

Veritas Storage Foundation™ Release Notes

AIX

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/283990>

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.

Storage Foundation Release Notes

This document 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 Maintenance Pack 3 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™ 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/282024>

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*.

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.

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* ([sfcfs_install.pdf](#)) and the *Veritas Storage Foundation Cluster File System Release Notes* ([sfcfs_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](#)).

Asynchronous I/O fast path for CIO with Veritas File System

AIX asynchronous I/O (AIO) subsystem provides two different paths for I/O to a regular file. In the first path, I/O requests for a regular file are queued to an AIO queue. The `kproc` kernel process called `aio_server` controls each request from the time a request is taken off the queue until it completes. The number of `aio_server`s limits the number of disk I/O operations that can be in progress in the system simultaneously. The second path where `aio_server` threads are avoided is referred to as fast path. Here, the AIO subsystem routes all requests for a file that is open in CIO mode to the file system for service. The request needs to go through non-blocking interface and it must initiate the I/O and return without waiting for I/O. This results in less context switching and immediate start of the I/O leading to performance improvement. This is particularly helpful for relational databases

that use regular files as database objects. VxFS supports AIO fast path for CIO for local mounted file system. In case of cluster mounted file system, AIO subsystem follows the normal kproc path for all regular files. The AIO fast path for CIO is available in AIX 5.3 TL07 or later.

Partition mobility support

VxVM, VxFS and SFCFS are Live/Active partition migration safes. The Live/Active partition migration process moves a running logical partition and all of its running applications from one physical server to another without disrupting operation of the logical partition. Partition mobility is available in AIX 5.3 TL07 or later.

Storage key support

This release supports the storage key feature introduced in the AIX 6.1 operating system. Storage keys provide a mechanism to catch unauthorized access to memory. The keys operate at a page-level granularity and can be used to read or write protect pages. Once a key is associated with a page, a thread that must access the memory must explicitly get permission to access the memory. Unauthorized access to such memory by a thread results in an exception. Storage keys provide a way to improve the reliability of code.

Storage keys are available only with the AIX 6.1 operating system.

Veritas Concurrent I/O support for Sybase ASE

Veritas Concurrent I/O improves the performance of regular files on a VxFS file system without the need for extending namespaces and presenting the files as devices. This simplifies administrative tasks and permits database applications to access database files concurrently.

In the 5.0MP3 release of Storage Foundation, there is support for Concurrent I/O with Sybase ASE databases. Concurrent I/O is supported for Sybase servers running ASE 12.5.4 and ASE 15.0.2 databases.

Enabling Concurrent I/O

You can enable Concurrent I/O on regular files because you do not need to extend name spaces and present the files as devices.

Before enabling Concurrent I/O on your system, you must review the following pre-requisites:

- To use the Concurrent I/O feature, the file system must be a VxFS file system.
- You must ensure the mount point on which you plan to mount the file system exists.

- You must ensure that the Database administrator can access the mount point.

To enable Concurrent I/O on a file system using mount with the -o cio option

- ◆ Mount the file system using the `mount` command as follows:

```
# /usr/sbin/mount -V vxfs -o cio special mount_point
```

where *special* is a block special device, and *mount_point* is the directory where the file system will be mounted.

Disabling Concurrent I/O

If you need to disable Concurrent I/O, you must unmount the VxFS file system and mount it again without the `mount` option.

To disable Concurrent I/O on a file system using the mount command

- 1 Shutdown the Sybase instance.
- 2 Unmount the file system using the `umount` command.
- 3 Mount the file system again using the `mount` command without using the `-o cio` option.

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.

Workload Partitions support

This release provides support for the Workload Partition (WPAR) feature introduced in the AIX 6.1 operating system. Support for WPARs is limited in scope as follows:

- All Storage Foundation packages must be installed and configured in the global partition of AIX.
- Storage Foundation can administered only in the global partition.
- The storage resources of Storage Foundation, primarily file system mount points, can be used to create and configure system or application WPARs.
- File system administration within a WPAR is limited in scope and matches the scope of administration that is available for JFS2 as of this release.
- The execution of Storage Foundation administrative tools may fail without appropriate error messages.

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.

Virtual SCSI device support

Veritas Volume Manager 5.0MP1 introduced support for virtual SCSI (vSCSI) devices.

This release, 5.0MP3, integrates the administration commands for vSCSI devices into the `vxddladm` utility and the `vxddladm` utility. These utilities replace the `vxddladm` utility that was introduced in VxVM 5.0MP1. The `vxddladm` utility lets you enable and disable DMP on vSCSI devices, add and remove supported arrays, and list supported arrays. The `vxddladm` utility controls the I/O policy and the path policy for vSCSI devices.

In this release, DMP is enabled on VIO clients by default. All VSCSI devices are under DMP control. Use the `vxddladm disablevscsi` command followed by a reboot to disable DMP and enable MPIO on these vSCSI devices.

Note the following limitations of support for vSCSI devices:

- This feature is supported on AIX 5.3 TL7 SP2 and later, and on AIX 6.1 TL0 SP4 and later.
- Cross-platform Data Sharing (CDS) (also known as Portable Data Containers (PDC)) is not supported for vSCSI devices. Such devices must be initialized as `aixdisk` format disks, and be added to non-CDS disk groups. By default, VxVM tries to initialize disks with the `cdsdisk` format. To avoid this behavior, create the file `/etc/defaults/vxdisk`, and add the following line to this file:

```
format=aixdisk
```

To turn off CDS as the default for disk groups, create the file `/etc/defaults/vxdg`, and add the following line:

```
cds=off
```

With the defaults files set up in this way, VxVM initializes disks with the `aixdisk` format in non-CDS disk groups by default.

- vSCSI devices do not support the SCSI-3 Persistent Group Reservation feature. This limitation prevents SFRAC functionality from being used.

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 `vxmpadm 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 `vxmpadm getdmpnode` command to display any devices that have fewer than the minimum number of paths. You can also configure DMP to notify you when the number of active paths falls below the configured minimum.

Enhanced subpaths listing

The `vxmpadm getsubpaths` command now provides the ability to list all subpaths known to DMP, subpaths of an enclosure, or subpaths through an array port or

`pwwn`. To list the paths through an array port, specify either a combination of enclosure name and array port id, or the array port WWN.

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

Enhanced I/O statistics

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

Queued and Erroneous I/O counts

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

Filter zero entries

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

Specifying units for statistics data

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

Cumulative I/O statistics

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

Miscellaneous improvements to DMP I/O statistics

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

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

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

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

The `vxddmpadm iostat show` command now shows full names for disks with target WWN. If the output exceeds 80 columns, then the output is wrapped.

Making DMP restore options persistent

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

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

New log file location for DMP events

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

Extended device attributes displayed in vxdisk list

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

Display use_all_paths attribute for an enclosure

Display value of `use_all_paths` attribute for an enclosure.

Viewing information about the ASLs installed on the system

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

Displaying the count of LUNs in an enclosure

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

Displaying LUN serial number

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

Displaying HBA details

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

New exclude and include options for the vxddmpadm command

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

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

New command for reporting DMP node information

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

Setting attributes for all enclosures

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

Support for ALUA JBOD devices

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

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

Enhancements to the Cluster Volume Manager

Cluster Volume Manager has been enhanced.

