are visible in `lsscsi`.

Workaround

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

Veritas Storage Foundation Cluster File System 5.0 MP3 known issues

The following are new additional Veritas Storage Foundation Cluster File System known issues in this 5.0 MP3 release.

CFSMount resource may fault

During cluster startup in a SFRAC or SFCFS environment, a CFSMount resource may fault even though the underlying CVMVolDg resource becomes online successfully. If the CVMVolDg resource contains layered VxVM volumes, the reason for the fault could be that the CVMVolDg resource went online before all the subvolumes of the layered volume could be started.

Workaround

In order to ensure that a CVMVolDg resource containing layered volumes becomes online only after all the subvolumes are enabled, the CVMVolume attribute in the `main.cf` file should be populated with the names of the layered volumes under that CVMVolDg resource.

See the *Veritas Cluster Server User's Guide*.

Veritas Storage Foundation Cluster File System 5.0 MP2 known issues

There are no known issues in the 5.0 MP2 release of Veritas Storage Foundation Cluster File System.

Veritas Storage Foundation Cluster File System 5.0 MP1 known issues

There are no additional known issues in the 5.0 MP1 release of Veritas Storage Foundation Cluster File System.

Veritas Volume Replicator known issues

Known issues in the Veritas Volume Replicator 5.0 release are listed in the *Veritas Volume Replicator 5.0 Release Notes*, which is available at the following URL:

<http://entsupport.symantec.com/docs/283839>

Veritas Volume Replicator 5.0 MP3 known issues

The following are new additional known issues in this 5.0 MP3 release of Veritas Volume Replicator.

Mirrors are not synchronized when volumes created using `init=active` option (1287111)

For volumes created using `init=active` option, the mirrors (plexes) are not synchronized. The `vradmin verifydata` command could incorrectly report differences for such volumes. To rectify this situation, synchronize mirrors (plexes) and resynchronize the secondary by doing Automatic Synchronization, Full Synchronization, or Difference-based Synchronization.

VVR support with previous disk group version (1371981)

After upgrading to 5.0MP3 from a earlier version of VVR, the disk group version remains at the previous version unless you explicitly upgrade the disk group version. In a shared disk group environment with an older disk group version (110

or 120), the following issue may occur. If the master node on the primary CVM-VVR cluster fails, the RLINK may go into DETACHED and STALE state after the RVGLogowner service group fails over to the surviving node. Additionally, the `passthru` and `srl_header_err` flags are set on the primary RVG.

This issue may also occur in a private disk group environment after a reboot, when SRL recovery takes place.

Workaround

Upgrade the disk group version to the latest (140) using the following command:

```
# vxvg upgrade diskgroup
```

Veritas Volume Replicator 5.0 MP2 known issues

There are no additional Veritas Volume Replicator known issues in the 5.0 MP2 release.

Veritas Volume Replicator 5.0 MP1 known issues

The following are additional known issues in the 5.0 MP1 release of Veritas Volume Replicator.

Issue with VVR VEA in the Japanese locale (616709)

In the Japanese locale, the Add Bunker wizard page has truncated text. When you add a bunker using VVR VEA, the description text for the Bunker DG and Protocol fields is truncated.

The incomplete text should read as follows:

- Bunker DG: If protocol is Storage the Bunker DG is expected to have been imported on the Primary host.
- Protocol: Protocol should be set to Storage when Bunker storage is directly accessible from the Primary host.

Veritas Enterprise Administrators known issues

Veritas Enterprise Administrator includes the following additional known issues in 5.0 MP3:

Issue with SNMP settings through Tools (1209747)

This issue is specific to Rule Manager window in the VEA Client which is running on a Linux system. When `Default SNMP Settings` dialog is launched in Rule

Manager using menu `Tools->Default SNMP Settings`, the `Hostname` field does not allow any characters to be entered.

Workaround

To resolve this issue

- 1 Click on the `Port` field in `Default SNMP Settings` dialog.
- 2 Enter the port number in the `Port` field.
- 3 Click on `Hostname` field to enter hostname details.

Veritas Storage Foundation Graphical User Interface known issues

The following new additional known issues exist in this 5.0 MP3 release of Veritas Storage Foundation Graphical User Interface.

Cannot convert mountpoint to volume set using the Java Graphical User Interface (1176531)

When you use the Java Graphical User interface to convert a mountpoint into a volume set, the operation is not successful.

Workaround

You must use the `dbdst_convert` command line interface to convert a mountpoint volume to a volume set. You must not use the Java Graphical User Interface to convert volumes.

Adding two storage classes consecutively through VEA Java Graphical User Interface fails (1231856)

When you use the VEA Java Graphical User Interface to add the first class, the operation is successful. However, when you use the VEA Java Graphical User Interface to add a second class, you may see the following error:

```
SFORA dbdst_admin V-81-6212
```

```
Do not add or remove class in a single command.
```

This issue is encountered only with Graphical User Interface and not with the command line interface.

