

Veritas Storage Foundation™ Release Notes

Linux

5.0 Maintenance Pack 3



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 MP3

Document version: 5.0MP3.1

Legal Notice

Copyright © 2008 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.

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the *Veritas Storage Foundation 5.0 Release Notes*.

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

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

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
20330 Stevens Creek Blvd.
Cupertino, CA 95014

<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:

<http://www.symantec.com/techsupp/>

Contacting Technical Support

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

<http://www.symantec.com/business/support/index.jsp>

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

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

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

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

Licensing and registration

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

<http://www.symantec.com/techsupp/>

Customer service

Customer service information is available at the following URL:

<http://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	contractsadmin@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportolutions@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:

<http://www.symantec.com>

Select your country or language from the site index.

Contents

Technical Support	4
Chapter 1 Storage Foundation Release Notes	11
Overview of this release	11
About the Simple Admin utility	12
Changes in Storage Foundation	13
Installation and upgrade	13
Storage Foundation	14
Veritas Volume Manager	15
Veritas File System	21
Veritas Storage Foundation for Oracle	23
Veritas Storage Foundation for DB2	23
Veritas Storage Foundation Cluster File System	24
Required patches	24
Storage Foundation for Databases supported features	24
Supported features for Linux	25
System requirements	26
Software and hardware requirements	26
Supported Linux operating systems	26
Persistent network interface names on SUSE clusters	28
Storage Foundation Cluster File System node requirements	28
Storage Foundation memory requirements	28
Storage Foundation Cluster File System memory requirements	28
Storage Foundation supported DB2 versions	28
Storage Foundation supported Oracle versions	28
Storage Foundation Cluster File System supported Oracle versions	29
Mandatory patch required for Oracle Bug 4130116	29
VxVM licenses	29
Cross-Platform Data Sharing licensing	30
Component product release notes	30
Software limitations	30
Veritas Storage Foundation software limitations	30
Veritas Volume Manager software limitations	32
Veritas File System software limitations	32

Veritas Storage Foundation Cluster File System software limitations	33
Veritas Storage Foundation Cluster File System for Oracle RAC software limitations	34
Veritas Storage Foundation for Databases software limitations	36
Veritas Volume Replicator software limitations	41
Fixed issues	44
Veritas Storage Foundation fixed issues	44
Veritas Volume Manager fixed issues	44
Veritas File System fixed issues	51
Veritas Storage Foundation Cluster File System fixed issues	53
Veritas Volume Replicator fixed issues	55
Veritas Storage Foundation for Databases fixed issues	59
Known issues	61
Veritas Storage Foundation known issues	62
Veritas Storage Foundation Basic known issues	65
Veritas Volume Manager known issues	65
Veritas File System known issues	71
Veritas Storage Foundation Cluster File System known issues	74
Veritas Storage Foundation Cluster File System for Oracle RAC known issues	78
Veritas Volume Replicator known issues	78
Veritas Storage Foundation for Oracle known issues	80
Veritas Storage Foundation for DB2 known issues	87
Veritas Enterprise Administrators known issues	89
Veritas Storage Foundation Graphical User Interface known issues	90
No longer supported	91
Documentation	91
Veritas documentation disc	91
Relevant component product release notes	92
Veritas Storage Foundation guides	92
Manual Pages	95
Documentation errata	96
Web GUI help errata	96
Manual pages errata	96
Veritas Storage Foundation for DB2 Administrator's Guide errata	97

Appendix A	Veritas Storage Foundation 5.0 MP3 for Xen	99
	Supported Platforms for Xen	99
	Supported Features	101
	Unsupported Features	101
	Packages Required for Dom0 Installation	101
	Installing the Veritas Software in Dom0	102
	Verifying Software Versions	103
	Using the Veritas Software under Xen	103
	Uninstalling the Veritas Software from Dom0	104

Storage Foundation Release Notes

This chapter includes the following topics:

- [Overview of this release](#)
- [Changes in Storage Foundation](#)
- [Required patches](#)
- [Storage Foundation for Databases supported features](#)
- [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 MP3 Maintenance Pack 3 (MP3) product line:

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

- Veritas Storage Foundation™ for Oracle (Standard, Enterprise, and HA Editions)
- Veritas Storage Foundation™ for DB2 (Standard, Enterprise, and HA Editions)
- Veritas™ Volume Manager (VxVM)
- Veritas™ File System (VxFS)
- Veritas™ Volume Replicator (VVR)
- Veritas Storage Foundation™ Cluster File System (SFCFS)
- Veritas Storage Foundation™ Cluster File System for Oracle RAC (SFCFS for Oracle RAC)

Veritas Storage Foundation™ QuickStart is not available in this release.

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*.

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/283161>

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&pvaid=203_1

Changes in Storage Foundation

This section describes the changes in Veritas Storage Foundation 5.0 MP3.

Installation and upgrade

Storage Foundation installation and upgrade includes the following changes in 5.0 MP3:

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 or upgrade. 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.

Simplified upgrade procedures

Upgrade procedures have been simplified and automated, based on usability testing.

Support for Red Hat Enterprise Linux 5 and SUSE Linux Enterprise Server 10

The Storage Foundation software stack has been enhanced in 5.0 MP3 to support Red Hat Enterprise Linux 5 and SUSE Linux Enterprise Server 10.

Support for Red Hat Enterprise Linux 5, Security-Enhanced Linux

The Storage Foundation software stack has been enhanced in 5.0 MP3 to operate normally with Security-Enhanced Linux (SE Linux) enabled and used in enforcing mode with the Red Hat Enterprise Linux 5 targeted policy. At the moment, only the targeted policy in the `unconfined_t` domain is supported. No support is given for `strict`, `mls`, or user-customized security policies, nor is there support present in Veritas File System (VxFS) for SE Linux file labels or context mounts.

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.

Storage Foundation

Storage Foundation includes the following changes in 5.0 MP3:

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](#)).

Quick I/O and Veritas ODM

Quick I/O and Veritas Extension for Oracle Disk Manager (ODM) are now enabled by default for Storage Foundation and Storage Foundation Cluster File System.

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 this release, the feature is turned off by default. To turn on and use the full functionality the SmartMove feature, see the following URL:

<http://entsupport.symantec.com/docs/306479>

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 MP3:

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.

Support iSCSI devices are only for RHEL 5 and SLES 10.

Support on Xen

Veritas Volume Manager is now supported for Xen architecture on Linux. However, not all the features of VxVM are supported and some constraints apply. Before

installing Veritas Storage Foundation on Xen, review the installation information in this document.

See “[Supported Platforms for Xen](#)” on page 99.

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.

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 `vxdmpadm 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 `vxddmpadm getddmpnode` 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 `vxddmpadm 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 `vxddmpadm 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 `vxddmpadm 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 `vxddmpadm 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 `vxddmpadm iostat` command now has a `groupby` clause to provide cumulative I/O statistics listing per `dmnode`, 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/vx/dmpevents.log`. For backward compatibility, `/etc/vx/dmpevents.log` is a soft link to `/var/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 `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 listenenclosure` 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 MP3:

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 `fsppadm` command now supports UID, GID, and TAG elements in the placement policy XML file.
- Improved `fsppadm` 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, `fsppmk`, 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 for Oracle

Veritas Storage Foundation for Oracle includes the following change in 5.0 MP3.

Support for new database

Veritas Storage Foundation for Oracle provides support for the Oracle™ 11g database in the 5.0 MP3 release.

The following features of Veritas Storage Foundation for Databases are supported in the Oracle 11g environment:

- Storage Checkpoints
- Storage mapping
- Database cloning (`clonedb`)
- Database Flashsnap
- Tiered storage for databases (DBDST)
- Quick I/O
- Veritas Extension for Oracle Disk Manager (ODM)

Veritas Storage Foundation for DB2

Veritas Storage Foundation for DB2 includes the following change in 5.0 MP3.

Support for new database

The following features of Veritas Storage Foundation for Databases are supported in the DB2 9.5 environment:

- Storage checkpoints
- Storage mapping
- Database cloning (`clonedb`)
- Database Flashsnap
- Tiered storage for databases (DBDST)
- Concurrent I/O

Veritas Storage Foundation Cluster File System

Veritas Storage Foundation Cluster File System includes the following changes in 5.0 MP3:

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

Quick I/O and Veritas ODM

Quick I/O and Veritas Extension for Oracle Disk Manager (ODM) are now enabled by default for Storage Foundation and Storage Foundation Cluster File System. Existing licenses will enable these features after 5.0 MP3 is installed or upgraded on the system. This will not require the addition of a new license.

Required patches

There are no required patches for this release.

Storage Foundation for Databases supported features

The following sections list the supported features for each platform in the 5.0 MP3 release of Storage Foundation for Databases.

Supported features for Linux

The following features of Storage Foundation for Databases are supported on the following Linux platforms in this release.

Storage Foundation for Oracle features

This release of Storage Foundation for Oracle is supported on the following Linux platform.

Table 1-1 Supported Linux platforms for Storage Foundation for Oracle

OS, Platform, Version	OEL4	OEL5	RHEL4	RHEL5	SLES9	SLES10
Oracle9i R2	X	X	Full Support	X	Full Support	X
Oracle10g R1	X	X	Full Support	X	Full Support	X
Oracle10g R2	Full Support	Full Support	Full Support	Full Support	Full Support	Full Support
Oracle11g	Full Support	Full Support	Full Support	Full Support	X	Full Support*

Note: For Oracle 11gR1, there is full support for Storage Foundation for Oracle on SLES 10 in this release.

Note: The table displays Linux platforms supported for Storage Foundation for Oracle at the time of this release.

For late-breaking details about supported Linux platforms for Storage Foundation for Oracle, refer to the support website:

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

Storage Foundation for DB2 features

This release of Storage Foundation for DB2 is supported on the following Linux platform.