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 `vxdg 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 `vxdg` 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 `fspadm` command now supports UID, GID, and TAG elements in the placement policy XML file.
- Improved `fspadm` command scan performance.
- Suppressed the processing of the chosen RULE.
- Parser support for UID, GID, and TAG elements in a DST policy.
- What-if support for `analyze` and `enforce` without requiring the policy to be assigned.
- Storage Checkpoint data placement support in a DST policy.
- Upgrade to SQLite 3.3.9
- Shared DB thread handle support
- CPU and I/O throttling support for DST scans.
- New command, `fstag`, for file tagging.
- New command, `fspm`, for creating XML policies.

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 `umount` command reverses the `mntlock` option if you previously locked the file system.

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

Veritas Storage Foundation for DB2 provides full support for the DB2 9.5 database in the 5.0 MP3 release.

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 the *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 AIX

The following sections list the supported features for AIX 5.3 and AIX 6.1 in the 5.0 MP3 release of Storage Foundation for Databases.

Supported features for AIX 5.3

For AIX 5.3, the following features of Storage Foundation for Databases are supported in this release.

Storage Foundation for Oracle features

For Oracle 9i, Oracle 10g, and Oracle 11g, all features of Storage Foundation for Oracle are supported in this release.

Storage Foundation for DB2 features

For DB2 8.1, DB2 8.2, DB2 9.1 and DB2 9.5, all features of Storage Foundation for DB2 are supported in this release.

Supported features for AIX 6.1

For AIX 6.1, the following features of Storage Foundation for Databases are supported in this release.

Storage Foundation for Oracle features

For Oracle 9i, Oracle 10g and Oracle 11g, all features of Storage Foundation for Oracle are supported in this release.

Storage Foundation for DB2 features

For DB2 8.1, DB2 8.2, DB2 9.1, and DB2 9.5, all features of Storage Foundation for DB2 are supported in this release.

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 AIX operating systems

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

Product installation scripts verify the required update levels. The installation process terminates if the target systems do not meet the maintenance level requirements.

For any Veritas cluster product, all nodes in the cluster must have the same operating system version and update level.

The following minimum system requirements for this release are:

- AIX 5.3 TL7 with SP2
- AIX 6.1 TL0 with SP 4

For the latest information on updates, patches, and software issues regarding this release, refer to the Symantec Technical Support website:

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

Storage Foundation Cluster File System node requirements

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

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/283990>

5.0 MP3 Veritas Storage Foundation software limitations

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

qiomkfile and odmmkfile commands must be run as a privileged user

The `qiomkfile` and `odmmkfile` commands must be run as a privileged user. This behavior has changed from the previous releases where these commands could be run by any user.

5.0 MP1 Veritas Storage Foundation software limitations

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

Advanced POWER Virtualization

The Cross-platform Data Sharing (CDS) feature is not supported for Virtual SCSI devices. By default, all Virtual SCSI devices are supported as native type devices of type `aixdisk`.

SCSI-3 Persistent Group Reservation (PGR) is not supported for Virtual SCSI devices, which prevents SFRAC functionality being applied to such devices.

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/283990>

5.0 MP3 Veritas Volume Manager software limitations

The following are additional software limitations in this 5.0 MP3 release of Veritas Volume Manager.

How DMP interacts with MPIO

MPIO is a "path suppressing driver". Therefore, when it is enabled only one path is visible to DMP. MPIO supersedes DMP, controlling failover and load balancing. Path failure is transparent to DMP. However, when DMP metanodes are created on a path, and that path fails, the path is disabled because DMP uses the controller device number for the SCSI inquiry. To avoid this, set the tunable parameter `dmp_fast_recovery` to off.

The `dmp_fast_recovery` tunable parameter controls whether DMP should attempt to obtain SCSI error information directly from the HBA interface. The default setting is on.

To set the tunable parameter `dmp_fast_recovery` to off, use the following commands:

```
# vxddmpadm settune dmp_fast_recovery=off
# vxddmpadm gettune all | grep dmp_fast
dmp_fast_recovery      off      on
```


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/283990>

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 MP1 Veritas File system software limitations

The following are additional software limitations in the 5.0 MP1 release of Veritas File System.

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, nor can that file be concurrently opened with the `o_cio` flag. A file opened with the `o_cio` flag or 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.

Veritas Storage Foundation Cluster File System software limitations

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

<http://entsupport.symantec.com/docs/283986>

5.0 MP3 Veritas Storage Foundation Cluster File System software limitations

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

Workload partition environment

The Veritas Storage Foundation Cluster File System is not supported in a workload partition (WPAR) environment.

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 4.0.

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, nor can that file be concurrently opened with the `O_CIO` flag. A file opened with the `O_CIO` flag or 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 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 for Oracle 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/283990>

5.0 MP3 Veritas Storage Foundation for Oracle software limitations

There are no additional Veritas Storage Foundation for Oracle software limitations in the MP3 release.

5.0 MP1 Veritas Storage Foundation for Oracle software limitations

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

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)

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`

Multiple archive log destinations with RAC (795617)

Multiple archive log locations are not supported in RAC configurations.

Repository hostnames are case insensitive (851129)

Since DNS hostname 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 hostname differs from another hostname only by case.

Database FlashSnap limitations (34570)

This section describes Database FlashSnap limitations.

The following are Database FlashSnap limitations:

- The Database FlashSnap feature does not support RAID-5 volumes.
- When cloning a database using Database FlashSnap, the Oracle database must have at least one mandatory archive destination, otherwise dbed_vmchecksnap results in this error message.[270905]

```
SFORA dbed_vmchecksnap ERROR V-81-5677 Could not find a
mandatory, primary and valid archive destination for
database PROD.
```

```
Please review the LOG_ARCHIVE_DEST_n parameters and check
v$archive_dest.
```

This example shows how to establish a mandatory archive destination using SQL*Plus.

```
alter system set log_archive_dest_1 =
'LOCATION=/ora_mnt/oracle/oradata/PROD/archivelogs MANDATORY
[REOPEN]' [scope=both];
```

See your Oracle documentation.

- Existing snapshot plexes created by the vxassist command are not supported. A combination of snapshot plexes created by vxassist and vxsnap is also not supported.

Veritas Storage Foundation for DB2 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/283990>

5.0 MP3 Veritas Storage Foundation for DB2 software limitations

There are no additional Veritas Storage Foundation for DB2 software limitations in the 5.0 MP3 release.

5.0 MP1 Veritas Storage Foundation for DB2 software limitations

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

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.

Cannot restore if tablespace is converted from Quick I/O to regular file after backup (25272)

If you convert a tablespace from a Quick I/O file to a regular file after backing up the database, you will not be able to restore the tablespace from that backup. For example, if you take a backup of a database that has a DMS tablespace with Quick I/O files as containers, and later convert the Quick I/O files to regular files, restoring the database from that backup will fail.

Selected utilities require setuid (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, 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`

Repository hostnames are case-insensitive (851129)

Since DNS hostname 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 hostname differs from another hostname only by case.

Veritas Volume Replicator software limitations

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

<http://entsupport.symantec.com/docs/283988>

5.0 MP3 Veritas Volume Replicator software limitations

There are no new additional Veritas Volume Replicator software limitations in this 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 recommended 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 (592349)

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

VRW limitations

The following VRW limitations exist.

Increasing font size

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

IBC messaging support

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

VCS Agents for VVR limitations

The following VCS Agents for VVR limitations exist.

