

Veritas Storage Foundation™ Release Notes

Solaris

5.0 Maintenance Pack 1



Veritas Storage Foundation Release Notes

Copyright © 2007 Symantec Corporation. All rights reserved.

Storage Foundation 5.0 Maintenance Pack 1

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

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THIS 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” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

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/283886>

Solaris is a trademark of Sun Microsystems, Inc.

Licensing and registration

Veritas Storage Foundation is a licensed product. See the *Veritas Storage Foundation Installation Guide* for license installation instructions.

Technical support

For technical assistance, visit

http://www.symantec.com/enterprise/support/assistance_care.jsp and select Product Support. Select a product and use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Introduction	7
New features	9
Veritas File System	9
System requirements	9
Solaris operating system requirements	9
Component product release notes	10
Installing for the first time	11
Installing Storage Foundation 5.0 MP1	11
Installing Storage Foundation on a cluster	12
Upgrading Storage Foundation on a standalone system	16
Changing permissions for Storage Foundation for Databases	20
Enabling Storage Mapping in Veritas Storage Foundation for Sybase	21
Removing patches	22
Removing an installed patch	22
Removing an installed patch from a cluster	22
Software limitations	23
Veritas File System software limitations	23
Veritas Storage Foundation for Databases software limitations	23
No longer supported	27
Fixed issues	28
Veritas Volume Manager fixed issues	28
Veritas File System fixed issues	31
Veritas Volume Replicator fixed issues	33
Veritas Storage Foundation for Databases fixed issues	35
Known issues	41
Veritas Volume Manager known issues	41
Veritas File System known issues	47
Veritas Volume Replicator known issues	47
Veritas Storage Foundation for Databases known issues	48
Documentation errata	55
Veritas Storage Foundation Installation Guide errata	55
Web GUI help errata	57
Manual page errata	58
Veritas Volume Manager Administrator's Guide errata	58
Veritas Storage Foundation for Oracle Administrator's Guide errata	59
Veritas Storage Foundation for Oracle Graphical User Interface Guide errata	61

Veritas Storage Foundation for DB2 Administrator's Guide errata 62
Veritas Storage Foundation for DB2 Graphical User Interface Guide errata
63

Veritas Storage Foundation Release Notes

Introduction

This document provides release information about the products in the Veritas Storage Foundation 5.0 Maintenance Pack 1 (MP1) Solaris (SPARC Platform) 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 Storage Foundation for Sybase (Standard, Enterprise, and HA Editions)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)

Note: 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*.

Note: 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/281987>

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*.

New features

The following new features have been incorporated into this release of Veritas Storage Foundation.

Veritas File System

The following new features have been incorporated into Veritas File System.

Directory name lookup cache support for long filenames

This enhancement allows the VxFS directory name lookup cache (DNLC) to cache filenames with a length up to the maximum filename length supported by the operating system. Previously, VxFS DNLC only cached filenames that were 32 bytes or shorter.

System requirements

This section describes the system requirements for this release.

Solaris operating system requirements

The Veritas Storage Foundation 5.0 MP1 product line operates on the following Solaris operating systems:

- Solaris 8 (SPARC Platform 32-bit and 64-bit)
- Solaris 9 (SPARC Platform 32-bit and 64-bit)
- Solaris 10 (SPARC Platform 64-bit)

Storage Foundation verifies that the target system is running a required version of the Solaris operating system. Storage Foundation installation will fail if the product discovers an incorrect Solaris version.

Component product release notes

Release notes for component products in all versions of the Veritas Storage Foundation are located under the `storage_foundation/release_notes` directory of the Veritas Storage Foundation disc. It is important that you read the relevant component product release notes before installing any version of Veritas Storage Foundation:

Veritas Storage Foundation Release Notes (sf_notes.pdf)

Veritas Cluster Server Release Notes (vcs_notes.pdf)

Because product release notes are not installed by any packages, Symantec recommends that you copy them to the `/opt/VRTSproduct_name/doc` directory after the product installation so that they are available for future reference.

Installing for the first time

If you are installing Veritas Storage Foundation for the first time you must first install version 5.0 and then upgrade to MP1.

Installing Storage Foundation 5.0 MP1

If you install Veritas Storage Foundation using the installation script, the patches are added for you. However, if you choose to install the product using `pkgadd`, you must manually add the patches after installing.

Note: If you manually install a language pack, you must reboot after installing the pack:

```
# /usr/sbin/shutdown -y -i6 -g0
```

You must have superuser (`root`) privileges to install the Veritas software.

Note: If using JumpStart or Live Update to install this MP1 release of Veritas Storage Foundation (including Veritas Storage Foundation for RAC and Storage Foundation Management System version 1.0 MP1), you must first execute the `dbed_patch_50ga` file located in the `/mount_point/product-directory/scripts` directory before starting JumpStart or Live Update.