Workaround

You must logout and disconnect from VEA. Then you must start a new session to add a second class through the Graphical User Interface.

Alternatively, you may use the `dbdst_admin` command in the command line interface to add a new class:

```
# /opt/VRTS/bin/dbdst_admin -D DB4 -o addclass=NEWCLASS:"newclass"
```

No longer supported

This section describes Veritas Storage Foundation features that will not be supported in future releases. Symantec advises customers to minimize the use of these features.

The following features will not be supported in the next release of Veritas Storage Foundation products:

- The `seconly` feature in SFCFS
- Storage Expert
- The use of the `vxvoladm` command line utility
- Intelligent Storage Provisioning (ISP)

Documentation

Veritas Storage Foundation and High Availability Solutions 5.0 Release Update 1 adds support for SuSE Linux Enterprise Server 11. The Getting Started Guide and the Release Notes have been updated for this release. Refer to the Veritas Storage Foundation 5.0 MP3 documentation set for all other information.

The 5.0 MP3 documentation is available at the following URL:

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

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 5.0 Release Update 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 document is the relevant component product release notes:

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

Veritas Storage Foundation guides

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

Table 1-17 describes the guides in the Veritas Storage Foundation documentation set.

Table 1-17 Guides in Veritas Storage Foundation documentation set

Guide Title	Filename
<i>Third-party Legal Notices</i>	3rdpartyattributions.pdf
<i>Veritas FlashSnap Point-In-Time Copy Solutions Administrator's Guide</i>	flashsnap_admin.pdf
<i>Veritas Storage Foundation and High Availability Getting Started Guide</i>	getting_started.pdf
<i>Read me end user license agreement documentation</i>	README_EULA
<i>Veritas Storage Foundation read me first</i>	readme_first.txt
<i>Veritas Storage Foundation Release Notes (this document)</i>	sf_notes.pdf
<i>Veritas Storage Foundation Installation Guide</i>	sf_install.pdf
<i>Veritas Storage Foundation for DB2 Administrator's Guide</i>	sf_db2_admin.pdf
<i>Veritas Storage Foundation for Oracle Administrator's Guide</i>	sf_ora_admin.pdf
<i>Veritas Storage Foundation for DB2 Graphical User Interface Guide</i>	sf_db2_gui.pdf
<i>Veritas Storage Foundation for Oracle Graphical User Interface Guide</i>	sf_ora_gui.pdf
<i>Veritas Storage Foundation Intelligent Storage Provisioning Administrator's Guide</i>	sf_isp_admin.pdf
<i>Veritas Storage Foundation Intelligent Storage Provisioning Solutions Guide</i>	sf_isp_solutions.pdf

Table 1-17 Guides in Veritas Storage Foundation documentation set (*continued*)

Guide Title	Filename
<i>Veritas Storage Foundation Cross-Platform Data Sharing Administrator's Guide</i>	sf_cds_admin.pdf
<i>Veritas Enterprise Administrator User's Guide</i>	vea_users.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 1-18](#) describes the Veritas Storage Foundation Cluster File System (CFS) documentation set.

Table 1-18 Guides in Veritas Storage Foundation Cluster File System documentation set

Guide Title	Filename
<i>Veritas Storage Foundation Cluster File System Administrator's Guide</i>	sfcfs_admin.pdf

Veritas Cluster Server documentation

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

[Table 1-19](#) describes the Veritas Cluster Server documentation set.

Table 1-19 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 User's Guide</i>	vcs_users.pdf

Table 1-19 Guides in Veritas Cluster Server documentation set (*continued*)

Guide Title	Filename
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf
<i>Veritas Cluster Server Bundled Agent's Reference Guide</i>	vcs_bundled_agents.pdf
<i>VCS Enterprise Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_install.pdf
<i>VCS Enterprise Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_install.pdf
<i>VCS Enterprise Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_install.pdf

Veritas Volume Replicator documentation

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

[Table 1-20](#) describes the Veritas Volume Replicator documentation set.

Table 1-20 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 Web Console Administrator's Guide</i>	vvr_web_admin.pdf
<i>Veritas Volume Replicator Advisor User's Guide</i>	vvr_advisor_users.pdf
<i>Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide</i>	vvr_agents_config.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
```

Note: Installing documentation and online manual pages is optional.

Documentation errata

The following describes documentation errata.

Web GUI help errata

The Web GUI help is updated in this Maintenance Pack to include corrections for several help screens.

The Web GUI help is updated in the 5.0 RU1 release to include corrections for several help screens.

Manual pages errata

Several manual pages have been updated in the 5.0 RU1 release to include corrections for several errors or omissions.