Onlining and offlining the RVG resources

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

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 Volume Manager fixed issues

The following sections list Veritas Volume Manager fixed issues for Storage Foundation 5.0 Maintenance Packs. 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/283990>

Veritas Volume Manager 5.0 MP3 fixed issues

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

Table 1-1 Veritas Volume Manager fixed issues

Incident	Description
536853	If the <code>vxdiskadm</code> command is used to replace a disk in site-consistent disk group, the new disk was expected to be tagged with the same site name as the disk that is being replaced.
592685	Fixed a <code>vxddm adm getddmnode help</code> usage error.
602953	In the VEA client for Microsoft Windows systems, existing volume tags were not displayed when adding a new volume tag.
607096	Incorrect error message when importing a disk group.
779060	Enhanced VxVM to perform faster failover/plex detach for mirrored volumes.
840452	Search did not return any objects for non-Administrator users.
898339	The <code>vxconvert</code> utility failed to convert VG into CDS disk group.
934142	Fixed a failure that occurred when importing a version 60 diskgroup.
971002	Fixed a failure that occurred when using the <code>vxddm adm setattr</code> command to update the dynamic multipathing events.log.
973256	VxVM displayed the same dynamic multipathing node for two different LUNs.

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

Incident	Description
990003	<code>vxddt1 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.
1034591	Changed DMP get path state code to check only for write protection for Symmetrix.
1035027	Fixed a memory leak in the VxVM and DMP plugin of VxMS.
1046926	The same DMP node was used for two different disks when the LUN number of a DMP node was deallocated and reassigned to another physical device.
1050994	Using the <code>vxassist grow</code> command on a striped-mirror, log volume failed with the "Cannot update volume" error message.
1053089	Enhanced the <code>vxrecover</code> command to handle multiple recoveries.
1057146	DCO maps were cleared due to a failed attempt to attach a plex, which led to data corruption.
1064373	For HDS Tagmetastore array, ecopy probing lead to loss of storage access.
1065478	Fixed a panic in DMP due to stale DMP nodes.
1067870	A DCO volume moved to another disk when the volume was grown, due to <code>vxassist grow</code> not honoring the allocation specification.
1078615	Added support for initialization of LUNs greater than 2 TB.
1084199	The Oracle recovery process would loop when trying to read the next mirror.
1085134	VxVM device discovery will detect SAS disks.
1086982	Fixed a false serial split brain detection during DR.
1087183	The <code>vxconfigd</code> daemon core dumped on the primary during flashsnap.
1092390	Fixed a memory leak in <code>vxconfigd</code> .
1097834	Added a failover mode for Clarion arrays.

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

Incident	Description
1104933	Fixed an incorrect calculation of disk capacity done by the <code>vxdisksetup</code> command.
1107078	The <code>vxconfigd</code> daemon core dumped at startup in CBR.
1107156	Added instantaneous device suppression support in DMP.
1120131	Fixed a system panic in DCO detach during a link failover.
1150763	Fixed an issue in which running <code>vxdisk updateudid</code> on an imported diskgroup disk could render the diskgroup unable to be imported.
1152033	The <code>vxconfigd</code> daemon core dumped during a BCV diskgroup import.
1157509	The <code>vxconvert</code> utility failed to convert for VPATH devices.
1164610	Needed to set VxVM ownership on all paths in AIX ODM database when more paths were added after initializing disk from VxVM.
1175654	Major I/O performance hit after installing MPIO driver.
1176580	Fixed a CVM node panic that occurred when cables were unplugged from an array.
1183331	The <code>vxconfigrestore -p</code> command failed with a syntax error.
1187723	<p><code>vxconfigd</code> dumped core while setting debug level 0 using the command:</p> <pre>vxctl debug 0</pre>
1194738	Fixed a panic in <code>vxio</code> due to an incorrect calculation of number of DCOs.
1195666	<p>After removing SDD, the <code>vxconvert</code> utility failed in the analysis phase with the following error message:</p> <pre>vxicap pvptr = NULL before reading VGRA</pre>
1201755	The <code>vxrelocd</code> daemon failed to work during a simulated disk failure when using a layered volume.
1203835	VxVM was not claiming vSCSI devices properly. VxVM was adding all vSCSI luns as paths to a single LUN.

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

Incident	Description
1204200	The <code>vx dg join</code> command failed due to a transaction abort.
1204961	The <code>vx resize</code> command failed to resize with the following error message: "Cannot determine file system current size for volume."
1204981	Fixed an error that occurred with the <code>vx disk scandisk</code> command for <code>/dev/cciss/*</code> devices.
1205065	Fixed an incorrect error message for disabling the root disk or a failed disk.
1210160	Different <code>iotime</code> computation units used by <code>vx io</code> and <code>vx dmp</code> resulted in a false insane DMP node.
1210609	The entries in <code>/dev/vx/[r]dsk</code> are not cleared during reboot/system start-up after a temporary import of a disk group.
1213564	Fixed a panic that occurred during a DCO update.
1213627	<code>vx configd</code> dumped core on AIX.
1214356	The <code>vx configd</code> daemon transactions failed after a minor number conflict and temporary remapping of shared volumes.
1224737	Deported diskgroup disks from Clariion A/PF LUNs that were not on the preferred path were not accessible.
1224780	EMC SRDF devices which are configured as read-only were not displayed by the <code>vx disk list</code> command.
1228536	Running <code>vx dg flush</code> on a slave node in a CVM cluster disabled the disk.
1230360	Fixed a system panic when one of the mirrored volume's plexes went down due to LUN access failure.
1232013	The <code>vx configd</code> daemon core dumped while importing a diskgroup.
1245037	The "Operations requires transaction" error occurred during snapshot creation.
1248303	Fixed a panic in <code>vx dmp</code> that was caused by an invalid CPU table address.
1264340	The <code>vx dmpadm getattr arrayname name recoveryoption</code> command core dumped when more than one enclosure of a given type was attached.

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

Incident	Description
1265972	After a crash or reboot, a node failed to re-join the CVM cluster.
1267134	The <code>vxconfigd</code> daemon core dumped due to stack corruption.
1280387	Fixed a panic that occurred when using the <code>vxsnap admir</code> command to take a link-breakoff snapshot.
1288418	The current primary path was not getting set on a DMP node in a CVM environment when all of the ports on a switch were disabled from the host side, and then the ports were enabled.
1293922	Enhanced VxVM to include an online plex comparison utility.
1299252	VxVM commands hung because <code>vxconfigd</code> was waiting for a configuration update to finish.
1321195	Fixed a panic on CVM nodes.

Veritas Volume Manager 5.0 MP1 fixed issues

[Table 1-2](#) describes fixed issues in the Veritas Volume Manager 5.0 MP1 release.

Table 1-2 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.
603522	If a VPATH metadvice was opened with the <code>O_NDELAY</code> flag, <code>ioctl</code> s or I/O to the device failed with <code>ENXIO</code> .