Also run the `dbed_patch_50ga` file if you are manually adding patches with `patchadd` to any of the the following packages: `VRTSdbcom`, `VRTSdbed`, `VRTSdb2ed`, `VRTSorgui`, `VRTSd2gui`, `VRTSmapro`, `VRTSorweb`, or `VRTSd2web`.
[783875, 800190]

The following sections describe how to upgrade a cluster and a standalone system:

- [“Installing Storage Foundation on a cluster”](#) on page 12.
- [“Upgrading Storage Foundation on a standalone system”](#) on page 16.

Installing Storage Foundation on a cluster

Caution: Existing data could be destroyed on any disks that are touched by upgrading the operating system. While upgrading, do not reconfigure any disks other than the root disk. To ensure the integrity of your data, it is recommended that you back it up before starting the upgrade.

To upgrade to Storage Foundation 5.0 MP1 on a cluster

- 1 Log in as superuser.
- 2 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 3 Use `hastop` command to stop the cluster:

```
# /opt/VRTSvcs/bin/hastop -all
```

Note: Do not use the `-force` option when executing `hastop`. This will leave all service groups online and shut down Veritas Cluster Server (VCS), causing undesired results while upgrading the packages.

All nodes in a cluster must currently be running the Veritas Storage Foundation 5.0 software, and the correct licenses must be present on these system.

- 4 If you are upgrading the HA version of the Veritas Storage Foundation 5.0 software, follow the guidelines given in the *Veritas Cluster Server 5.0 Installation Guide* and *Veritas Cluster Server 5.0 Release Notes* for information on preserving your VCS configuration across the upgrade procedure.
- 5 Perform any necessary preinstallation checks and configuration. See the *Veritas Storage Foundation 5.0 Installation Guide* for more information.

- 6 For each node in the cluster, check if its root disk is under VxVM control:

```
# df -k /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- a Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- b Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices:

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

- 7 Check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -F vxfs
```

- 8 Unmount all Storage Checkpoints and VxFS file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

- 9 Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean  
flags 0 mod 0 clean clean_value
```

A `clean_value` value of `0x5a` indicates the file system is clean, `0x3c` indicates the file system is dirty, and `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.

- a If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem  
# mount -F vxfs filesystem mountpoint  
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system  
file system still in progress.
```

- b If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
 - c Repeat [step 9](#) to verify that the unclean file system is now clean.
- 10 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps. Otherwise, proceed to [step 11](#).
- a Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
 - b Use the `vrxvg stop` command to stop each RVG individually:

```
# vxrvvg -g diskgroup stop rvg_name
```
 - c On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

Caution: To avoid data corruption, do not proceed until all RLINKs are up-to-date.

- d If you are going to use cross-version replication between release 5.0 and release 4.0, perform the following step on the host that is running version 5.0 to make the tunable value persistent. Otherwise, proceed to [step 11](#).

```
# vxvoltune vol_vvr_use_host_byte_order 1
```
- 11 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
- 12 Stop all VxVM volumes by entering the following command for each disk group:

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

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

```
# vxprint -Aht -e v_open
```
- 13 Comment out any mount points in the `/etc/vfstab` file for VxFS file systems, or for any file systems that are configured on VxVM volumes.

- 14 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```
 - 15 Insert the disc containing the Veritas software into the DVD-ROM drive, and mount the disc on a suitable mount point, for example:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```
 - 16 Move to the top-level directory on the DVD:

```
# cd /mnt/cdrom
```
 - 17 To upgrade the Storage Foundation software, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:
 - To install on the local system, enter the following command:

```
# ./installmp
```
 - To install on more than one system using secure shell (SSH) utilities, enter the following command:

```
# ./installmp system_name1 system_name2 ...
```
 - To install on more than one system using remote shell (RSH) utilities, enter the following command:

```
# ./installmp system_name1 system_name2 ... -rsh
```
 - 18 After the initial system checks have completed successfully, press **Enter** to start the requirements checks for the upgrade.
 - 19 After the requirement checks have completed successfully, press **Enter** to begin upgrading the packages.
-
- Note:** As you are upgrading multiple systems in a cluster, select to upgrade the systems simultaneously.
-
- 20 After the upgrade of the packages is complete, use the following command to shut down the cluster nodes:

```
# shutdown -y -i6 -g0
```
 - 21 Reboot the cluster nodes, and reinstate any missing mount points in the `/etc/vfstab` file.
 - 22 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```
 - 23 Restart each RVG (if any were stopped in [step 10](#)):

```
# vxrvg -g diskgroup start rvg_name
```