Table 1-2 Supported Linux platforms for Storage Foundation for DB2

OS, Platform, Version	OEL4	OEL5	RHEL4	RHEL5	SLES9	SLES10
DB2 8.1	X	X	Full Support	Full Support	Full Support	Full Support
DB2 8.2	X	X	Full Support	Full Support	Full Support	Full Support
DB2 9.1	X	X	Full Support	Full Support	Full Support	Full Support
DB2 9.5	X	X	Full Support	Full Support	Full Support	Full Support

Note: The table displays Linux platforms supported for Storage Foundation for DB2 at the time of this release.

For late-breaking details about Linux platforms supported for Storage Foundation for DB2, refer to the support website:

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

System requirements

This section describes the system requirements for this release.

Software and hardware requirements

For information on hardware requirements, see the hardware compatibility list. The hardware compatibility list (HCL) is available at:

<http://entsupport.symantec.com/docs/283161>

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

Supported Linux operating systems

This section lists the supported Linux operating systems for this release of Veritas Storage Foundation products, including Veritas Storage Foundation Cluster File System.

Read the Technical Support TechNote for the latest information on updates, patches, and software issues regarding this release.

<http://www.symantec.com/techsupp/>

The Veritas 5.0 MP3 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 4 (RHEL 4) with Update 3 (2.6.9-34 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64).
- Red Hat Enterprise Linux 5 (RHEL 5) with Update 1 (2.6.18-53.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 9 (SLES 9) with SP3 (2.6.5-7.244 kernel or later) or SP4 (2.6.5-7.308 kernel or later) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 10 (SLES 10) with SP1 (2.6.16.46-0.12 kernel or later) or SP2 (2.6.16.60-0.21 kernel or later) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux 4.4 or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux 5.1 or later on AMD Opteron or Intel Xeon EM64T (x86_64)

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

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

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

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

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

The Veritas 5.0 MP3 release is also supported on the Xen platform for Linux, with some restrictions.

See the Appendix, Veritas Storage Foundation 5.0 MP3 for Xen, in the *Storage Foundation Release Notes* for more information.

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

Persistent network interface names on SUSE clusters

On SUSE systems, network interfaces can change their names following a reboot. To configure persistent network interface names, add a `PERSISTENT_NAME=ethX` entry to the `/etc/sysconfig/network/ifcfg-eth-id-mac_address` network interface configuration file for each interface on each node of the cluster, where `X` and `mac_address` correspond to the interface number and MAC address.

Storage Foundation Cluster File System node requirements

All nodes in a Cluster File System must have the same operating system version and update level.

Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Storage Foundation Cluster File System memory requirements

2 GB of memory is required.

Storage Foundation supported DB2 versions

DB2 9.1 and DB2 8.2, also referred to as DB2 8.1 with FixPak 7, or later is supported on the Linux operating systems listed above.

DB2 9.5 support for Concurrent I/O is provided in the 5.0MP3 release. There is no DB2 9.5 support for Quick I/O in this release.

You can also refer to the *5.0MP3 Veritas Storage Foundation Installation Guide* for detailed information on supported Linux platforms versions for Storage Foundation for DB2.

Storage Foundation supported Oracle versions

Oracle versions 9.2.0.6, 10g, 10gR2 and 11g are all supported on the Linux operating system.

For Oracle 11gR1, there is full support for Storage Foundation for Oracle on SLES 10 in this release.

You can also refer to the *5.0MP3 Veritas Storage Foundation Installation Guide* for detailed information on supported Linux platforms versions for Storage Foundation for Oracle.

Storage Foundation Cluster File System supported Oracle versions

Oracle versions 10g Release 2 and 11g Release 1 are supported for use with Storage Foundation Cluster File System for Oracle RAC.

Mandatory patch required for Oracle Bug 4130116

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

VxVM licenses

The following table shows the levels of licensing in Veritas Volume Manager and the features supported at each level.

[Table 1-3](#) describes the levels of licensing in Veritas Volume Manager and supported features.

Table 1-3 Levels of licensing in Veritas Volume Manager and supported features

VxVM License	Description of Supported Features
Full	Concatenation, spanning, rootability, volume resizing, multiple disk groups, co-existence with native volume manager, striping, mirroring, DRL logging for mirrors, striping plus mirroring, mirroring plus striping, RAID-5, RAID-5 logging, Smartsync, hot sparing, hot-relocation, online data migration, online relayout, volume snapshots, volume sets, Intelligent Storage Provisioning, FastResync with Instant Snapshots, Storage Expert, Device Discovery Layer (DDL), Dynamic Multipathing (DMP), and Veritas Enterprise Administrator (VEA).
Add-on Licenses	Features that augment the Full VxVM license such as clustering functionality (cluster-shareable disk groups and shared volumes) and Veritas Volume Replicator.

Note: You need a Full VxVM license to make effective use of add-on licenses to VxVM.

To see the license features that are enabled in VxVM

- ◆ Enter the following command:

```
# vxctl license
```

Cross-Platform Data Sharing licensing

The Cross-Platform Data Sharing (CDS) feature is also referred to as Portable Data Containers.

The ability to import a CDS disk group on a platform that is different from the platform on which the disk group was last imported is controlled by a CDS license. CDS licenses are included as part of the Veritas Storage Foundation license.

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 MP3 Veritas Storage Foundation software limitations

The following are additional software limitations in the 5.0 MP3 release of Veritas Storage Foundation.

Support for Red Hat Enterprise Linux 5, Security-Enhanced Linux (1133511)

The Storage Foundation software stack has been enhanced in 5.0 MP3 to operate normally with Security-Enhanced Linux (SE Linux) enabled and used in enforcing mode with the RedHat Enterprise Linux 5 targeted policy. At the moment, only the targeted policy in the `unconfined_t` domain is supported. No support is given for `strict,mls`, or user-customized security policies, nor is there support present in Veritas File System (VxFS) for SE Linux file labels or context mounts.

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.

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 could be run by any user.

5.0 MP2 Veritas Storage Foundation software limitations

The following are additional software limitations in the 5.0 MP2 release of Veritas Storage Foundation.

Security-enhanced Linux on Red Hat Enterprise Linux 4 and SUSE Linux Enterprise Server 9

On Red Hat Enterprise Linux 4 (RHEL4) and SUSE Linux Enterprise Server 9 (SLES9), security-enhanced Linux (SELinux) support is provided for evaluation purposes only. Security policy files are not currently available for the Veritas Storage Foundation product stack.

To disable SELinux at boot time on both RHEL4 and RHEL5

- 1 Set the kernel boot parameter `selinux` to 0 (`selinux=0`) and reboot the machine.
- 2 Assuming the system has been configured for booting from SELinux, edit the `/boot/grub/grub.conf` file for RHEL 4 or RHEL 5 or edit the `/boot/grub/menu.lst` file for SLES 9 or SLES 10 to include `selinux=0` on the kernel line.
- 3 Reboot the machine to ensure the setting takes effect.

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 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 MP3 Veritas Volume Manager software limitations

There are no new additional software limitations in this 5.0 MP3 release of Veritas Volume Manager.

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 MP3 Veritas File system software limitations

The following are new additional software limitations in this 5.0 MP3 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 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

Veritas Storage Foundation Cluster File System software limitations 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>

5.0 MP3 Veritas Storage Foundation Cluster File System software limitations

The following are new additional software limitations in this 5.0 MP3 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

The following are additional software limitations in the 5.0 MP2 release of Veritas Storage Foundation Cluster File System.

Security-enhanced Linux on Red Hat Enterprise Linux 4 and SUSE Linux Enterprise Server 9

On Red Hat Enterprise Linux 4 (RHEL4) and SUSE Linux Enterprise Server 9 (SLES9), security-enhanced Linux (SELinux) support is provided for evaluation purposes only. Security policy files are not currently available for the Veritas product stack.

To disable SELinux at boot time on both SLES9 and RHEL4

- 1 Set the kernel boot parameter `selinux=0` (`selinux=0`) and reboot the machine.
- 2 Assuming the system has been configured for booting from SELinux, edit the `/boot/grub/menu.lst` file to include `selinux=0` on the kernel line.
- 3 Then reboot the machine to ensure the setting takes effect.

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 Storage Foundation Cluster File System for Oracle RAC software limitations

The following are additional software limitations for Veritas Storage Foundation Cluster File System for Oracle RAC.

5.0 MP3 Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

The following are software limitations in the 5.0 MP3 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

I/O Fencing

Veritas Storage Foundation Cluster File System (SFCFS) for Oracle RAC is not licensed to support the I/O Fencing feature. If you use the product installer or installation scripts, you are not given the option during installation to enable I/O

Fencing, and the I/O Fencing driver will be disabled. If you install the product manually, use the `/etc/vxfen.d/vxfenmode_disabled` file as a template for the `/etc/vxfenmode` configuration file.

If you subsequently enable I/O Fencing, importing a shared disk group will fail with the following error message:

```
VxVM vxdg ERROR V-5-1-10978 Disk group shared_dg: import failed:  
License has expired or is not available for operation
```

To verify the state of the I/O Fencing driver, use the `/sbin/vxfenadm -d` command.

Storage Foundation Cluster File System for Oracle RAC does not support Symantec's implementation of SCSI-3 PGR based I/O fencing and Oracle Clusterware (CRS) is expected to handle any split-brain situations. More information is available at the following URL:

<http://entsupport.symantec.com/docs/306411>

Avoiding the serial split brain condition

Veritas Storage Foundation Cluster File System (SFCFS) for Oracle RAC does not include support for the I/O Fencing feature. As a result, SFCFS requires that two heartbeat links be configured during installation. When a node is reduced to a single heartbeat connection, SFCFS can no longer discriminate between the loss of a system and the loss of the final network connection. This "jeopardy" state affects all applications that use the cluster file system mount points. (Jeopardy cannot be avoided as the I/O Fencing feature is not available.)