Table 1-2 Veritas Veritas Volume Manager fixed issues (*continued*)

Incident	Description
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 vxddmpadm disable command was applied to a primary path on one node in a CVM cluster, the other nodes did not fail over to the secondary path.
611333, 622508	DMP could not obtain the correct serial number for a device if its LUN serial number contained a comma. This problem was seen on EMC Symmetrix arrays with more than 8096 LUNs.
614061, 614787	Adding cache volumes (used by space-optimized instant snapshots) to volume sets could cause data corruption and system panics.
617331, 631334	I/O was not restored on a path that was re-enabled after a failback or a non-disruptive upgrade (NDU) operation.
618317	A system crash could occur while bringing up cluster if I/O were performed on a unopened path.
621832	Immediately after installation, the vxesd daemon had the DVD mount point as its current working directory, which prevented the DVD from being unmounted.
625877	The error "/etc/vx/vxvm-startup: line 241: /usr/sbin/vxddladm: No such file or directory" was seen at boot time.
643089	Relayout from mirror-stripe to concat-mirror did not work for site-consistent volumes.
645749	Growing a volume by a specified amount did not work for a site-consistent volume with more than 2 disks per site.
793159	Automatic reattachment of a remote site did not work correctly.
801445	The DMP feature to detect and respond to intermittently failing paths was turned off by default in the 5.0 release, and the values of the dmp_health_time and dmp_path_age tunables were both set to 0. This feature is now enabled by default in 5.0 MP1. The default values of dmp_health_time and dmp_path_age are 60 and 300 seconds respectively.

Veritas Enterprise Administrator 5.0 MP3 fixed issues

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

Veritas Enterprise Administrator 5.0 MP1 fixed issues

[Table 1-3](#) describes fixed issues in the Veritas Enterprise Administrator 5.0 MP1 release.

Table 1-3 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 fixed issues

[Table 1-4](#) describes fixed issues in the Veritas Web GUI 5.0 MP1 release.

Table 1-4 Veritas Web GUI fixed issues

Incident	Description
564455	Removing a volume from a volume set returned a Java exception.
565072	Creating a file system on a disabled volume returned both success and failure messages.
566619	The Scan Disks By Controller View did not list the available controllers.
574410	Attempting to create a volume without an existing disk group produced a misleading error.
575262	Disabling a path to a SENA storage array produced an erroneous message.

Table 1-4 Veritas Web GUI fixed issues (*continued*)

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

Veritas File System fixed issues

The following sections list Veritas File System fixed issues for Storage Foundation 5.0 Maintenance Packs. 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/283990>

Veritas File System 5.0 MP3 fixed issues

Table 1-5 describes fixed issues in the Veritas File System 5.0 MP3 release.

Table 1-5 Veritas File System fixed issues

Incident	Description
851252	Enhanced performance of using the <code>ps -efa</code> command on a cluster node when there are several SFCFS mmap files involved.

Table 1-5 Veritas File System fixed issues (*continued*)

Incident	Description
995917	Fixed an inode table overflow that occurred when upgrading disk layout Version 5 to Version 6 with the <code>vxupgrade</code> command.
1005256	Improved the performance of <code>fsck</code> log replay.
1022679	Fixed a performance degradation that occurred during DLPAR memory movement.
1022702	Fixed a system panic that occurred during DB2 fast preallocation.
1027137	Fixed a system hang that occurred because <code>vm_releasep()</code> was being called in interrupt context from <code>pvn_write_done()</code> .
1048438	Fixed a performance degradation in <code>vx_logbuf_clean()</code> .
1051116	Improved the performance of the <code>dd</code> command.
1088153	Fixed a system panic that occurred in <code>vx-untimeout</code> .
1093636	Backed out changes to <code>vx_extfree()</code> introduced in incident 544066.
1099215	Fixed a possible infinite loop in <code>vx_write_alloc3()</code> due to a fragmented file system.
1143812	Removed the check for a Concurrent I/O license.
1167380	Fixed an incorrect use of the <code>fse_funmounted</code> flag that resulted in EIO errors.
1201080	Fixed a null pointer panic that occurred when <code>vx_populate_attrdata()</code> called <code>bcopy()</code> with a NULL destination.
1207239	Fixed a panic in <code>vx_getblk_cmn()</code> that could be caused when a Storage Checkpoint was mounted after the <code>ls -@</code> command was run against the primary file system.
1210247	Produced a quick fix for a situation in which a VxFS file system fails to mount and outputs the EAGAIN error on a system with 64 physical CPUs and SMT enabled.
1213253	The <code>fsadm -b</code> command sometimes failed to shrink a file system.
1254405	Fixed a panic in <code>vx_dnlc_purge_ip()</code> that was caused by a dereference to a NULL <code>d_flhead</code> pointer.
1255517	The <code>df</code> command had different output depending on if a file system was local or part of a cluster.

Table 1-5 Veritas File System fixed issues (*continued*)

Incident	Description
1256243	Fixed an infinite loop that occurred in <code>vx_bmap_enter()</code> during a clone removal truncate or push operation.
1275508	Fixed a performance degradation in <code>vx_logbuf_clean()</code> .
1283708	Fixed an issue in which records specified with the <code>vxdump -B</code> command were incorrectly calculated if the value was larger than 2 gigabytes.
1363642	Fixed an exception <code>vx_vnmap_min()</code> that was due to force unmounting the file system.
1364995	Reduced the delay from the time that a file is deleted to the time that the <code>du</code> and <code>bdf</code> commands report the free space from the deleted files.

Veritas File System 5.0 MP1 fixed issues

[Table 1-6](#) describes fixed issues in the Veritas File System 5.0 MP1 release.

Table 1-6 Veritas File System fixed issues

Incident	Description
596307	Dramatically decreased the time required to freeze file systems during a dynamic reconfiguration. This is a performance enhancement.
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.
770916	Fixed the problem in which <code>vxfsconvert</code> failed to convert the file system if it was invoked with <code>-s estimated_value</code> option, and the estimated value happened to be an exactly the required space for conversion and not a multiple of 8KB.
770917	Inode ownership issues detected in large directory related code paths have been fixed.

Table 1-6 Veritas File System fixed issues (*continued*)

Incident	Description
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>rwrx-rx-x</code> permissions, which is consistent with the behavior of <code>mkfs</code> .
771085	An uninitialized minor device number in <code>vxuncfg</code> command could cause a system panic. This was fixed by initializing the variable before making system calls.
771086	Fixed an <code>fsck</code> problem in which users could end up creating multiple lost+found directories when running the <code>fsck -o full</code> command and answering <code>fsck</code> questions interactively. Now, <code>fsck</code> creates only one. It also checks for multiple lost+found entries and removes duplicate directory entries.
771996	Enhanced VxFS to use less CPU when doing administrative tasks on the devices of multi-volume file systems.
772013	Enhanced the <code>fsck</code> command to enforce the lost+found file name in the root directory of the file system to be a directory file type.
777012	If the system crashed or there was a metadata I/O error, after the <code>fsadm</code> command reorganized the lost+found directory, running the <code>fsck -o full</code> command may not have been able to clean the file system with regard to names that needed to be added to the lost+found directory. The problem happened on single-volume and multi-volume file systems.
785649	A situation where <code>vxfsconvert</code> of a dusty file system loops forever in user-level code when an inode with pending truncation operation is encountered has been fixed.
793022	<code>vxfs_nattr_open()</code> API interface has been fixed to shrink files, as appropriate, when invoked with <code>O_TRUNC</code> flag.
793030	<code>vxfsutil.h</code> uses <code>struct timeval</code> in one of the function declarations, but does not include <code>time.h</code> . This causes user applications to report warnings during compilation. This issue has been fixed by including <code>time.h</code> in <code>vxfsutil.h</code> .