24 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem  
# mount /checkpoint_name
```

25 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

26 If you set the value of the `vol_vvr_use_host_byte_order` tunable to 1 in [step d](#) of [step 10](#), reboot the system.

There are several optional configuration steps that you may perform:

- If you want to use features of Veritas Storage Foundation 5.0 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.
- To encapsulate and mirror the boot disk, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager 5.0 Administrator’s Guide*.
- Restore any VCS configuration files as described in the *Veritas Cluster Server 5.0 Installation Guide* and *Veritas Cluster Server 5.0 Release Notes*.

Upgrading Storage Foundation on a standalone system

Caution: Existing data could be destroyed on any disks that are touched by upgrading the operating system. While upgrading, do not reconfigure any disks other than the root disk. To ensure the integrity of your data, it is recommended that you back it up before starting the upgrade.

To upgrade to Storage Foundation 5.0 MP1 on a standalone system

- 1 Log in as superuser.
- 2 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 3 Perform any necessary preinstallation checks and configuration. See the *Veritas Storage Foundation 5.0 Installation Guide*.

4 Check if the system's root disk is under VxVM control:

```
# df -k /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- a** Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- b** Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices:

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

5 Check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -F vxfs
```

6 Unmount all Storage Checkpoints and VxFS file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

7 Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean  
flags 0 mod 0 clean clean_value
```

A `clean_value` value of `0x5a` indicates the file system is clean, `0x3c` indicates the file system is dirty, and `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.

- a** If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem  
# mount -F vxfs filesystem mountpoint  
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

- 12 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```
 - 13 Insert the disc containing the Veritas software into the DVD-ROM drive, and mount the disc on a suitable mount point, for example:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```
 - 14 Move to the top-level directory on the DVD:

```
# cd /mnt/cdrom
```
 - 15 To upgrade the Storage Foundation software, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:
 - To install on the local system, enter the following command:

```
# ./installmp
```
 - To install on more than one system using secure shell (SSH) utilities, enter the following command:

```
# ./installmp system_name1 system_name2 ...
```
 - To install on more than one system using remote shell (RSH) utilities, enter the following command:

```
# ./installmp system_name1 system_name2 ... -rsh
```
 - 16 After the initial system checks have completed successfully, press **Enter** to start the requirements checks for the upgrade.
 - 17 After the requirement checks have completed successfully, press **Enter** to begin upgrading the packages.
-
- Note:** If you are upgrading multiple standalone systems, you can choose to upgrade the systems simultaneously.
-
- 18 After the upgrade of the packages is complete, use the following command to shut down the system:

```
# shutdown -y -i6 -g0
```
 - 19 Reboot the system, and reinstate any missing mount points in the `/etc/vfstab` file.
 - 20 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```
 - 21 Restart each RVG (if any were stopped in [step 8](#)):

```
# vxrvg -g diskgroup start rvg_name
```

- 22 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem  
# mount /checkpoint_name
```

- 23 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

- 24 If you set the value of the `vol_vvr_use_host_byte_order` tunable to 1 in [step d](#) of [step 8](#), reboot the system.

There are several optional configuration steps that you may perform:

If you want to use features of Veritas Storage Foundation 5.0 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.

To encapsulate and mirror the boot disk, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager 5.0 Administrator’s Guide*.

Changing permissions for Storage Foundation for Databases

After installing the Veritas Storage Foundation 5.0 MP1 patches, follow these post-installation steps to ensure Veritas Storage Foundation for Oracle and Veritas Storage Foundation for DB2 commands work correctly. This does not apply to Sybase.[772592]

Note: Do not recursively change permissions, groups, or owners.

To change permissions

- 1 Change permissions for the following directory, depending on which product you have installed:

For Veritas Storage Foundation for Oracle:

```
# chmod 550 /opt/VRTSdbed
```

For Veritas Storage Foundation for DB2:

```
# chmod 550 /opt/VRTSdb2ed
```

- 2 Reset owner and group settings to the appropriate owner and group for the database administrators on your system.

For example, in Veritas Storage Foundation for Oracle, to change owner to the user `oracle` and the group `dba`, run the following command:

```
# chown oracle:dba /opt/VRTSdbed
```

In Veritas Storage Foundation for DB2, for example, to change owner to the user `db2` and the group `db2grp`, run the following command:

```
# chown db2:db2grp /opt/VRTSdb2ed
```

- 3 Upgrade the repository.

In a standalone instance, run `sfua_db_config` once:

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

In a cluster environment, follow these steps:

- a Unconfigure the SFUA repository from the VCS configuration:

```
# /opt/VRTSdbcom/bin/sfua_db_config -o unconfig_cluster
```
- b Mount the repository file system manually.
- c Run the repository upgrade command again with no options:

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

Enabling Storage Mapping in Veritas Storage Foundation for Sybase

After installation, the root user must run the following command to permit Sybase users to use the Storage Mapping feature:

```
# /opt/VRTSobc/pal33/bin/veaconfig -c add_user -o host \  
-r Operator -n user@host.unixpwd
```

For example:

```
# /opt/VRTSobc/pal33/bin/veaconfig -c add_user \  
-o rockfowl.veritas.com -r Operator \  
-n sybase@rockfowl.veritas.com.unixpwd
```

Removing patches

The following procedures remove the patches if you want to uninstall MP1.
See [“Patch issues”](#) on page 41.

Removing an installed patch

To remove an installed patch

- 1 Log into the machine as root.
- 2 Remove the patch using the `patchrm` command.