If a node fails after the jeopardy state has been notified, all the cluster nodes cease to be members of the shared disk group, but the mounted cluster file systems are not disabled on all the nodes in the cluster. This is fine because Oracle's native fencing will evict the problem node from the cluster and prevent potential data corruption. The surviving cluster nodes will be unaffected and continue to operate.

Oracle's native fencing will evict the problem node from the cluster if either of the following events occur:

- Simultaneous failure of both heartbeat links to a node.
- A node hangs and is unable to respond to heartbeat messages.

5.0 MP2 Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

The following are software limitations in the 5.0 MP2 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

Storage Foundation Management Server is not supported for centralized servers

Storage Foundation Management Server is only supported for managed hosts; it is not supported for centralized servers.

5.0 MP1 Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

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

Veritas Storage Foundation for Databases software limitations

The following are Veritas Storage Foundation for Databases software limitations.

5.0 Veritas Storage Foundation for Databases 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 MP3 Veritas Storage Foundation for Databases software limitations

The following are new additional software limitations in this 5.0 MP3 release of Veritas Storage Foundation for Databases software limitations.

Security-enhanced Linux on Red Hat Enterprise Linux 5

Due to limitations with Oracle's and DB2's libraries, Storage Foundation cannot be used with security-enhanced Linux (SE Linux) enabled. Disable SE Linux in these environments.

5.0 MP2 Veritas Storage Foundation for Databases software limitations

The following are additional software limitations in the 5.0 MP2 release of Veritas Storage Foundation for Databases.

Deep mapping unsupported

Deep mapping on EMC SYMCLI is currently not supported.

Use of buffered I/O

VxFS provides sequential consistency among the read and write accesses to a file. That is, the results of these reads and writes appear as if they occurred in a serial order consistent with program order, and each access appears to be atomic. This is consistent with traditional UNIX file system semantics for reads and writes.

Other Linux file systems do not guarantee atomicity of reads and writes, which allows more efficient implementation, but also requires that applications use other mechanisms to achieve the same level of consistency if they require it.

VxFS file consistency can be relaxed in several ways. In case of a cluster mount, reads from and writes to a file are not considered conflicting unless they operate on an overlapping byte range. On a local mount, the user can request that the Concurrent I/O option be used on a specific file. This will mean that reads and writes are not guaranteed to be atomic, which can be desirable behavior for some applications.

5.0 MP1 Veritas Storage Foundation for Databases software limitations

The following are additional software limitations in the 5.0 MP1 release of Veritas Storage Foundation for Databases.

Some features stop working after a Global Cluster failover (563603)

Some Storage Foundation for Databases features do not work correctly after a Global Cluster (GCO) Failover. In 5.0, the Storage Foundation for Database (SFDB) repository and tools do not manage virtual hostnames correctly in a Global Cluster environment. The SFDB repository does not correctly adjust to the secondary host after the failover.

Features like Storage Checkpoint, Database FlashSnap, the scheduler, and Database Dynamic Storage Tiering (DBDST) will not function as normal after a failover. However, features such as Oracle Disk Manager (ODM), Quick I/O, and Concurrent I/O (CIO) will continue to work after a failover. This issue will be fixed after the next release.

Differing locales produces unintelligible characters in GUI (605487)

The GUI does not support database users having a different locale than the superuser's locale. The GUI will display unintelligible characters if the SFDB repository server starts with a locale that is different from the database user locale (client).

DBDST limitations with non-English filenames and placement class names (599164)

DBDST does not work on non-English database filenames or non-English placement class names, due to limitations in VxFS Dynamic Storage Tiering and VxVM volume tags. VxFS Dynamic Storage Tiering does not support placement of non-English filenames. The VxVM volume tag feature does not support non-English volume tag names.

Avoid using UNIX VEA via PC-based UNIX emulators

There can be problems displaying deep mapping topology in PC-based UNIX emulators like Exceed. Use the Windows VEA client instead of running the UNIX VEA client via emulators.

CLI database state changes are delayed in GUI (604685)

If you use the command line to start or stop the database, the state change is not immediately shown in the GUI. This delay can take up to 60 minutes.

Workaround

Start or stop the database from the GUI, or do a manual rescan from the GUI after starting or stopping with CLI.

Deep mapping unsupported

Deep mapping on EMC SYMCLI is currently not supported.

Command Line Interface limitations

This section describes command line interface limitations.

The following commands do not support multipartition databases on SMP architectures:

- `db2ed_clonedb`
- `db2ed_vmchecksnap`
- `db2ed_vmclonedb`

Database cloning limitation (285139)

If you clone a database and mount it, ensure that the directory where the mount point resides is owned by the instance owner of the cloned database.

If the directory where the mount point resides is not owned by the instance owner, an error message is displayed when you attempt to remove and unmount the cloned database.

For example:

```
$ db2ed_clonedb -I inst01 -S prod -T clone -c \  
Checkpoint_1105997700 -m /mnt
```

where /mnt is created by root and the owner has been changed to inst01.

When you attempt to remove and unmount the clone database, you will get the following error message:

```
$ db2ed_clonedb -T clone -o umount -d  
rm: Unable to remove directory /mnt Permission denied
```

This error message does not affect the functionality of db2ed_clonedb. The clone database has been removed and unmounted even when you receive this error message.

To avoid this error, create a directory under / as root and change the owner of the directory to the instance owner. Then, specify a mount point under the newly created directory. For example, instead of using the mount point /mnt as in the above example, specify a mount point under /mnt, such as /mnt/clone:

```
$ db2ed_clonedb -I inst01 -S prod -T clone -c \  
Checkpoint_1105997700 -m /mnt/clone
```

Note: If your mount point is under the directory /tmp, you will not encounter this problem.

No support for Intelligent Storage Provisioning

The Standard, Standard HA, Enterprise, and Enterprise HA versions of Veritas Storage Foundation for DB2 do not support Intelligent Storage Provisioning (ISP).

Using VxVM volumes as DB2 DEVICE containers

For this release, Symantec recommends that the VxVM raw device interface be used when configuring a raw device as a DB2 container. Using VxVM raw interface ensures fast and reliable I/O support. Also, VxVM's 5.0 raw device interface supports full Linux native asynchronous I/O.

Make sure the DB2 instance owner has read/write access to the VxVM volume. Also, the DB2 instance owner should have permissions to change the directory to /dev/vx. By default, only the root user can access the /dev/vx directory.

The following example shows how a DMS tablespace is created on a VxVM volume.

```
CREATE REGULAR TABLESPACE USERRAW \  
PAGE SIZE 4K \  
MANAGED BY DATABASE \  
ON VOLUME vxvol1
```

```
USING (DEVICE '/dev/vx/dsk/db2dg/db2vol1' 128000) \  
EXTENTSIZE 32 PREFETCHSIZE 32 \  
BUFFERPOOL IBMDEFAULTBP
```

An alternative method to use a VxVM volume is to go through the Linux raw binding interface. For further information, refer to the man page for raw commands on Linux. Using Linux raw binding is not a recommended method for VxVM.

Deep mapping on hard disk arrays

This release of Veritas Storage Foundation 5.0 for DB2 will not offer the deep mapping feature because partner disk array software on Linux platforms have been unavailable. The storage mapping display will show disk arrays only to the LUN level.

During installation, the `VRTSvail` package is not installed. However, this package is included in this release, so that this issue can be corrected in a future maintenance pack release.

Use of buffered I/O

VxFS provides sequential consistency among the read and write accesses to a file – that is, the results of these reads and writes appear as if they occurred in a serial order consistent with program order, and each access appears to be atomic. This is consistent with traditional UNIX file system semantics for reads and writes.

Other Linux file systems do not guarantee atomicity of reads and writes, which allows more efficient implementation, but also requires that applications use other mechanisms to achieve the same level of consistency if they require it.

VxFS file consistency can be relaxed in several ways. In case of a cluster mount, reads from and writes to a file are not considered conflicting unless they operate on an overlapping byte range. On a local mount, the user can request that the Concurrent I/O option be used on a specific file. This will mean that reads and writes are not guaranteed to be atomic, which can be desirable behavior for some applications.

Veritas recommends the use of Concurrent I/O with DB2. In DB2 UDB version 8.2.2 or above, Concurrent I/O is used with VxFS when DBAs choose "Direct I/O" for a given tablespace. When Direct I/O is specified this way, the DB2 server activates Concurrent I/O semantics in a way that is transparent to users.

See the *Veritas Storage Foundation for DB2 Administrator's Guide* to learn how to use Concurrent I/O with DB2 UDB on Linux, UNIX, and Windows.

DBDST class names limited to 29 characters (601746)

The `dbdst_admin -o rmclass` command fails when attempting to remove a class name of 30 characters or more. The maximum class name length is 29 characters.

Selected utilities require setuid (643964, 643964)

Some Veritas Storage Foundation for Databases programs are setuid binaries because they are meant to be run as a database administrator and the APIs used are root access-only Symantec internal APIs. The affected binaries are used mainly for information query purposes.

For these reasons, the following programs are setuid-enabled in Veritas Storage Foundation for Oracle:

- `/opt/VRTSdbed/.dba/dbed_analyzer`
- `/opt/VRTSdbed/.dba/vxckptplan`
- `/opt/VRTSdbcom/bin/vxstorage_stats`
- `/opt/VRTSdbcom/.dba/vxdbd_start`
- `/opt/VRTSdbcom/.dba/vxckpt_ismounted`

In Veritas Storage Foundation for DB2, the following programs are setuid-enabled:

- `/opt/VRTSdb2ed/.dba/vxdb2adm`
- `/opt/VRTSdbcom/bin/vxstorage_stats`
- `/opt/VRTSdbcom/.dba/vxdbd_start`
- `/opt/VRTSdbcom/.dba/vxckpt_ismounted`

Multiple archive log destinations with RAC (795617)

Multiple archive log locations are not supported in RAC configurations.