Table 1-6 Veritas File System fixed issues (*continued*)

Incident	Description
795073	The increased CPU utilization when writing to a file system that is almost full due to more background processing threads than are actually required being enqueued has been fixed.

Veritas Storage Foundation Cluster File System fixed issues

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

<http://entsupport.symantec.com/docs/283986>

Veritas Storage Foundation Cluster File System 5.0 MP3 fixed issues

Table 1-7 describes fixed issues in the Veritas Storage Foundation Cluster File System 5.0 MP3 release.

Table 1-7 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
1074981	Removed an unnecessary delay call that degraded SFCFS performance.
1155353	Fixed an issue in which upgrading from disk layout Version 6 to Version 7 caused <code>fsck</code> to fail on the upgraded file system.
1167284	Fixed a live lock issue in <code>vx_iupdat_msg()</code> that resulted in the ENOTOWNER error.
1178621	Fixed an SFCFS multi-transaction server inode leak.
1213262	The <code>cbuf</code> caches were consuming too much memory.
1242909	Fixed an issue in which the default ACL was sometimes not inherited correctly.
1300990	Fixed an ODM issue that produced the "Resource temporarily unavailable" error in Veritas Storage Foundation for Oracle 5.0 MP1.
1314995	Fixed a leak of attribute inodes that caused the attribute <code>ilist</code> to become very large.

Veritas Storage Foundation Cluster File System 5.0 MP1 fixed issues

[Table 1-8](#) describes fixed issues in the Veritas Storage Foundation Cluster File System 5.0 MP1 release.

Table 1-8 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.
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.

Table 1-8 Veritas Storage Foundation Cluster File System fixed issues
(continued)

Incident	Description
771999	A node in a cluster was hanging under heavy CFS activity in the cluster. When all the worker threads servicing callbacks from a queue were blocked, the callbacks that were still lying in the same queue faced a deadlock. The problem has been resolved by queuing the callbacks that need to be processed to unblock other callbacks into a different queue, and processing them at a higher priority.
775611	Upon encountering an I/O error, the file system is correctly disabled but an internal inconsistency check fails. The check failure in this instance is quite benign and cannot cause customer data corruption since the file system is safely disabled upon encountering an I/O error of this nature. To further diagnose this condition, a new consistency check has been added.
782024	Under heavy inode cache stress, file system recovery on a cluster file system was hanging after an I/O error. The problem has been resolved by ensuring that the code paths that attempt to reuse an inode do not wait for a dirty/bad/inactive inode to become free if recovery initiated by a disabled file system is currently in progress. If there are no inodes available for re-use, a new inode will be allocated to ensure that the recovery proceeds to completion.
793030	The vxfsutil.h was using 'struct timeval' in one of the function declarations without including time.h resulting in a warning message during compilation. The problem has been resolved by including time.h in vxfsutil.h.

Veritas Volume Replicator fixed issues

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

<http://entsupport.symantec.com/docs/283988>

Veritas Volume Replicator 5.0 MP3 fixed issues

[Table 1-9](#) describes fixed issues in the Veritas Volume Replicator 5.0 MP3 release.

Table 1-9 Veritas Volume Replicator Fixed Issues

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

Veritas Volume Replicator vradmin 5.0 MP1 fixed issues

[Table 1-10](#) describes fixed issues in the Veritas Volume Replicator vradmin 5.0 MP1 release.

Table 1-10 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 1-10 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 fixed issues

[Table 1-11](#) describes fixed issues in the Veritas Volume Replicator Web GUI 5.0 MP1 release.

Table 1-11 Veritas Volume Replicator Web GUI fixed issues

Incident	Description
576729	When a user without sufficient privileges tried to use the Create Primary wizard, the operation resulted in an incorrect error message.
611792	When a user tried to create a Primary without specifying the RVG name, the operation resulted in an RVG with an invalid configuration being created.
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.
768497	In the Create Primary wizard, sometimes non-free volumes were displayed in the volume list as free volumes.

Table 1-11 Veritas Volume Replicator Web GUI fixed issues (*continued*)

Incident	Description
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 fixed issues

[Table 1-12](#) describes fixed issues in the Veritas Volume Replicator VEA 5.0 MP1 release.

Table 1-12 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.
776622	The Creating a Primary wizard failed when the list of data volumes was too long.

Table 1-12 Veritas Volume Replicator VEA fixed issues (*continued*)

Incident	Description
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 sections list Veritas Storage Foundation for Databases fixed issues for Storage Foundation 5.0 Maintenance Packs. 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/283990>

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.

Table 1-13 Veritas Storage Foundation for Oracle fixed issues

Incident	Description
900779	If an Oracle query is implemented during the time period when the database is suspended, it no longer causes Oracle to hang and crash. [Titan case: 320-018-260]
1111678	In the 5.0MP1 release, <code>VRTSdbcom</code> and <code>VRTSdbed</code> patches could not be installed on AIX 5300-05-CSP, which is a OS security, trusted computing base (TCB) enabled machine. [Titan case: 230-407-473]

Veritas Storage Foundation for Oracle 5.0 MP1 fixed issues

[Table 1-14](#) describes fixed issues in the Veritas Storage Foundation for Oracle 5.0 MP1 release.

Table 1-14 Veritas Storage Foundation for Oracle fixed issues

Incident	Description
567342	An unmounted checkpoint clone database reappears in the Java GUI tree after rescanning.

Table 1-14 Veritas Storage Foundation for Oracle fixed issues (*continued*)

Incident	Description
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="534 435 1022 487">([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="534 618 1150 670">Error V-39-53246-8 Get EntryPoint failed. Please check the manifest related information</pre>
600431	<p>Storage Checkpoint operations are currently not supported for databases cloned with Database FlashSnap.</p>
608697	<p>You can now refresh the View Statistics wizard in the Firefox browser.</p>
613681	<p>The sfua_db_config -o startdb command does not print a message to indicate whether the SFDB repository database was successfully started or not.</p>
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="565 1216 1157 1269"># /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="565 1468 1180 1520"># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh \ -a agent_name -s 0</pre> <p>where agent_name is the agent name output from the -l 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.

Veritas Storage Foundation for DB2 5.0 MP1 fixed issues

[Table 1-15](#) describes fixed issues in the Veritas Storage Foundation for DB2 5.0 MP1 release.

Table 1-15 Veritas Storage Foundation for DB2 fixed issues

Incident	Description
567342	An unmounted checkpoint clone database reappears in the Java GUI tree after rescanning.
582069	SFDB commands no longer fail when executed within a locale that differs from the locale in use when the SFDB server was started.
582416	Clicking the Help button on a GUI wizard produces the following error message: Error V-39-53246-8 Get EntryPoint failed. Please check the manifest related information
600490, 600492	The db2ed_clonedb and db2ed_vmclonedb commands now support automatic storage databases.
608697	You can now refresh the View Statistics wizard in the Firefox browser.
613681	The sfua_db_config -o startdb command does not print a message to indicate whether the SFDB repository database was successfully started or not.

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/283990>

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.

Required operating system library for the event source daemon (vxesd)

On AIX, the Storage Foundation 5.0MP3 event source daemon process (vxesd) depends on the operating system library `/usr/lib/libHBAAPI.a`. The process fails to start if the library is not installed.

Workaround

If the library is not installed, install the AIX fileset

`devices.common.IBM.fc.hba-api` before you install Storage Foundation 5.0MP3.