```
# patchrm 117080
```
- 3 Shut down the system.

```
# shutdown -i6 -g0 -y
```

Removing an installed patch from a cluster

To remove an installed patch from a cluster

- 1 Log into one of the nodes as superuser.
- 2 Separate the node from the cluster.

```
# hastop -local
```
- 3 Stop the `vxconfigd` daemon on that node.

```
# vxdctl stop
```
- 4 Remove the patches.

```
# patchrm -G 117080
```
- 5 Shut down the node.

```
# shutdown -i6 -g0 -y
```
- 6 Repeat these steps for each node in the cluster.

Software limitations

The following sections describe Veritas Storage Foundation software limitations that exist in this 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/283886>

Veritas Storage Foundation for Databases software limitations

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

<http://entsupport.symantec.com/docs/283886>

The following are new software limitations in this MP1 release of Veritas Storage Foundation for Databases:

Veritas Storage Foundation for Oracle software limitations

DBDST limitations with non-English filenames and placement class names

DBDST does not work on non-English database filenames or non-English placement class names, due to limitations in VxFS Dynamic Storage Tiering and VxVM volume tags. VxFS Dynamic Storage Tiering does not support placement of non-English filenames. The VxVM volume tag feature does not support non-English volume tag names. [599164]

Differing locales produces unintelligible characters in GUI

The GUI does not support Oracle users having a different locale than the superuser's locale. The GUI will display unintelligible characters if the SFDB repository server starts with a locale that is different from the Oracle user locale (client). [605487]

Some features stop working after a GCO failover

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

Features like Storage Checkpoint, Database FlashSnap, the scheduler, and Database Dynamic Storage Tiering (DBDST) will not function as normal after a failover. However, features such as Oracle Disk Manager (ODM), Quick I/O, and Concurrent I/O (CIO) will continue to work after a failover. This issue will be fixed after the next release. [563603]

Veritas Enterprise Administrator limitations

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

DBDST class names limited to 29 characters

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. [601746]

Selected utilities require setuid

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`

[643964]

Multiple archive log destinations with RAC

Multiple archive log locations are not supported in RAC configurations. [795617]

Repository hostnames are case insensitive

Because DNS host name lookup queries are by definition case insensitive, make sure the SFDB repository is running on a host with a name that is truly unique -- regardless of case -- within the local subnet. Errors may occur if the repository host name differs from another host name only by case. [859863]

Veritas Storage Foundation for DB2 software limitations

DBDST limitations with non-English filenames and placement class names

DBDST does not work on non-English database filenames or non-English placement class names, due to limitations in VxFS Dynamic Storage Tiering and VxVM volume tags. VxFS Dynamic Storage Tiering does not support placement of non-English filenames. The VxVM volume tag feature does not support non-English volume tag names. [599164]

Some features stop working after a GCO failover

Some Veritas Storage Foundation for DB2 features do not work correctly after a Global Cluster (GCO) Failover. In 5.0, the Storage Foundation for Database (SFDB) repository and tools do not manage virtual hostnames correctly in a Global Cluster environment. The SFDB repository does not correctly adjust to the secondary host after the failover.

Features like Storage Checkpoint, Database FlashSnap, the scheduler, and Database Dynamic Storage Tiering (DBDST) will not function as normal after a failover. However, features such as Oracle Disk Manager (ODM), Quick I/O, and Concurrent I/O (CIO) will continue to work after a failover. This issue will be fixed after the next release. [563603]

Avoid using UNIX VEA via PC-based UNIX emulators

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

CLI database state changes are delayed in GUI

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

Workaround: Start or stop the database from the GUI, or do a manual rescan from the GUI after starting or stopping with CLI. [604685]

DBDST class names limited to 29 characters

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. [601746]

Selected utilities require setuid

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 DB2:

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

[643964]

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

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.

Workaround: Use the `qio_recreate` command to re-create the necessary Quick I/O files before you restore the database. [25272]

Repository hostnames are case insensitive

Because DNS host name lookup queries are by definition case insensitive, make sure the SFDB repository is running on a host with a name that is truly unique -- regardless of case -- within the local subnet. Errors may occur if the repository host name differs from another host name only by case. [859863]

Veritas Storage Foundation for Sybase software limitations

Poor Sybase performance with Concurrent I/O

Due to the way Sybase Adaptive Server Enterprise manages devices that use operating system files, the performance of Concurrent I/O files may not be as good as the performance of Quick I/O. Sybase has opened a feature request (CR 444156 - Enhance CIO on Veritas Storage Foundation) to fully exploit the Veritas Concurrent I/O feature in ASE. Until that support is delivered, we recommend customers continue to use Quick I/O for their Sybase database files. [836921]

Repository hostnames are case insensitive

Because DNS host name lookup queries are by definition case insensitive, make sure the SFDB repository is running on a host with a name that is truly unique -- regardless of case -- within the local subnet. Errors may occur if the repository host name differs from another host name only by case. [859863]

No longer supported

This section describes Veritas Storage Foundation features that are no longer supported in this release and future end of support notices.

- The use of the `vxvoladm` command line utility will not be supported in the next major release of Veritas Storage Foundation.

Fixed issues

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

For a list of additional issues fixed in this release, see the following TechNote:
<http://entsupport.symantec.com/docs/285869>

Veritas Volume Manager fixed issues

The following issues have been fixed in this release of VxVM.

Incident	Description
528677	Volume relayout is now supported for site-confined volumes and for site-consistent volumes.
540351	Reattaching a site when the disks were in the serial-split brain condition gave an error message.
540523	Under some circumstances, DMP nodes could be incorrectly enabled.
563524	Split, join and move operations failed on a source disk group that had any site-confined volumes.
584200	The <code>vxmake</code> command could not be used to recreate site records. This is now supported if the <code>-d</code> option is used to read from a description file.
601274	In a CVM cluster, DMP did not fail over to a secondary path when the primary paths were disconnected.
605743	If a disk group were split from a source disk group, volumes in the split-off disk group did not retain their volume tags.
609199	When the <code>vxmpadm disable</code> command was applied to a primary path on one node in a CVM cluster, the other nodes did not fail over to the secondary path.
611333, 622508	DMP could not obtain the correct serial number for a device if its LUN serial number contained a comma. This problem was seen on EMC Symmetrix arrays with more than 8096 LUNs.