Repository hostnames are case insensitive (859863)

Because DNS host name lookup queries are by definition case insensitive, make sure the SFDB repository is running on a host with a name that is truly unique -- regardless of case -- within the local subnet. Errors may occur if the repository host name differs from another host name only by case.

Veritas Volume Replicator 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/283839>

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.

Replication not supported between Solaris and HP-UX for shared disk groups

Replication in a shared disk group is not supported between Solaris and HP-UX. (592349)

VRW limitations

VRW limitations

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

VCS Agents for VVR limitations

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.

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 release.

Veritas Storage Foundation fixed issues

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 MP3 Veritas Storage Foundation fixed issues

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

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 describe Veritas Volume Manager (VxVM) fixed issues.

Veritas Volume Manager 5.0 MP3 fixed issues

[Table 1-4](#) describes fixed issues in the Veritas Volume Manager 5.0 MP3 release.

Table 1-4 Veritas Volume Manager fixed issues

Incident	Description
592685	Fixed a <code>vxdmpadm getdmpnode help</code> usage error.
779060	Enhanced VxVM to perform faster failover/plex detach for mirrored volumes.
934142	Fixed a failure that occurred when importing a version 60 diskgroup.
971002	Fixed a failure that occurred when using the <code>vxdmpadm setattr</code> command to update the dynamic multipathing events.log.
973256	VxVM displayed the same dynamic multipathing node for two different LUNs.
990003	<code>vxctl 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>vxvoltage</code> command could not set <code>maxiosize</code> to a value greater than 512.
1053089	Enhanced the <code>vxrecover</code> command to handle multiple recoveries.
1084199	The Oracle recovery process would loop when trying to read the next mirror.
1087183	The <code>vxconfigd</code> daemon core dumped on the primary during flashsnap.
1092390	Fixed a memory leak in <code>vxconfigd</code> .
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.
1194738	Fixed a panic in <code>vxio</code> due to an incorrect calculation of number of DCOs.
1201755	The <code>vxrelocd</code> daemon failed to work during a simulated disk failure when using a layered volume.
1204200	The <code>vxdg join</code> command failed due to a transaction abort.

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

Incident	Description
1210160	Different <code>iotime</code> computation units used by <code>vxio</code> and <code>vxdmp</code> resulted in a false insane DMP node.
1213564	Fixed a panic that occurred during a DCO update.
1224737	Deported diskgroup disks from Clariion A/PF LUNs that were not on the preferred path were not accessible.
1228536	Running <code>vx dg flush</code> on a slave node in a CVM cluster disabled the disk.
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>vxdmp</code> that was caused by an invalid CPU table address.
1264340	The <code>vxdmpadm getattr arrayname name recoveryoption</code> command core dumped when more than one enclosure of a given type was attached.
1267134	The <code>vxconfigd</code> daemon core dumped due to stack corruption.
1280387	Fixed a panic that occurred when using the <code>vxsnap addmir</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.
1299252	VxVM commands hung because <code>vxconfigd</code> was waiting for a configuration update to finish.
1321195	Fixed a panic on CVM nodes.
1034591	Changed DMP get path state code to check only for write protection for Symmetrix.
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.

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

Incident	Description
1057146	DCO maps were cleared due to a failed attempt to attach a plex, which led to data corruption.
1067870	A DCO volume moved to another disk when the volume was grown, due to <code>vxassist grow</code> not honoring the allocation specification.
1078615	Added support for initialization of LUNs greater than 2 TB.
1086982	Fixed a false serial split brain detection during DR.
1097834	Added a failover mode for Clarion arrays.
1107078	The <code>vxconfigd</code> daemon core dumped at startup in CBR.
1107156	Added instantaneous device suppression support in DMP.
1183331	The <code>vxconfigrestore -p</code> command failed with a syntax error.
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>vxdisk scandisk</code> command for <code>/dev/cciss/*</code> devices.
1205065	Fixed an incorrect error message for disabling the root disk or a failed disk.
1214356	The <code>vxconfigd</code> daemon transactions failed after a minor number conflict and temporary remapping of shared volumes.
1230360	Fixed a system panic when one of the mirrored volume's plexes went down due to LUN access failure.
1265972	After a crash or reboot, a node failed to re-join the CVM cluster.
1035027	Fixed a memory leak in the VxVM and DMP plugin of VxMS.
1152127	Enforced cleaning up of obsolete external ASL-APMs after upgrading to VxVM 5.0.
1176580	Fixed a CVM node panic that occurred when cables were unplugged from an array.
1050994	Using the <code>vxassist grow</code> command on a striped-mirror, log volume failed with the "Cannot update volume" error message.

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

Incident	Description
1104933	Fixed an incorrect calculation of disk capacity done by the <code>vxdisksetup</code> command.
1065478	Fixed a panic in DMP due to stale DMP nodes.

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 <code>vxmake</code> command could not be used to recreate site records. This is now supported if the <code>-d</code> 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 <code>vxdmadm disable</code> 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.

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

Incident	Description
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.
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 <code>"/etc/vx/vxvm-startup: line 241: /usr/sbin/vxddladm: No such file or directory"</code> 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 <code>dmp_health_time</code> and <code>dmp_path_age</code> tunables were both set to 0. This feature is now enabled by default in 5.0 MP1. The default values of <code>dmp_health_time</code> and <code>dmp_path_age</code> are 60 and 300 seconds respectively.

Veritas Storage Foundation Enterprise Administrator 5.0 MP3 fixed issues

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

Veritas Storage Foundation 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.

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

Incident	Description
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.
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 describe Veritas File System (VxFS) fixed issues.

Veritas File System 5.0 MP3 fixed issues

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

Table 1-8 Veritas File System fixed issues

Incident	Description
823590	Fixed a deadlock involving POSIX threads when the threads memory mapped VxFS files.
841017	Fixed an erroneous EFAULT error that was returned for Linux's asynchronous I/O against VxFS files.
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.

Veritas File System 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-9](#) describes fixed issues in the Veritas File System 5.0 MP1 and 5.0 MP2 releases.

Table 1-9 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 (atime) or modification time (mtime) 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> .
770964	<code>fsck</code> has been enhanced to replay file systems created with earlier log versions on volume sets.

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

Incident	Description
771086	Fixed an fsck problem in which users could end up creating multiple lost+found directories when running the fsck -o full command and answering fsck questions interactively. Now, fsck 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 fsck 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 fsadm command reorganized the lost+found directory, running the fsck -o full 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 vxfsconvert of a dusty file system loops forever in user-level code when an inode with pending truncation operation is encountered has been fixed.
793022	vxfs_nattr_open() API interface has been fixed to shrink files, as appropriate, when invoked with O_TRUNC flag.
793030	vxfsutil.h uses struct timeval in one of the function declarations, but does not include time.h. This causes user applications to report warnings during compilation. This issue has been fixed by including time.h in vxfsutil.h.
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

Veritas Storage Foundation Cluster File System fixed 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 MP3 fixed issues

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

Table 1-10 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 running <code>fsck</code> on the failed file system to fail.

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

[Table 1-11](#) describes fixed issues in the Veritas Storage Foundation Cluster File System 5.0 MP1 and 5.0 MP2 releases.

Table 1-11 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.
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.

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

Incident	Description
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.
771998	Add CFS support for solaris 10 local zones as of the 5.0 MP1 release. VRTSvxfs and VRTSodm now support Clustered File Systems mounted within Solaris 10 local zones. This includes LOFS and direct mounts.
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 describe Veritas Volume Replicator (VVR) fixed issues.

Veritas Volume Replicator 5.0 MP3 fixed issues

Table 1-12 describes fixed issues in the Veritas Volume Replicator 5.0 MP3 release.

Table 1-12 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.

Veritas Volume Replicator vradmin 5.0 MP1 and 5.0 MP2 fixed issues

Table 1-13 describes fixed issues in the Veritas Volume Replicator vradmin 5.0 MP1 and 5.0 MP2 releases.

Table 1-13 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.

Table 1-13 Veritas Volume Replicator vradmin fixed issues (*continued*)

Incident	Description
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 1-14](#) describes fixed issues in the Veritas Veritas Volume Replicator Web GUI 5.0 MP1 and 5.0 MP2 releases.

Table 1-14 Veritas Volume Replicator Web GUI fixed issues

Incident	Description
516812	On HP-UX, uninstalling did not remove all VRW files and directories.
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.
612565	VRW sometimes displayed the sizes of the SRL and data volumes incorrectly for an RDS replicating between VVR 4.1 and VVR 5.0 on the HP-UX operating system.
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.
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.

Table 1-14 Veritas Volume Replicator Web GUI fixed issues (*continued*)

Incident	Description
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 1-15](#) describes fixed issues in the Veritas Volume Replicator VEA 5.0 MP1 and 5.0 MP2 releases.

Table 1-15 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.
612565	VVR VEA sometimes displayed the sizes of the SRL and data volumes incorrectly for an RDS replicating between VVR 4.1 and VVR 5.0 on the HP-UX operating system.
616709	In the Japanese locale, the Add Bunker wizard page showed truncated text.

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

Incident	Description
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.

Veritas Storage Foundation for Databases fixed issues

The following are Veritas Storage Foundation for Databases fixed issues.

Veritas Storage Foundation for Oracle 5.0 MP3 fixed issues

The following new additional issues for Veritas Storage Foundation for Oracle have been fixed in the 5.0 MP3 release.

Using Web GUI to view snapplan log successful (861696)

When trying to view a snapplan log using the Web GUI, you no longer get the error message "Unable to load operation". You can load and view a snapplan log successfully in the Web GUI.