Using ODM with Storage Foundation or Storage Foundation Cluster File System

Starting with 5.0MP3, 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.

In order to use this functionality, you must install a hotfix for the Veritas licensing package. See the late breaking news Tech Note for the location of the hotfix.

To use this functionality, you may also need to manually install the required packages and patches. Refer to the chart below for your product and install scenario.

Table 1-16 Installation and upgrade scenarios

Product	Installation/Upgrade scenario	Instructions
Storage Foundation Cluster File System	Upgrading SFCFS from 5.0 or a previous 5.0MP	Install ODM and GMS before upgrading SFCFS to 5.0MP3. See “Installing ODM before upgrading from Storage Foundation or Storage Foundation Cluster File System 5.x” on page 53.

Table 1-16 Installation and upgrade scenarios (*continued*)

Product	Installation/Upgrade scenario	Instructions
	Upgrading SFCFS from 4.x version.	The product installer installs ODM and GMS. No manual steps are required.
	Installing SFCFS for the first time.	The product installer installs ODM and GMS. No manual steps are required.
Storage Foundation	Upgrading SF from 5.0 or a previous 5.0MP	Install ODM before upgrading Storage Foundation to 5.0MP3. See “Installing ODM before upgrading from Storage Foundation or Storage Foundation Cluster File System 5.x” on page 53.
	Upgrading SF from 4.x version.	Install Storage Foundation 5.0MP3, and then manually install ODM. See “Installing ODM after installing Storage Foundation 5.0MP3 or Storage Foundation Cluster File System 5.0MP3” on page 54.
	Installing SF for the first time.	Install Storage Foundation 5.0MP3, and then manually install ODM. See “Installing ODM after installing Storage Foundation 5.0MP3 or Storage Foundation Cluster File System 5.0MP3” on page 54.

Installing ODM before upgrading from Storage Foundation or Storage Foundation Cluster File System 5.x

Starting with 5.0MP3, ODM is licensed with the Storage Foundation license key. Use the procedure in this section if you are upgrading from Storage Foundation or Storage Foundation Cluster File System 5.0 or a previous 5.0MP.

To use ODM with Storage Foundation, install ODM before upgrading Storage Foundation to 5.0MP3.

To use ODM with Storage Foundation Cluster File System, install ODM and also GMS before upgrading Storage Foundation Cluster File System to 5.0MP3.

To install ODM when you are upgrading from 5.0 or a previous 5.0MP

- 1 Find the 5.0 ODM package, `VRTSodm`, located in the `storage_foundation_for_oracle/pkgs` directory in a compressed format.

If you are installing Storage Foundation Cluster File System, you also need the `GMS` package, `VRTSgms`, located in the `sfrac/pkgs` directory.

- 2 Copy the ODM package to a temporary directory. If you are installing Storage Foundation Cluster File System, also copy the `GMS` package.

```
cp storage_foundation_for_oracle/pkgs/VRTSodm.bff.gz /tmp
cp sfrac/pkgs/VRTSgms.bff.gz /tmp
```

- 3 Unzip the packages in the temporary directory.

```
cd /tmp
gunzip *.gz
```

- 4 Install the `VRTSodm` package and, if applicable, the `VRTSgms` package.

```
installp -acd VRTSodm.bff all
installp -acd VRTSgms.bff all
```

- 5 After the 5.0 packages are installed, upgrade to Storage Foundation or Storage Foundation Cluster File System to 5.0MP3, using the `installmp` script.

The `installmp` script also upgrades the `VRTSodm` package and, if applicable, the `VRTSgms` package to the 5.0MP3 level.

See the *Storage Foundation Installation Guide* for more information about the `installmp` script.

- 6 After installing 5.0MP3, configure ODM.

For information about configuring ODM, see the *Storage Foundation for Oracle Administrator's Guide*.

Installing ODM after installing Storage Foundation 5.0MP3 or Storage Foundation Cluster File System 5.0MP3

Starting with 5.0MP3, ODM is licensed with the Storage Foundation license key. To use ODM with Storage Foundation 5.0MP3, you must install ODM separately. Use the procedure in this section if Storage Foundation 5.0MP3 or Storage Foundation Cluster File System 5.0MP3 is already installed.

To install ODM if Storage Foundation 5.0MP3 is already installed

- 1 Find the 5.0 ODM package, `VRTSodm`, located in the `storage_foundation_for_oracle/pkgs` directory in a compressed format.
If you are installing Storage Foundation Cluster File System, you also need the GMS package, `VRTSgms`, located in the `sfrac/pkgs` directory.

- 2 Copy the ODM package to a temporary directory. If you are installing Storage Foundation Cluster File System, also copy the GMS package.

```
# cp storage_foundation_for_oracle/pkgs/VRTSodm.bff.gz /tmp
# cp sfrac/pkgs/VRTSgms.bff.gz /tmp
```

- 3 Unzip the packages in the temporary directory.

```
# cd /tmp
# gunzip *.gz
```

- 4 Copy the 5.0 MP3 patch for the ODM package to a temporary directory. If you are installing CFS, you also need the GMS patch.

```
# cp storage_foundation_for_oracle/patches/VRTSodm.bff.gz /tmp
# cp sfrac/patches/VRTSgms.bff.gz /tmp
```

- 5 Unzip the patches.

```
# cd /tmp
# gunzip *.gz
```

- 6 Install the ODM patch and, if required, the GMS patch.

```
# installp -acd VRTSodm.bff all
# installp -acd VRTSgms.bff all
```

- 7 After installing 5.0MP3, configure ODM.

For information about configuring ODM, see the *Storage Foundation for Oracle Administrator's Guide*.

- 8 If you reboot the machine during or after the installation, GMS is automatically loaded. If you do not reboot, you must manually load GMS.

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

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

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

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

Workaround

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

Product installer does not support upgrading to a different Storage Foundation product and upgrading the product version in a single process

The Veritas product installer does not support upgrading to a different product in the Storage Foundation product line and upgrading the product version in a single process. For example, you cannot upgrade directly from Storage Foundation 4.0 to Storage Foundation Cluster File System 5.0 MP3. First, you must upgrade the Storage Foundation version from 4.0 to 5.0 MP3. Then, upgrade the product from Storage Foundation 5.0 MP3 to Storage Foundation Cluster File System 5.0 MP3.

Installing mksysb image on a different host includes stale configuration information (154124)

On AIX, you can use the `mksysb` utility to back up the system image. You can install the image on another host. For example, you can use the utility to set up a disaster recovery site. If you install the `mksysb` image of one host on a different host, some configuration information about the original host remains, in host-specific files.

Workaround

You can use the procedure in this section to clean stale configuration information and files.

This document uses the following terms:

- Source host: the host from which the image is taken.
- Target host: the host on which the image is installed

To clean the configuration records on the target machine

- 1 Reboot the target machine.
- 2 Stop Veritas Volume Manager:

```
# vxctl stop
```

If necessary, this can be forced with the following command:

```
# vxctl -k stop
```

- 3 If any of the following files or directories are present, delete them.

The following files contain stale information:

```
/etc/vx/darecs  
/etc/vx/disk.info  
/etc/vx/array.info  
/etc/vx/jbod.info
```

The following directory contains stale information:

```
/etc/vx/tempdb
```

- 4 Issue the following command:

```
# vxconfigd -kr reset
```

This command recreates the files deleted in step 3 , if required.