614061, 614787	Adding cache volumes (used by space-optimized instant snapshots) to volume sets could cause data corruption and system panics.

Incident	Description
617331, 631334	I/O was not restored on a path that was re-enabled after a failback or a non-disruptive upgrade (NDU) operation.
618317	A system crash could occur while bringing up cluster if I/O were performed on a unopened path.
621832	Immediately after installation, the <code>vxesd</code> daemon had the DVD mount point as its current working directory, which prevented the DVD from being unmounted.
625877	The error <code>"/etc/vx/vxvm-startup: line 241: /usr/sbin/vxddladm: No such file or directory"</code> was seen at boot time.
643089	Relayout from <code>mirror-stripe</code> to <code>concat-mirror</code> 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 <code>dmp_health_time</code> and <code>dmp_path_age</code> tunables were both set to 0. This feature is now enabled by default in 5.0 MP1. The default values of <code>dmp_health_time</code> and <code>dmp_path_age</code> are 60 and 300 seconds respectively.

Veritas Enterprise Administrator fixed issues

The following issues have been fixed in this release of VEA.

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.

Incident	Description
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 fixed issues

The following issues have been fixed in this release of the Web GUI.

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.
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.

Incident	Description
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.
805595, 807387	Migrating from Central mode to Standalone mode caused the Action Agent package (VRTSaa) to be removed.

Veritas File System fixed issues

The following issues have been fixed in this release of VxFS:

Incident	Description
616323	For WebGUI online help, the following issues have been fixed: For the Remount Storage Checkpoint operation, the More info link on the second wizard page did not function properly for cluster file systems. For the Unmount Storage Checkpoint operation, the More info link on the second wizard page did not function properly for cluster file systems.
770917	Inode ownership issues detected in large directory related code paths have been fixed.
770935	Prevented the system from panicking when setting access time (<i>atime</i>) or modification time (<i>mtime</i>) of named data streams by calling <code>vxfv_nattr_utimes()</code> API on 32-bit kernel.
770953	<code>fsck</code> used to create the <code>lost+found</code> directory with the <code>rwrxrwxrwx</code> permissions if it decided to create one. Now, it creates the directory with the <code>rwxr-xr-x</code> permissions, which is consistent with the behavior of <code>mkfs</code> .
770964	<code>fsck</code> has been enhanced to replay file systems created with earlier log versions on volume sets.
771086	Fixed an <code>fsck</code> problem in which users could end up creating multiple <code>lost+found</code> 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 <code>lost+found</code> entries and removes duplicate directory entries.

Incident	Description
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 <code>lost+found</code> 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 <code>lost+found</code> 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 <code>lost+found</code> directory. The problem happened on single-volume and multi-volume file systems.
785649	A situation where <code>vxfsconvert</code> of a dusty file system loops forever in user-level code when an inode with pending truncation operation is encountered has been fixed.
793022	<code>vxfs_nattr_open()</code> API interface has been fixed to shrink files, as appropriate, when invoked with <code>O_TRUNC</code> flag.
793030	<code>vxfsutil.h</code> uses <code>struct timeval</code> in one of the function declarations, but does not include <code>time.h</code> . This causes user applications to report warnings during compilation. This issue has been fixed by including <code>time.h</code> in <code>vxfsutil.h</code> .
795073	The increased CPU utilization when writing to a file system that is almost full due to more background processing threads than are actually required being enqueued has been fixed.
858866	Resolved a failure in VxFS that caused installation to fail on an alternate disk when using Solaris Live Upgrade driven through CPI.

Veritas Volume Replicator fixed issues

Veritas Volume Replicator vradmin fixed issues

The following table contains information about fixed issues for VVR vradmin in this release:

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

Veritas Volume Replicator Web GUI fixed issues

The following table contains information about fixed issues in this release of VVR Web GUI:

Incident	Description
516812	On HP-UX, uninstalling did not remove all VRW files and directories.
576729	When a user without sufficient privileges tried to use the Create Primary wizard, the operation resulted in an incorrect error message.
611792	When a user tried to create a Primary without specifying the RVG name, the operation resulted in an RVG with an invalid configuration being created.
612565	VRW sometimes displayed the sizes of the SRL and data volumes incorrectly for an RDS replicating between VVR 4.1 and VVR 5.0 on the HP-UX operating system.
615758	If a CVM master node contains a private disk group, a Primary RVG created in that disk group was not displayed in the VVR Web GUI.

Incident	Description
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.
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.

Veritas Volume Replicator VEA fixed issues

The following table contains information about fixed issues in this release of VVR VEA:

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.

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

Veritas Storage Foundation for Oracle fixed issues

The following issues have been fixed in this release of Veritas Storage Foundation for Oracle:

Incident	Description
567342	An unmounted checkpoint clone database no longer 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 no longer produces the following error message: Error V-39-53246-8 Get EntryPoint failed. Please check the manifest related information
584044	The GUI now shows the mount point information correctly if the user mounts a read/write checkpoint again.
597257	Removing VRTSdbms3 may leave behind /opt/VRTSdbms2/bin32/servername or /opt/VRTSdbms3 which can cause a new installation of VRTSdbms3 to fail.
604849	You need not run <code>dbed_update</code> from the command line before starting an Oracle database from the GUI. Previously, the tablespace folder was empty when starting the Oracle database from the GUI, unless first running <code>dbed_update</code> from the command line.
604858	In the datafile statistic wizard, if there are no statistics for a datafile, a warning message confirming that no statistics are available is now shown instead of showing only an empty table. Also, in the datafile statistic wizard, if a user does not select a statistic, the wizard now shows a warning message and will not advance to next page.
605776	Previously, in the VEA GUI main window, the toolbar used the same hotkey of “ V ” for both View and View Mapping commands. Now View Mapping uses “ M ” as its hotkey.

Incident	Description
605583	The repository database server now handles the ja_JP.PCK locale correctly.
607001	Repository changes resulting from executing SFDB Storage Checkpoint CLIs are no longer delayed in the SFDB GUI.
607082, 610519	A problem was fixed that sometimes caused Web GUI requests to produce the error message Error V-40-49408-54.
607618	Previously, when scheduling a start up clone database task from the GUI, after all the clone database information was entered in the Start Up Snapshot database page, the Next button was disabled, preventing the user from finishing the scheduled task creation.