This issue has been fixed.

dbdst_file_move, dbdst_makelbfs, and dbed_vmsnap commands successful on OEL platform (1055246, 1055249, 1065248)

The `dbdst_file_move`, `dbdst_makelbfs`, and `dbed_vmsnap` commands no longer fail on the OEL platform because the `grep`, `mkfs`, and `sed` utilities are located in different directories than on the RHEL platform.

This issue has been fixed.

Veritas Storage Foundation for Oracle 5.0 MP1 and 5.0 MP2 fixed issues

[Table 1-16](#) describes fixed issues in the Veritas Storage Foundation for Oracle 5.0 MP1 and 5.0 MP2 releases.

Table 1-16 Veritas Storage Foundation for Oracle fixed issues

Incident	Description
567342	An unmounted checkpoint clone database reappears in the Java GUI tree after rescanning.
582069	<p>If SFDB commands are executed with a locale that differs from the locale in use when the SFDB server was started, the commands may fail with the following message:</p> <pre data-bbox="534 517 1022 569">([Sybase][ODBC Driver][Adaptive Server Anywhere]Syntax error).</pre>
582416	<p>Clicking the Help button on a GUI wizard produces the following error message:</p> <pre data-bbox="534 699 1150 751">Error V-39-53246-8 Get EntryPoint failed. Please check the manifest related information</pre>
600431	Storage Checkpoint operations are currently not supported for databases cloned with Database FlashSnap.
607001	Repository changes resulting from executing SFDB Storage Checkpoint CLIs are no longer delayed in the SFDB GUI.
608697	You can now refresh the View Statistics wizard in the Firefox browser.
609682	Specifying with dbed_vmsnap a snapplan that does not exist no longer produces the error SFORA dbed_vmsnap ERROR V-81-6518 Could not find snapplan 'snap_plan' in repository.
786989	The qio_getdbfiles_ora script now detects when an Oracle instance is in Standby mode.

Table 1-16 Veritas Storage Foundation for Oracle fixed issues (*continued*)

Incident	Description
853363	<p>The I/O performance of EMC Symmetrix arrays has been improved in this release.</p> <p>To enable these changes, after upgrading to this release, set the discovery mode of the VAIL provider to discover only those Symmetrix devices that are visible to the host, in the following order:</p> <ul style="list-style-type: none"> ■ Determine the agent name under which the Symmetrix provider is configured: <pre data-bbox="615 578 1206 630"># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh -l</pre> <p>The agent name will be "VAILAgent" for installations of Veritas Storage Foundation, Veritas Storage Foundation for Databases, or Veritas Storage Foundation for RAC. It will be "StorageAgent" if VxFAS is configured.</p> ■ Set the discovery mode to discover host-visible devices only: <pre data-bbox="615 829 1231 881"># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh \ -a agent_name -s 0</pre> <p>where <code>agent_name</code> is the agent name output from the <code>-l</code> option in the previous step.</p>

Veritas Storage Foundation for DB2 5.0 MP3 fixed issues

The following additional issues for Veritas Storage Foundation for DB2 have been fixed in the 5.0 MP3 release.

Using Web GUI to view snapplan log (861696)

When trying to view a snapplan log using the Web GUI, you no longer get the error message "Unable to load operation".

This issue has been fixed.

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 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

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)

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

There is no version 8.01.06 of the HBA driver for RHEL5.

Installation issue with RHEL5 and rsh

Due to known limitations with the targeted policy shipped with RHEL5 and `rsh`, the product may not install correctly. Use `ssh` when using the CPI to install the product on cluster nodes.

Product installer does not support upgrading directly from Storage Foundation 5.0 to Storage Foundation Cluster File System 5.0 MP3

The Storage Foundation product installer does not support upgrading directly from Storage Foundation 5.0 to Storage Foundation Cluster File System 5.0 MP3. Instead, the recommended procedure is to first install the Storage Foundation 5.0 MP3 packages, and then upgrade to Storage Foundation Cluster File System.

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.

Upgrade issues

The following are additional upgrade 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 `--no-preun` 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.

Upgrading the DCLI package

The `--no-preun` option should be specified if you use the `rpm -U` command to upgrade the distributed command line interface (`VRTSdcli`) package.

This issue does not apply 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 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.

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 destroy of LVM partitions. For the Physical Volumes created on CCISS device disk partitions, `vxvmconvert` will succeed.

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.

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.

Upgrading external ASLs and APMs

The Storage Foundation 5.0 MP3 release only supports upgrading external ASL and APM packages from release 5.0 or later.

If your system uses any existing ASLs and APMs from 4.x release, you may need to download a 5.0 version from Symantec. Check the latest array support information to determine whether an updated ASL or APM package is available for your arrays.

See the 5.0 MP3 Hardware Compatibility List for information about supported arrays.

Before upgrading a Storage Foundation product, you must manually remove any existing external ASL or APM packages which are lower than release 5.0.

After completing the upgrade, obtain the required updated ASLs or APMs to ensure the array is claimed correctly.

See the upgrading chapter of the *Storage Foundation Installation Guide* for details.

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

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

Rootability issues of Veritas Volume Manager

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

Encapsulated root disk cannot be mirrored (797829)

If persistent device naming is configured, and a new array is added to a system, the names of DMP nodes may no longer correspond to the operating system-based device names. If you then attempt to mirror the encapsulated root disk, this can result in errors such as the following.

```
VxVM vxrootmir ERROR V-5-2-1969
Mirror disk sdb is smaller than Root disk sda
Cannot mirror root disk to sdb
VxVM vxrootmir ERROR V-5-2-1970 Mirror disk and Root disk should
have same geometry
```

In this example, the DMP node, `sdb`, no longer corresponds to the OS-based device following the addition of a new disk array.

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.

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.

Fabric monitoring (784343)

The new Fabric Monitoring feature controls whether the Event Source daemon (`vxesd`) uses the Storage Networking Industry Association (SNIA) HBA API. This API allows DMP to improve the performance of failover by collecting information about the SAN topology and by monitoring fabric events.

Note: The vendor-provided ASL must also support the use of the SNIA HBA API.

Fabric monitoring may be turned on or off by using the following `vxdmppadm` commands:

```
# vxdmppadm settune monitor_fabric=on
# vxdmppadm settune monitor_fabric=off
```

The current setting of `monitor_fabric` can be displayed by using the following command:

```
# vxdmppadm gettune monitor_fabric
```

The default setting of `monitor_fabric` is `on`.

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`.

Replacing a disk in a site-consistent disk group (536853)

If the `vxdiskadm` command is used to replace a disk in site-consistent disk group, the new disk is expected to be tagged with the same site name as the disk that is being replaced. If the sites do not match, `vxdiskadm` cannot complete the replacement without disabling site-consistency on the volume.

To avoid this, tag the replacement disk with same site name as the disk that is being replaced:

```
# vxdisk settag replacement_disk site=sitename
```

After tagging the replacement disk, you can use `vxdiskadm` to replace the failed disk.

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.

Volume tags not displayed (602953)

In the VEA client for Microsoft Windows systems, existing volume tags are not displayed when adding a new volume tag.

Search does not return any objects for non-Administrator users (840452)

A search that is performed by a user in a non-Administrator group should return an access-denied error and not an empty list of objects.

Veritas Web GUI known issues

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

Incorrect error message when importing a disk group (607096)

An incorrect error message such as the following may be displayed when importing a disk group.

```
<!--td align="center" height="287" valign="midd
```

Solaris x64 hosts cannot be managed (615554)

The Web GUI cannot be used to manage Solaris x64 for Opteron hosts that are running the Storage Foundation 4.1 software.

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.
```

Veritas File System known issues

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 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:

```
ioctl132(mkfs.vxfs:7471): Unknown cmd fd(3) cmd(564f4c02){16}  
arg(fff51138) on /dev/sdc2
```

You can safely ignore these messages.

SLES 10 boot sequence displays error messages due to Veritas File System entries in /etc/fstab (1074101)

The boot sequence on a SLES 10 system displays error messages due to Veritas File System (VxFS) entries in the `/etc/fstab` file. The errors display because VxFS does not support the `-f` option of the operating system's mount command, which is passed during the boot sequence.

Workaround

Issue the `mount -i -a -f -t vxfs` command to update the `/etc/mtab` file with all missing mounted VxFS entries in the `/etc/fstab` file after rebooting the system.

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>

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 MP3 known issues

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

During system startup when on RHEL 5 and OEL 5, with SELinux enabled and with the VRTSdcli rpm installed, an error message may appear (1303170)

During system startup when on RHEL 5 and OEL 5, with SELinux enabled and with the VRTSdcli rpm installed, the following error message may appear:

```
Starting vxdcli.sh:
/usr/sbin/vxadm: error while loading shared libraries:
/opt/VRTSdcli/lib/libvxdcli.so:
cannot restore segment prot after reloc: Permission denied
```

or in the `/var/log/audit/audit.log`

```
type=AVC msg=audit(1216288241.978:197) :
avc: denied { execmod } for pid=8762
comm="vxdclid" path="
/opt/VRTSdcli/lib/libvxdcli.so" dev=sda3 ino=4570006
scontext=root:system_r:unconfined_t:s0-s0:c0.c1023
tcontext=system_u:object_r:lib_t:s0 tclass=file
```

Performing an upgrade from 4.1 MP4 to 5.0 MP3 failed on RHEL5 U2 (1300743)

After installing Storage Foundation Cluster File System 4.1 MP4 on RHEL5, and after upgrading the OS from RHEL 5 to RHEL5 U2, the SFCFS stack upgrade failed.

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.

Converting a standalone UNIX host in a cluster to a Storage Foundation Manager managed host fails (1361639)

If a Storage Foundation (SF) host on UNIX, that is added to a cluster, has not been configured using the `configure` option of the `installsf` command, it cannot be migrated to be managed by SF Manager. The migration fails with a message stating that the Veritas Provider Access layer (VxPAL) agents are not installed.

Veritas Storage Foundation Cluster File System 5.0 MP2 known issues

The following are additional known issues in the 5.0 MP2 release of Veritas Storage Foundation Cluster File System.

Oracle Disk Manager is incompatible with auditing on SLES9

Oracle Disk Manager (ODM) is incompatible with the auditing subsystem on SUSE Linux Enterprise Server 9 systems.

Veritas Storage Foundation Cluster File System 5.0 MP1 known issues

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

Oracle-Linux (840486)

Problems uninstalling or upgrading Veritas Storage Foundation for Oracle when Veritas Storage Foundation Cluster File System is installed on the same system.

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine, do not use the installer to uninstall if you are planning to uninstall only one product.

If you want to uninstall the product, you must uninstall the Veritas Storage Foundation for Oracle packages manually.

To uninstall the Veritas Storage Foundation for Oracle packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 Stop the repository database and unmount the repository volume.

In a stand-alone configuration:

Stop the database repository:

```
# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb
```

Unmount the database repository:

```
# /opt/VRTSdbcom/config/sfua_rep_mount stop
```

In an HA configuration:

Stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Remove the Veritas Storage Foundation for Oracle packages using the `rpm -e` command.