- 5 Re-initialize the VxVM configuration:

```
# vxctl init  
# vxctl initdmp
```

- 6 Install a new license key if prompted:

```
# vxlicinst -k key
```

7 Check the default disk layout attributes. The `/etc/default/vxdisk` file stores information about the default format (aixdisk or cdsdisk) for initializing disks. Use the Default disk layout attributes option from the `vxdiskadm` menu to change the default format, if required.

8 Check DMP I/O policies with the following command:

```
# vxdmpadm getattr enclosure enc_name iopolicy
```

Change the I/O policies if required.

```
# vxdmpadm setattr enclosure enc_name iopolicy=iopolicy
```

9 Check and set the VVR tunables with the `vxtune` command, if required.

10 Update the following files, which contain stale CVM-related configuration information. Otherwise, CVM is not able to start.

■ `/etc/llthosts`

The `llthosts` file contains node id information of all the nodes in the CVM as follows:

```
0 node1
1 node2
```

Update the `llthosts` file with the correct nodes in the CVM and their respective node ids.

■ `/etc/llttab`

The `llttab` file contains information like node name, cluster id and network information as follows:

```
set-verbose 1
set-node nodename
set-cluster cluster-id

link en0 /dev/dlpi/en:0 - ether - -
link en1 /dev/dlpi/en:1 - ether - -
start
```

Update the node name and the cluster id correctly. Update network devices, represented by the link field, correctly.

Veritas Storage Foundation 5.0 MP1 known issues

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

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/283990>

<http://entsupport.symantec.com/docs/289317>

Veritas Storage Foundation Basic 5.0 MP3 known issues

There are no additional known issues in this 5.0 MP3 release of Veritas Storage Foundation.

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/283990>

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
```

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

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

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

Workaround:

Tag disks manually; do not use autotagging.

Upgrading from 5.0 MP1 Update 1 on AIX 6.1 with SANBoot to 5.0 MP3 may fail (1302279)

When upgrading from 5.0 MP1 Update 1 on AIX 6.1 with SANBoot to 5.0 MP3 `vxdmproot uninstall` command may fail with the following error message:

```
# vxdmproot uninstall
"VxVM ERROR V-5-1-0 bosboot on all the subpaths bosboot on all the subpaths
bosboot on all the subpaths bosboot on all the subpaths failed
please run /usr/sbin/bosboot -ad /dev/hdisk0 manually to build NVRAM
VxVM ERROR V-5-1-0 Failed to build boot partition"
```

Workaround

To resolve the vxdmproot uninstall failure

- 1 Create a link to DMP metanode corresponding to the boot device in `/dev` such as `asIBM_VSCSI0_0` is the DMP metanode. For example,

```
# ln /dev/vx/rdmp/IBM_VSCSI0_0 /dev/tmp_dmp_boot
```

- 2 Run `bosboot`, enter:

```
# bosboot -a
```

- 3 Upgrade to 5.0 MP3.

See the *Veritas Storage Foundation Installation Guide* for upgrading to 5.0 MP3.

- 4 Install `vxdmproot`, enter:

```
# vxdmproot install
```

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 vxdm 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.

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.

Using the `vxconvert` command on Big VolumeGroups results in data loss on AIX 6.1 platforms (1230431)

After the conversion of Big VolumeGroup using the `vxconvert` command, you cannot mount the file system on the resultant VxVM volume. There is no workaround.

Manually installing the VRTSvxvm patch requires a reboot (1180992)

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

Plex disabled in DS4000 disk array (924680)

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

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 MP1 known issues

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

DMP issues of Veritas Volume Manager

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

Limitations of DMP support for LVM root disks

The following are limitations of DMP support for LVM root disks.

DMP support for LVM root disks has the following limitations:

- DMP can be enabled only on the boot disk. The root volumes on disks other than boot disks cannot be controlled by DMP. In addition, the mirror disk cannot be controlled by DMP. However, when you reboot on the mirrored disk, it will come up with DMP enabled.
- The `vxdmproot add` and `vxdmproot remove` commands cannot be applied on non-bootable boot disks. In other words, these commands can only be applied to the mirror rootdisks.
- DMP may not operate correctly if Dynamic Reconfiguration (DR) or similar operations are performed that shuffle device numbers or device names. If such a reconfiguration occurs, use the following commands to reconfigure DMP before the system is next rebooted.

```
# cfgmgr
# vxdmproot install
```

PowerPath disks and vxconvert (857504)

If a PowerPath disk is formatted as a simple disk (a foreign device), then LVM to VxVM conversion will fail using the `vxconvert` command. The format of the disk can be seen using the `vxdisk list` command. Also, make sure there is no `hdiskpower` disk entry in the file `/etc/vx/darecs`.

The problem of having an entry in the `/etc/vx/darecs` file is because the PowerPath disks were configured as foreign disks in Storage Foundation 4.0, and this was not changed after subsequent upgrades.

Handling intermittently failing paths

The `dmp_health_time` and `dmp_path_age` tunables control how DMP handles intermittently failing paths. The default values of `dmp_health_time` and `dmp_path_age` are 60 and 300 seconds respectively. The value of `dmp_health_time` represents the minimum time in seconds for which a path must stay healthy. If

a path changes state between enabled and disabled on a shorter time scale than this, DMP marks the path as intermittently failing and disables I/O on the path. I/O is not re-enabled on an intermittently failing path until `dmp_path_age` seconds have elapsed without further outage.

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

Cluster issues of Veritas Volume Manager

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

Handling intermittently failing paths in a Campus Cluster

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

Domain controller mode in CVM clusters (603213)

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

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

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

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

Veritas Enterprise Administrator known issues

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

Veritas Web GUI known issues

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

Error when creating a volume set (615960)

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

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

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/283990>

Veritas File System 5.0 MP3 known issues

There are no additional known issues in this 5.0 MP3 release of Veritas Storage Foundation.

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_malloc=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/283986>

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.

Upgrade warning messages (1399374)

If you are upgrading from SF CFS 4.0 to 5.0MP3, the following warning messages will be displayed:

```
SF CFS configuration could not be completed during installation
/VRTS/bin/cfscluster config -s> must be run after the system reboot
```

These messages may be disregarded. Do not run `cfscluster config` if you are performing an upgrade.

CFSMount resource may fault

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

Workaround

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

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

Veritas Storage Foundation Cluster File System 5.0 MP1 known issues

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

Limited Support for Advanced Power Virtualization

Storage Foundation Cluster File System (SFCFS) support for Advanced Power Virtualization (APV) of AIX 5.3 is limited to micro-partitioning of CPUs and Virtual Ethernet. As applications' workload vary across different systems, the workload imposed on SFCFS, network interfaces, the network interconnect, and CPUs must be appropriate for the deployment of APV solutions within a virtual partition installed with SFCFS. Sharing of ethernet devices or CPUs across multiple partitions may be inappropriate with a high level of workload, which can result in degradation of performance or false failover depending on the degree of virtualization implemented within a partition. Partitioning of I/O devices is not currently supported with SFCFS.

vxvol command fails

Sometimes when the cluster starts with multiple nodes, some shared volumes might remain in DISABLED state due to the failure of the `vxvol` command. If this occurs, the following message appears on the console:

```
VxVM vxvol ERROR V-5-1-10128 Error in cluster processing
```

Workaround

Use the following procedure to recover.

To recover from this situation

- 1 For all the CVMVolDg and CFSSMount resources that are faulted, run the following command on one of the nodes:

```
# hares -clear resourcename -sys system01  
# hares -clear resourcename -sys system02  
# hares -clear resourcename -sys system03  
# hares -clear resourcename -sys system04
```

- 2 Start all volumes that are disabled by entering the following command on the CVM Master:

```
# vxvol -g diskgroup startall
```

- 3 For all the CVMVolDg resources whose faults were cleared in step 1, run the following command on one of the nodes:

```
# hares -online resourcename -sys system01  
# hares -online resourcename -sys system02  
# hares -online resourcename -sys system03  
# hares -online resourcename -sys system04
```

- 4 Mount the file systems by running the following command on any one node of the cluster:

```
# cfsmount filesystem
```

Oracle-AIX (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 Follow steps 1 through 2 in the uninstallation procedure in "Removing Storage Foundation packages" in the *Veritas Storage Foundation Installation Guide*.
- 3 Remove the Veritas Storage Foundation for Oracle packages using the `installp -u` command.