608667	The Rescan and Properties commands in the popup-menu have the same short cut key 'R'. In the Create Shortcut Wizard, the <code>Create checkpoint</code> command now uses the 't' key as a shortcut, and the <code>Retain this checkpoint...</code> command uses the 'e' key. Previously both commands used the same short cut key 't'.
608697	You can now refresh the View Statistics wizard in the Firefox browser.
608697	The Web GUI statistic scheduler no longer skips the first statistic collection.
609684	Specifying with <code>dbed_vmsnap</code> a snapplan that does not exist no longer produces the error <code>SFORA dbed_vmsnap ERROR V-81-6518 Could not find snapplan 'snap_plan' in repository.</code>
611154, 611152	Upgrading to SFDB 5.0 from a previous release's repository no longer produces the following error message after running <code>dbed_update</code> : [Sybase][ODBC Driver][Adaptive Server Anywhere]Index 'UniqueSnapshotTablespace' for table 'SnapshotTablespace' would not be unique (23000). <code>SFORA dbed_update ERROR V-81-3048 Could not upgrade repository.</code>
611944, 611942	The <code>sfua_db_config</code> command no longer creates an extra <code>/tmp</code> directory.
615819	Running statistic tasks from the Web GUI is now fully supported when connecting to a 4.1 host. No error messages are produced when running "Create Statistic Task."
786989	The <code>qio_getdbfiles_ora</code> script now detects when an Oracle instance is in Standby mode.

Incident	Description
853363	<p>The I/O performance of EMC Symmetrix arrays has been improved in this release. 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:</p> <ol style="list-style-type: none"> Determine the agent name under which the Symmetrix provider is configured: <pre># /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># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh \ -a agent_name -s 0</pre> <p>where <i>agent_name</i> is the agent name out put from the -l option in the previous step.</p>

Veritas Storage Foundation for DB2 fixed issues

The following are fixed issues in this release of Veritas Storage Foundation for DB2:

Incident	Description
417505	Executing an offhost <code>db2ed_vmclonedb</code> command with the <code>-o umount</code> option no longer produces a Segmentation Fault message for both <code>online_snapshot</code> and <code>offline</code> modes.
567342	An unmounted checkpoint clone database no longer 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.
584044	The GUI now shows the mount point information correctly if the user mounts a read/write checkpoint again.
597257	Removing <code>VRTSdbms3</code> may leave behind <code>/opt/VRTSdbms2/bin32/servername</code> or <code>/opt/VRTSdbms3</code> which can cause a new installation of <code>VRTSdbms3</code> to fail.
600490	The <code>db2ed_clonedb</code> and <code>db2ed_vmclonedb</code> commands now support automatic storage databases.

Incident	Description
604858	In the datafile statistic wizard, if there are no statistics for a datafile, a warning message confirming that no statistics are available is now shown instead of showing only an empty table. Also, in the datafile statistic wizard, if a user does not select a statistic, the wizard now shows a warning message and will not advance to next page.
604853	If a DB2 database is not in the SFDB repository and a user tries to start the instance from the GUI, the database is now properly displayed in the GUI.
605583	The repository database server now handles the ja_JP.PCK locale correctly.
605776	Previously, in the VEA GUI main window, the toolbar used the same hotkey of "V" for both View and View Mapping commands. Now View Mapping uses "M" as its hotkey.
607082, 610519	A problem was fixed that sometimes caused Web GUI requests to produce the error message <code>Error V-40-49408-54</code> .
607890, 588559	DBDST commands no longer occasionally produce output messages that contain references to Oracle (SFORA) instead of to DB2 (SFDB2).
608667	The Rescan and Properties commands in the popup-menu have the same short cut key 'R'. In the Create Shortcut Wizard, the <code>Create checkpoint</code> command now uses the 't' key as a shortcut, and the <code>Retain this checkpoint...</code> command uses the 'e' key. Previously both commands used the same short cut key 't'.
608697	You can now refresh the View Statistics wizard in the Firefox browser.
608697	The Web GUI statistic scheduler no longer skips the first statistic collection.
610283, 781751	If a checkpoint uses a mount prefix that was already used for another checkpoint, <code>db2ed_ckptmount</code> now produces an error message. Previously, the command would quit silently with return code 1.
610783, 610519, 607082	The DBEDAgent no longer fails to start if it is installed as a central managed host.
611944, 611942	The <code>sfua_db_config</code> command no longer creates an extra <code>/tmp</code> directory.
612220	The <code>db2ed_vmchecksnap</code> command no longer shows a <code>"grep: 0652-033 Cannot open - "</code> error multiple times in its output.
615819	Running statistic tasks from the Web GUI is now fully supported when connecting to a 4.1 host. No error messages are produced when running "Create Statistic Task."

Incident	Description
853363	<p>The I/O performance of EMC Symmetrix arrays has been improved in this release. 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:</p> <ol style="list-style-type: none"> Determine the agent name under which the Symmetrix provider is configured: <pre># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh -l</pre> 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. Set the discovery mode to discover host-visible devices only: <pre># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh \ -a agent_name -s 0</pre> where agent_name is the agent name out put from the -l option in the previous step.

Veritas Storage Foundation for Sybase fixed issues

The following are fixed issues in this release of Veritas Storage Foundation for Sybase:

Incident	Description
582069	SFDB commands no longer fail when executed within a locale that differs from the locale in use when the SFDB server was started.
413352, 634093	The <code>qio_convertdbfiles -u</code> command now converts any QIO files back to native files, even if the <code>mkqio.dat</code> file contains a file that does not reside on a VxFS filesystem.
634095	<p>Non-root users can now use the Storage Mapping feature, which was added to the 5.0 release of Veritas Storage Foundation for Sybase.</p> <p>After installation, the root user must run the following command to permit Sybase users to view storage mapping:</p> <pre>/opt/VRTSobc/pal33/bin/veaconfig -c add_user -o host \ -r Operator -n user@host.unixpwd</pre> <p>For example:</p> <pre>/opt/VRTSobc/pal33/bin/veaconfig -c add_user \ -o rockfowl.veritas.com -r Operator \ -n sybase@rockfowl.veritas.com.unixpwd</pre>

Incident	Description
853363	<p>The I/O performance of EMC Symmetrix arrays has been improved in this release. 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:</p> <ol style="list-style-type: none"><li data-bbox="472 406 1175 609">1 Determine the agent name under which the Symmetrix provider is configured: <pre data-bbox="518 470 1132 493"># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh -l</pre>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.<li data-bbox="472 621 1175 765">2 Set the discovery mode to discover host-visible devices only: <pre data-bbox="518 656 1118 708"># /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh \ -a agent_name -s 0</pre>where agent_name is the agent name out put from the -l option in the previous step.

Known issues

The following sections describe Veritas Storage Foundation issues that are known in this release.

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/283886>

See the following sections for information about known problems and issues in this release of VxVM.

Patch issues

Patch removal causes vxconfigd to dump core

If you use the `patchrm` command to remove the VxVM patch (122058-06), the `vxconfigd` daemon dumps core when it is restarted, and the following error message is displayed:

```
VxVM vxconfigd ERROR V-5-1-0 Bus error - core dumped.
```

The error can occur if:

- Any volume is open (for example, the root disk is encapsulated, or a file system on a VxVM volume is mounted).
- Any process is accessing VxVM drivers that cannot be unloaded. In this case, a workaround is to use the `kill vx` and `ps -ef | grep -i vx` commands to make sure that all `vx*` processes other than `vxconfigd` are stopped before removing the VxVM patch.

The error is harmless, and the patch is removed correctly. VxVM functions normally if the system is rebooted. [796270]

Device issues

Extensible Firmware Interface (EFI) support

The Solaris 9 and 10 64-bit kernel Operating Systems provide support for disks larger than 1 terabyte. Disks of this size are formatted with the EFI disk label rather than the VTOC disk label. EFI formatted disks are supported with Veritas Volume Manager on Solaris 10 only.

[303294, 834313, Sun Bug ID 6226760]

DMP issues

Fabric Monitoring

The new Fabric Monitoring feature controls whether the Event Source daemon (vxesd) uses the Storage Networking Industry Association (SNIA) HBA API. This API allows DMP to improve the performance of failover by collecting information about the SAN topology and by monitoring fabric events. Note that the vendor-provided ASL must also support the use of the SNIA HBA API.

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