```
# rpm -e VRTSorgui-common VRTSdbed-common VRTSdbcom-common \
VRTSdbdoc
```

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine and you are upgrading both products, use the installer to upgrade Veritas Storage Foundation Cluster File System first. Then, use the installer to upgrade Veritas Storage Foundation for Oracle.

If the second upgrade fails, remove the Veritas Storage Foundation for Oracle packages as described above, then run the installer to upgrade Veritas Storage Foundation for Oracle.

DB2-Linux (840486)

Problems uninstalling or upgrading Veritas Storage Foundation for DB2 when Veritas Storage Foundation Cluster File System is installed on the same system.

If Veritas Storage Foundation for DB2 and Veritas Storage Foundation Cluster File System are installed on the same machine, do not use the installer to uninstall if you are planning to uninstall only one product.

If you want to uninstall the product, you must uninstall the Veritas Storage Foundation for DB2 packages manually.

To uninstall the Veritas Storage Foundation for DB2 packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 Stop the repository database and unmount the repository volume.

In a stand-alone configuration: Stop the database repository:

```
# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb
```

Unmount the database repository:

```
# /opt/VRTSdbcom/config/sfua_rep_mount stop
```

In an HA configuration:

Stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hstop -local
```

To stop VCS processes on all systems:

```
# hstop -all
```

- 3 Remove the Veritas Storage Foundation for DB2 packages using the `rpm -e` command.

```
# rpm -e VRTSd2gui-common VRTSdb2ed-common VRTSdbcom-common VRTSdbdoc
```

If Veritas Storage Foundation for DB2 and Veritas Storage Foundation Cluster File System are installed on the same machine and you are upgrading both products, use the installer to upgrade Veritas Storage Foundation Cluster File System first. Then, use the installer to upgrade Veritas Storage Foundation for DB2.

If the second upgrade fails, remove the Veritas Storage Foundation for DB2 packages as described above, then run the installer to upgrade Veritas Storage Foundation for DB2.

Veritas Storage Foundation Cluster File System for Oracle RAC known issues

There are no additional software issues in 5.0 MP3 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

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:

```
# vxdg 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.

Upgrade issue (840927, 840955)

Replicated Volume Groups (RVGs) are not displayed by the Storage Foundation Web GUI as well as the VVR Java GUI after upgrade because of a bug in the Veritas Volume Replicator provider (VRTSvrpro) and Veritas Volume Replicator Web Console (VRTSvrw) packages in the 5.0 release. If the `rpm` command is used to upgrade to 5.0 MP1, the entries for the VVR provider are removed from the registry. This results in the Web GUI and Java GUI not being able to display any Replicated Volume Groups that are present on the system.

The workaround is to upgrade to the 5.0 MP1 VRTSvrpro and VRTSvrw packages by using the following forms of the `rpm -U` command:

```
# rpm -U --nopreun VRTSvrpro
# rpm -U --nopreun --nopostun VRTSvrw
```

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

Synchronizing volumes and RVG with large volumes (840217)

The `vradmin syncrvg` and the `vradmin syncvol` commands do not work correctly for volumes larger than 1TB. When either of these two commands is used to synchronize large volumes, the command still runs, but it reports wrong total size of the volumes being synchronized and it actually synchronizes only a portion of the volume having size larger than 1TB.

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 Storage Foundation for Oracle known issues

Known issues in the Veritas Storage Foundation for Oracle 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 for Oracle 5.0 MP3 known issues

The following new known issues exist in the 5.0 MP3 release of Veritas Storage Foundation for Oracle.

Datafile rollback in a Dynamic Storage Tiering environment fails to rollback the files (1227083)

In a Dynamic Storage Tiering setup, when you execute the `dbed_ckptrollback` command using the `-F <datafile>` option, the operation may fail and display the following error:

```
#$ dbed_ckptrollback -S $ORACLE_SID -H $ORACLE_HOME -c Checkpoint_name
-F datafile_name.dbf
```

```
Rolling back files using Checkpoint_XXXXXXXXX ...
VXDBA_PRODUCT vxckptadm WARNING V-81-4565
The following files/directories do not reside on VxFS file systems:
VXDBA_PRODUCT vxckptadm WARNING V-81-999
datafile_name.dbf
VXDBA_PRODUCT vxckptadm
ERROR V-81-4562 Storage Rollback failed.
SFORA rb.file
ERROR V-81-3046 Could not roll back datafiles.
```

This is only observed in the Dynamic Storage Tiering setup.

Workaround

You must restart the Veritas `vxdbd` daemon using the following procedure.

To restart the Veritas vxdbd daemon.

- 1 Check the Veritas vxdbd daemon status using the following command:

```
# /etc/init.d/vxdbdctrl status
Status of Veritas vxdbd
/opt/VRTSdbcom/bin/vxdbd ping SUCCESS
```

- 2 Stop the Veritas vxdbd daemon using the following command:

```
# /etc/init.d/vxdbdctrl stop
Stopping Veritas vxdbd
```

- 3 Start the Veritas vxdbd daemon using the following command:

```
# /etc/init.d/vxdbdctrl start
Starting Veritas vxdbd
/opt/VRTSdbcom/bin/vxdbd start SUCCESS
```

After restarting the Veritas vxdbd daemon, you may attempt the rollback operation again.

VRTSddlpr directory may exist in /opt/ after uninstalling the stack (1315258)

After uninstalling VRTSddlpr package, the /opt/VRTSddlpr directory may still exist.

Workaround

After uninstalling the VRTSddlpr package, you must manually clean up the /opt/VRTSddlpr directory.

VRTSdbcom directory may exist in /opt/ after uninstalling the stack (1315247)

After uninstalling the stack, the directory /opt/VRTSdbcom may still exist.

Workaround

After uninstalling the stack, you must manually clean up the /opt/VRTSdbcom directory .

VRTSdbms3 directory may exist in /opt/ after uninstalling the stack (1315262)

After uninstalling VRTSdbms3 package, the /opt/VRTSdbms3 directory may still exist.

Workaround

After uninstalling the `VRTSdbms3` package, you must manually clean up the `/opt/VRTSdbms3` directory.

VRTSvxms directory exists in /opt/ even after uninstalling the stack (1316534)

After uninstalling the stack, the `VRTSvxms` directory may still exist.

Workaround

After uninstalling the stack, you must check the `VRTSvxms` directory and all its subdirectories to see if they are empty.

You may manually clean up the `VRTSvxms` directory and all its subdirectories only if they are empty.

Reverse Resync not supported if database is created using Oracle Managed Files (1192729)

If an Oracle database is created with Oracle Managed Files (OMF), then `reverse_resync` operations would fail.

The following errors are displayed:

```
oracle@swlx07:~> /opt/VRTSdbed/bin/dbed_vmsnap -S $ORACLE_SID \  
-f sp4 -o
```

```
reverse_resync_begin  
dbed_vmsnap started at 2007-12-28 12:02:42
```

```
SFORA dbed_vmsnap WARNING V-81-5725 After reverse_resync_commit  
is performed, you need to recreate the Authentication Password  
File using the ORAPWD utility.
```

```
SFORA dbed_vmsnap ERROR V-81-4882 An error occurred while  
reconfiguring Oracle instance 'sfora'.
```

```
SFORA dbed_vmsnap ERROR V-81-4881 Log file is at  
/tmp/dbed_vmclonedb.12313/nomount.log.
```

```
SFORA dbed_vmsnap ERROR V-81-4918 Database sfora has not been  
correctly recovered.
```

```
SFORA dbed_vmsnap ERROR V-81-4881 Log file is at  
/tmp/dbed_vmclonedb.12313/recover.log.
```

Workaround

The `reverse_resync` operation for a database created with OMF is not supported in the 5.0 MP3 release.

There is no workaround for this issue.

Veritas Storage Foundation for Oracle 5.0 MP2 known issues

The following are known issues in the 5.0 MP2 release of Veritas Storage Foundation for Oracle.

Cannot unmount single-host clone in HA environment after failover (818522)

In an HA environment, after successfully taking a snapshot and cloning the database on the same host where primary is running, if a node failover happens then `dbed_vmclonedb -o umount` does not work.

Problems uninstalling or upgrading Veritas Storage Foundation for Oracle when Veritas Storage Foundation Cluster File System is installed on the same system (840486)

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine, do not use the installer to uninstall if you are planning to uninstall only one product.

You must uninstall the Veritas Storage Foundation for Oracle packages manually if you want to uninstall the product.

To uninstall the Veritas Storage Foundation for Oracle packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 Stop the repository database and unmount the repository volume.

In a stand-alone configuration:

Stop the database repository:

```
# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb
```

Unmount the database repository:

```
# /opt/VRTSdbcom/config/sfua_rep_mount stop
```

In an HA configuration:

Stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Remove the Veritas Storage Foundation for Oracle packages using the `rpm -e` command.