```
# installp -u VRTSorgui VRTSdbed VRTSdbdoc VRTSdbcom
```

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-AIX (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 Follow steps 1 through 2 in the uninstallation procedure in "Removing Storage Foundation packages" in the *Veritas Storage Foundation Installation Guide*.
- 3 Remove the Veritas Storage Foundation for DB2 packages using the `installp -u` command.

```
# installp -u VRTSd2gui VRTSdb2ed VRTSdbdoc VRTSdbcom
```

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 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/283988>

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:

```
# vxpdg upgrade diskgroup
```

Veritas Volume Replicator 5.0 MP1 known issues

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

Issue with VVR VEA in the Japanese locale (616709)

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

The incomplete text should read as follows:

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

Veritas 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/283990>

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.

dbed_ckptrollback may fail with error message on AIX 6.1 (1396168)

If you are running AIX 6.1 on your machine, and you execute the `dbed_ckptrollback` command, the command may fail and display the following error message:

```
ERROR V-81-4804 Thread 4 pwrite error.  
SFORA rollback ERROR V-81-4818  
Rollback failed for file /oracle/oradata/pat111/bmf.dbf.
```

Workaround

You must use the `dbed_ckptmount` command to mount the checkpoint. Then you must use the `cp` command to copy the files which failed to rollback.

Error message not displayed if sfua_db_config fails (1139077)

If the `sfua_db_config` utility fails for any reason, an error message is not displayed in the output. The output may display an empty NIC array.

Workaround

There is no workaround for this issue.

dbed_analyzer command displays incomplete information for multiple tables with same name (1121139)

When the database has more than one table with same name in different schema, the `dbed_analyzer` displays incomplete information for the tables.

Workaround

There is no workaround for this issue.

After using `dbed_ckptpolicy`, dbf files display in addition to the checkpoint policy (1186193)

After creating Checkpoint policies using the `dbed_ckptpolicy` command, when you display the created policies using the following command,

```
# dbed_ckptpolicy -S SID -n -o display
```

dbf files are displayed along with the created Checkpoint policies.

Workaround

There is no workaround for this issue.

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_xxxxxxxx ...
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.

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 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.

File fragmentation check in the `qio_convertdbfiles` command may report errors (819430)

The file fragmentation check in `qio_convertdbfiles` may report errors when run on multi-volume file systems. These errors are harmless and may be safely ignored. This issue also causes the `dbed_checkconfig` command to fail with an error if run on a database which uses one or more multi-volume file systems. The method used to determine fragmentation in `qio_convertdbfiles` has been deprecated. The preferred way to check and resolve file or file system fragmentation is through the use of the `fsadm` tool.

Refer to the *Veritas File System Administrator's Guide* for more information on using `fsadm` to display and resolve file system fragmentation.

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 Follow steps 1 through 2 in the uninstallation procedure in "Removing Storage Foundation packages" in the *Veritas Storage Foundation Installation Guide*.
- 3 Remove the Veritas Storage Foundation for Oracle packages using the `installp -u` command.

```
# installp -u VRTSorgui VRTSdbed VRTSdbdoc VRTSdbcom
```

Error message displayed after successfully creating a new repository (830629)

The following error message is displayed after using `sfua_db_config` to create a new repository:

```
Ping server failed -- Database server not found.
```

You can safely ignore the message.

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 snapplan 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/283990>

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.

Error message not displayed if sfua_db_config fails (1139077)

If the `sfua_db_config` utility fails for any reason, an error message is not displayed in the output. The output may display an empty NIC array.

Workaround

There is no workaround for this issue.

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 Follow steps 1 through 2 in the uninstallation procedure in "Removing Storage Foundation packages" in the *Veritas Storage Foundation Installation Guide*.
- 3 Remove the Veritas Storage Foundation for DB2 packages using the `installp -u` command.

```
# installp -u VRTSd2gui VRTSdb2ed VRTSdbdoc VRTSdbcom
```

db2ed_vmsnap gives error: unable to duplicate owner and mode after move (841766)

This move error message can be ignored. /tmp has gid bit set so the snapshot completed successfully. This warning will be suppressed or removed in a later release.

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 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

The 5.0MP3 release does not support AIX 5.2.

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 documents are relevant component product release notes:

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

Veritas Storage Foundation guides

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

[Table 1-17](#) describes the guides in the Veritas Storage Foundation documentation set.

Table 1-17 Guides in Veritas Storage Foundation documentation set

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

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

Guide Title	Filename
<i>Veritas Enterprise Administrator User's Guide</i>	vea_users.pdf
<i>Veritas File System Administrator's Guide</i>	vxfs_admin.pdf
<i>Veritas File System Programmer's Reference Guide</i>	vxfs_ref.pdf
<i>Veritas Volume Manager Administrator's Guide</i>	vxvm_admin.pdf
<i>Veritas Volume Manager Troubleshooting Guide</i>	vxvm_tshoot.pdf

Veritas Storage Foundation Cluster File System documentation

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

[Table 1-18](#) describes the Veritas Storage Foundation Cluster File System (CFS) documentation set.

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

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

Veritas Cluster Server documentation

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

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

Table 1-19 Guides in Veritas Cluster Server documentation set

Guide Title	Filename
<i>Veritas Cluster Server Release Notes</i>	vcs_notes.pdf
<i>Veritas Cluster Server Installation Guide</i>	vcs_install.pdf
<i>Veritas Cluster Server User's Guide</i>	vcs_users.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf

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

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

Veritas Storage Foundation for Oracle RAC documentation

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

[Table 1-21](#) describes the Storage Foundation for Oracle RAC documentation set.

Table 1-21 Guides in Storage Foundation for Oracle RAC documentation set

Guide Title	Filename
<i>Veritas Storage Foundation™ for Oracle RAC Installation and Configuration Guide</i>	sfrac_install.pdf
<i>Veritas Storage Foundation™ for Oracle RAC Release Notes</i>	sfrac_notes.pdf
<i>Veritas Storage Foundation™ for Oracle RAC Administrator's Guide</i>	sfrac_admin.pdf

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 1-22](#) indicates the PDF and document changes in this release.

Table 1-22 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 manual page errata (1361604)

For the `dbdst_report` manual page, one of the usage examples is incorrectly 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 manual pages errata (1121091)

You cannot access the correct manual 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
```