```
# rpm -e VRTSorgui-common VRTSdbed-common VRTSdbcom-common \
VRTSdbdoc
```

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine and you are upgrading both products, use the installer to upgrade Veritas Storage Foundation Cluster File System first. Then, use the installer to upgrade Veritas Storage Foundation for Oracle.

If the second upgrade fails, remove the Veritas Storage Foundation for Oracle packages as described above, then run the installer to upgrade Veritas Storage Foundation for Oracle.

Veritas Storage Foundation for Oracle 5.0 MP1 known issues

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

Cannot unmount single-host clone in HA environment after failover (818522)

In an HA environment, after successfully taking a snapshot and cloning the database on the same host where primary is running, if a node failover happens then `dbed_vmclonedb -o umount` does not work.

Problems uninstalling or upgrading Veritas Storage Foundation for Oracle when Veritas Storage Foundation Cluster File System is installed on the same system (840486)

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine, do not use the installer to uninstall if you are planning to uninstall only one product.

You must uninstall the Veritas Storage Foundation for Oracle packages manually if you want to uninstall the product.

To uninstall the Veritas Storage Foundation for Oracle packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 Stop the repository database and unmount the repository volume.

In a stand-alone configuration:

Stop the database repository:

```
# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb
```

Unmount the database repository:

```
# /opt/VRTSdbcom/config/sfua_rep_mount stop
```

In an HA configuration:

Stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hstop -local
```

To stop VCS processes on all systems:

```
# hstop -all
```

- 3 Remove the Veritas Storage Foundation for Oracle packages using the `rpm -e` command.

```
# rpm -e VRTSorgui-common VRTSdbed-common VRTSdbcom-common \  
VRTSdbdoc
```

dbed_vmclonedb -p failed to create clonedb with modified pfile (852188)

If you are running the `dbed_vmclonedb -p` or the `dbed_clonedb -p` command, the pfile modification will fail if there is an unquoted or unescaped special character in the primary instance's pfile. The following error will be displayed:

```
SFORA pfile_mod ERROR V-81-5781 Parse error in file  
/oracle/dbs/<pfile_name>. line 6: .
```

```
SFORA dbed_vmclonedb WARNING V-81-5788 Pfile modification failed.  
Clone instance <CLONE SID> may not start.
```

Workaround

To avoid this issue, make sure all special characters in the primary instance's pfile are either placed within quotes or escaped.

You can check the Oracle Reference Manual for a list of special characters which must be either placed within quotes or escaped when used as pfile parameter values. In some cases, Oracle will process pfile correctly at startup even if a parameter values contains unquoted special characters. However, the pfile parser we use strictly enforces the pfile specification contained in the Oracle Reference Manual.

Note: The primary instance's pfile is saved at the time of snapshot creation. If you attempt to clone the database using that snapshot you will be using the saved pfile, not the current pfile. Therefore you must create a new snapshot in order to ensure that the clone will use an updated pfile.

One-time scheduled tasks need Specific Date (861274)

When scheduling a one-time task from the GUI, the task may not be executed if a Specific Date (Include Date) is not set for it.

Database FlashSnap archive log destinations (862092, 862687)

With Oracle Release 10g and above, Database FlashSnap clones do not support `DB_RECOVERY_FILE_DESTINATION` as the sole mandatory archive log destination. This issue will not be detected by FlashSnap validation with `dbed_vmchecksnap`, or by the snapshot command `dbed_vmsnap`. However, recovery will fail when attempting to clone a database using the snapshot, and the message "ORA-01195: online backup of file 1 needs more recovery to be consistent" may appear in the log file.

Workaround

Define a mandatory log archive destination that is not `DB_RECOVERY_FILE_DESTINATION` and set the `ARCHIVELOG_DEST` parameter of the `snappan` to this value.

Veritas Storage Foundation for DB2 known issues

Known issues in the Veritas Storage Foundation for DB2 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 for DB2 5.0 MP3 known issues

The following new additional known issues exist in the 5.0 MP3 release of Veritas Storage Foundation for DB2.

Cannot completely rollback a 5.0GA checkpoint in 5.0 MP3 (1228785)

Storage checkpoints which were created in the 5.0 GA release environment do not work as expected when you rollback in 5.0 MP3 environment. When you rollback these checkpoints, the `dbed_ckptrollback` command completes successfully but gives an 'unexpected end of file' error.

An error message similar to the following is displayed :

```
gunzip: /tmp/11547203191175220319723820319/dbcfg_to_use.  
NODE0000.Z: unexpected end of file
```

Workaround

There is no workaround for this issue.

Incorrect usage for Database Storage Tiering commands (1233004)

In SUSE Linux Enterprise Edition 10 (SLES10), Database Storage Tiering commands for Storage Foundation for DB2 display incorrect usage information.

Workaround

You must uninstall the `orarun` package on your system. The `orarun` package is used by Novell/SUSE to automate several Oracle pre-install tasks. After uninstalling this package, the Database Storage Tiering commands for Storage Foundation for DB2 will execute correctly on your system.

Another workaround is to unset Oracle environment variables (`ORACLE_HOME` & `ORACLE_SID`) before executing any of the Database Storage Tiering commands.

Veritas Storage Foundation for DB2 5.0 MP2 known issues

The following are additional known issues in the 5.0 MP2 release of Veritas Storage Foundation for DB2.

Veritas Storage Foundation for DB2 5.0 MP1 known issues

The following are additional known issues in the 5.0 MP1 release of Veritas Storage Foundation for DB2.

Problems uninstalling or upgrading Veritas Storage Foundation for DB2 when Veritas Storage Foundation Cluster File System is installed on the same system (840486)

If Veritas Storage Foundation for DB2 and Veritas Storage Foundation Cluster File System are installed on the same machine, do not use the installer to uninstall if you are planning to uninstall only one product.

You must uninstall the Veritas Storage Foundation for DB2 packages manually if you want to uninstall the product.

To uninstall the Veritas Storage Foundation for DB2 packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 Stop the repository database and unmount the repository volume.

In a stand-alone configuration:

Stop the database repository:

```
# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb
```

Unmount the database repository:

```
# /opt/VRTSdbcom/config/sfua_rep_mount stop
```

In an HA configuration:

Stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hstop -local
```

To stop VCS processes on all systems:

```
# hstop -all
```

- 3 Remove the Veritas Storage Foundation for DB2 packages using the `rpm -e` command.

```
# rpm -e VRTSd2gui-common VRTSdb2ed-common VRTSdbcom-commonVRTSdbdoc
```

One-time scheduled tasks need Specific Date (861274)

When scheduling a one-time task from the GUI, the task may not be executed if a Specific Date (Include Date) is not set for it.

Veritas Enterprise Administrators known issues

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

Search does not return any objects for non-Administrator users (840452)

A search that is performed by a user in a non-Administrator group should return an access-denied error and not an empty list of objects. The workaround is to add the user to the Administrator group.

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).
-
- The following features will not be supported in the next release of Veritas Storage Foundation for Oracle and Veritas Storage Foundation for DB2:
 - Java GUI interface for all SF Oracle features
 - Mapping related functions like ORAMAP, `vxstorage_stats` and `dbed_analyzer` for all versions of ORACLE
 - DB2 tools like Flashsnap, checkpoint, and DBDST for all versions of DB2.

Documentation

Product guides are available on the documentation disc in PDF and HTML formats.

Veritas documentation disc

The Veritas documentation disc provides searchable, HTML documentation for each product in this release. Printable PDF documents are also included on the disc.

All documentation is organized by product groups and located on the Veritas documentation disc.

Symantec's Veritas Storage Foundation 5.0 MP3 documentation set is available at the following URL:

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

Note: The Veritas Storage Foundation 5.0 MP3 documentation will be posted at once the product is released.

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>	<code>3rdpartyattributions.pdf</code>
<i>Veritas FlashSnap Point-In-Time Copy Solutions Administrator's Guide</i>	<code>flashsnap_admin.pdf</code>
<i>Veritas Storage Foundation and High Availability Getting Started Guide</i>	<code>getting_started.pdf</code>
<i>Read me end user license agreement documentation</i>	<code>README_EULA</code>
<i>Veritas Storage Foundation read me first</i>	<code>readme_first.txt</code>
<i>Veritas Storage Foundation Release Notes</i> (this document)	<code>sf_notes.pdf</code>
<i>Veritas Storage Foundation Installation Guide</i>	<code>sf_install.pdf</code>

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

Guide Title	Filename
<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
<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>	sf_cfs_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
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents.pdf
<i>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

Table 1-20 Guides in Veritas Volume Replicator documentation set (*continued*)

Guide Title	Filename
<i>Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide</i>	vvr_agents_config.pdf

Veritas Storage Foundation for Databases documentation Changes

The 5.0 release introduces changes to the documentation for Veritas Storage Foundation for Databases. The `VRTSordoc` and `VRTSd2doc` packages no longer exist. All database guides are now in one package, `VRTSdbdoc`. New guides have been created for the GUI documentation and the PDF names have changed for existing guides.

[Table 1-21](#) indicates the PDF and document changes in this release.

Table 1-21 Veritas Storage Foundation for Databases guides

Guide Title	Previous PDF Name	Current PDF Name
<i>Veritas Storage Foundation for DB2 Administrator's Guide</i>	sf_db2_dbag.pdf	sf_db2_admin.pdf
<i>Veritas Storage Foundation for Oracle Administrator's Guide</i>	sf_ora_dbag.pdf	sf_ora_admin.pdf
<i>Veritas Storage Foundation for DB2 Graphical User Interface Guide (NEW)</i>	N/A	sf_db2_gui.pdf
<i>Veritas Storage Foundation for Oracle Graphical User Interface Guide (NEW)</i>	N/A	sf_ora_gui.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.

Manual pages errata

Several manual pages have been updated in this Maintenance Pack to include corrections for several errors or omissions

Dbdst_report man page errata (1361604)

For the `dbdst_report` man page, one of the usage examples is wrongly displayed as follows:

```
$ /opt/VRTS/bin/dbdst_show_fs -S $ORACLE_SID -o audit
```

The correct usage example is:

```
$ /opt/VRTS/bin/dbdst_report -S $ORACLE_SID -o audit
```

Workaround

There is no workaround for this issue.

Dbdst man pages errata (1121091)

You cannot access the correct man pages for the `dbdst` commands if you type in the command name as follows:

```
# man dsdst_command_name
```

Workaround

You must append the correct database suffix for each `dbdst` command in order to access the correct man page for that `dbdst` command.

To obtain correct usage notes for the `dbdst` command in an Oracle environment, you must use the following command format:

```
# man dsdst_command_name_ora
```


To obtain correct usage notes for the `dbdst` command in a DB2 environment, you must use the following command format:

```
# man dsdst_command_name_db2
```

Veritas Storage Foundation for DB2 Administrator's Guide errata

The following section is missing from the *Veritas Storage Foundation for DB2 Administrator's Guide*.

Reconfigure virtual IP address for repository configuration

When configuring a two-node cluster, use the following to change the virtual IP address.

In a standalone instance, first change the IP address. Then run the `sfua_db_config` once to update IP information for SFUA repository access.

```
# /opt/VRTSdbcom/bin/sfua_db_config
```

To configure a two-node cluster in a cluster environment

- 1 Change the IP address for the cluster.
- 2 Update the IP address for the repository configuration in HA environment by running the following set of commands.
- 3 Unconfigure the SFUA repository.

```
# /opt/VRTSdbcom/bin/sfua_db_config -o unconfig_cluster
```

- 4 Import the repository disk group.
- 5 Then, start then repository disk volume.
- 6 Mount the repository file system.
- 7 Then, run the command.

```
# /opt/VRTSdbcom/bin/sfua_db_config
```

- 8 When prompted, select the option to change the configuration parameters for the cluster configuration.
- 9 Enter the new cluster IP address for the cluster configuration.

Veritas Storage Foundation 5.0 MP3 for Xen

This appendix includes the following topics:

- [Supported Platforms for Xen](#)
- [Supported Features](#)
- [Unsupported Features](#)
- [Packages Required for Dom0 Installation](#)
- [Installing the Veritas Software in Dom0](#)
- [Verifying Software Versions](#)
- [Using the Veritas Software under Xen](#)
- [Uninstalling the Veritas Software from Dom0](#)

Supported Platforms for Xen

This release provides support for Veritas Storage Foundation 5.0 Maintenance Pack 3 (MP3) on the Xen platform for Linux.

Supported systems are:

- Red Hat Enterprise Linux 5 U1 (RHEL 5 U1)
- Red Hat Enterprise Linux 5 U2 (RHEL 5 U2)
- SUSE Linux Enterprise Server 10 SP1 (SLES 10 SP1).
- SUSE Linux Enterprise Server 10 SP2 (SLES 10 SP2).

For more information about installing and configuring Xen, read one of the following documents:

- RHEL 5 *Virtualization Guide*
- “Virtual Machine Server” chapter of the SLES 10 *Installation and Administration Guide*.

The packages provided in this release operate on the following hardware and operating systems:

- Intel Xeon 64-bit and AMD Opteron 64-bit systems (x86_64).
- For RHEL 5 U1 and SLES 10 SP1, use the Xen kernel that is included with the distribution.
- The disk arrays supported by this release are listed in the Hardware Compatibility List (HCL) at:

<http://entsupport.symantec.com/docs/283161>

This release provides specific versions of the Array Support Libraries (ASLs) and Array Policy Modules (APMs) to support the following arrays on the Xen platform:

- EMC CLARiiON CXx00 series
- Hitachi TagmaStore Adaptable Modular Storage (AMS)
- Hitachi TagmaStore Workgroup Modular Storage (WMS)
- Hitachi Thunder 9500 V series
- HP StorageWorks Enterprise Virtual Array (EVA) 3000, 4000, 5000, 6000 and 8000

The following ASLs and APMs support these arrays:

```
VRTSEMC-CLAR-asl-5.0.30.00-CAP_MP3_platform_name.x86_64.rpm  
VRTSEMC-CLAR-apm-5.0.30.00-CAP_MP3_platform_name.x86_64.rpm  
VRTSHDS-9500-asl-5.0.30.00-CAP_MP3_platform_name.x86_64.rpm  
VRTSHDS-9500-apm-5.0.30.00-CAP_MP3_platform_name.x86_64.rpm  
VRTSHP-EVA-asl-5.0.30.00-CAP_MP3_platform_name.x86_64.rpm  
VRTSHP-EVA-apm-5.0.30.00-CAP_MP3_platform_name.x86_64.rpm
```

where *platform_name* is RHEL5 or SLES10, as appropriate.

For the other supported arrays, you can use the ASLs that are listed in the HCL.

Symantec recommends a minimum of 2GB of memory, a minimum of 1GB configured for domain zero (Dom0), and 512MB for each unprivileged guest domain (DomU).

Supported Features

The following features of the Veritas Storage Foundation Enterprise product are provided in this release to support the Xen environment:

- Configuration of Veritas Volume Manager volumes in Dom0.
- Export of VxVM volumes to DomU as raw partition devices.
- Support for both SCSI devices and IDE disk devices (`/dev/sdx` and `/dev/hdx`).
- Relayout of a volume in Dom0 that has been exported to DomU is possible. Some reduction in I/O throughput may be experienced during the operation.
- Resizing a volume in Dom0 is not seen in DomU until the domain is rebooted. The file system must be resized independently of the volume. This limitation is imposed by Xen.
- Array Support Libraries (ASLs) and Array Policy Modules (APMs) that are supported in 5.0 MP3.
- Root encapsulation.
- SELinux

Unsupported Features

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

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

Packages Required for Dom0 Installation

The Veritas Storage Foundation 5.0 MP3 release includes the following packages to install in Dom0.

For RHEL 5:

```
VRTSvxvm-common-5.0.30.00-CAP_MP3_RHEL5.i686.rpm  
VRTSvxvm-platform-5.0.30.00-CAP_MP3_RHEL5.x86_64.rpm
```

For SLES 10 SP1 and SP2:

```
VRTSvxvm-common-5.0.30.00-CAP_MP3_SLES10.i586.rpm  
VRTSvxvm-platform-5.0.30.00-CAP_MP3_SLES10.x86_64.rpm
```

Apart from any required ASL and APM packages, there are no other Veritas Storage Foundation packages specific for the Xen platform.

Installing the Veritas Software in Dom0

If you plan to install the Veritas Storage Foundation 5.0 MP3 software for the first time and not upgrading an existing system, review the preinstallation instructions.

See the *Veritas Storage Foundation 5.0 MP3 Installation Guide*.

For important release information, review the *Veritas Storage Foundation 5.0 Release Notes*, and all documents in the `release_notes` directory.

To install the Veritas software in Dom0

- 1 Install the Veritas Storage Foundation 5.0 MP3 software from the software disc or downloaded image. Run the Veritas Volume Manager installation script instead of the generic `installer` script. Enter the following command from the `platform/volume_manager` directory of the mounted DVD-ROM or downloaded image:

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

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

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

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

You are prompted to supply a valid Veritas Storage Foundation product license.

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

- 3 If the `installvm` script did not reboot the Xen kernel in step 1, then boot the Xen kernel in Dom0.

- 4 Use the `rpm -ihv` command to install any required ASL and APM packages.
- 5 You can now configure VxVM volumes in Dom0. You can also export the volumes to DomU.

Verifying Software Versions

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

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

Using the Veritas Software under Xen

To enable access to a VxVM volume in DomU

- 1 Create a VxVM volume in Dom0, using a suitable command such as `vxassist` with the required options. For example, the following command creates a 2GB mirrored volume, `myvol1`, in the disk group `mydg`:

```
# vxassist -g mydg make volume myvol1 2g layout=mirror nmirror=2
```

- 2 Add a suitable entry to the Xen configuration file for the guest domain, for example:

```
disk = [ ..., 'phy:vx/dsk/mydg/myvol1,sdb,w' ]
```

This entry specifies that the device file for the volume `myvol1` in the disk group `mydg` is to be exported to the guest domain. The read-writable device `/dev/sdb` represents the volume in the guest domain. If more than one volume is to be exported, use commas to separate their definitions, as shown here:

```
disk=[..., 'phy:vx/dsk/mydg/myvol1,sdb1,w', 'phy:vx/dsk/mydg/myvol2,sdc1,w' ]
```

- 3 To start the guest domain, use the following command specifying the domain configuration file:

```
# xm create -c config_file
```

If the guest domain already exists, you can export a volume from Dom0 by using the following command:

```
# xm block-attach DomU_ID phy:vx/dsk/mydg/myvol1 /dev/sdb rw
```

where DomU_ID is the name or number of the guest domain. The domain number can be obtained by using the `xm list` command.

- 4 After the guest domain has been started, the volume can be accessed in DomU. Access the volume by using the disk or the disk partition to which it has been mapped.

The vbd hotplug script prevents you from mounting a file system as writable in both Dom0 and DomU.

Note: The `xm block-detach` command is currently documented as being broken. This issue may be fixed in a future release of Xen, but this functionality has not been tested in this release.

See the Xen documentation for more information about how to administer the Xen virtual machine.

Uninstalling the Veritas Software from Dom0

To uninstall the Veritas software from Dom0

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

- 4 Stop all VxVM volumes by entering the following command for each disk group:

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

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

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

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

```
# cd /opt/VRTS/install
```

```
# ./uninstallsf
```