```
# vxddladm settune monitor_fabric=on
# vxddladm settune monitor_fabric=off
```

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

```
# vxddladm gettune monitor_fabric
```

The default setting of `monitor_fabric` is on.

[784343]

Handling intermittently failing paths

The `dmp_health_time` and `dmp_path_age` tunables control how DMP handles intermittently failing paths. The default values in VxVM 5.0 MP1 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

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 `vxdlmpadm 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`.

Automatic site reattachment

A new automatic site reattachment daemon, `vxstited`, has been implemented to provide automatic reattachment of sites. `vxstited` uses the `vxnotify` mechanism to monitor storage coming back online on a site after a previous failure, and to restore redundancy of mirrors across sites.

If the hot-relocation daemon, `vxrelocd`, is running, `vxstited` attempts to reattach the site, and allows `vxrelocd` to try to use the available disks in the disk group to relocate the failed subdisks. If `vxrelocd` succeeds in relocating the failed subdisks, it starts the recovery of the plexes at the site. When all the plexes have been recovered, the plexes are put into the ACTIVE state, and the state of the site is set to ACTIVE.

If `vxrelocd` is not running, `vxstited` reattaches a site only when all the disks at that site become accessible. After reattachment succeeds, `vxstited` sets the site state to ACTIVE, and initiates recovery of the plexes. When all the plexes have been recovered, the plexes are put into the ACTIVE state.

Note: `vxstited` does not try to reattach a site that you have explicitly detached by using the `vxdiag detachsite` command.

The automatic site reattachment feature is enabled by default. The `vxstited` daemon uses email to notify `root` of any attempts to reattach sites and to initiate recovery of plexes at those sites. To send mail to other users, add the user name to the line that starts `vxstited` in the `/lib/svc/method/vxvm-recover` startup script and run the `svcadm refresh vxvm/vxvm-recover` command (for Solaris 10 onward), or `/etc/init.d/vxvm-recover` and reboot the system (for OS releases before Solaris 10).

If you do not want a site to be recovered automatically, kill the `vxstited` daemon, and prevent it from restarting. To kill the daemon, run the following command from the command line:

```
# ps -afe
```

Locate the process table entry for `vxstited`, and kill it by specifying its process ID:

```
# kill -9 PID
```

If there is no entry in the process table for `vxstited`, the automatic site reattachment feature is disabled.

To prevent the automatic site reattachment feature from being restarted, comment out the line that starts `vxstited` in the

```
/lib/svc/method/vxvm-recover
```

 startup script and run the `svcadm refresh vxvm/vxvm-recover` command (for Solaris 10 onward), or `/etc/init.d/vxvm-recover` (for OS releases before Solaris 10).

Replacing a disk in a site-consistent disk group

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

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

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

After tagging the replacement disk, you can use `vxdiskadm` to replace the failed disk. [536853]

Domain controller mode in CVM clusters

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. [603213]

Delays in systems with DS4800 storage

In a cluster with a shared IBM System Storage DS4800 disk storage system that is under a very heavy I/O load, opening the primary paths of a LUN or joining a node may take a long time. For example, it can take up to 15 minutes for a node to join a single-node cluster where approximately 90 LUNS are present. This behavior occurs even if the node that is opening the LUN is not involved in the I/O activity, and even if is not busy in any other way. [616166]

Localization issues

Upgrading language packages

You must uninstall the old version of the language packages before installing the SF 5.0 language packages, `VRTSmulic` and `VRTSmuvmp`. [625958]

View Mapping screen

The View Mapping button and the field to enter a file or directory name in the VEA are not completely visible when viewed in the Japanese locale. These objects appear correctly if the window is made larger. [631174]

Help for Symantec Product Authentication Services

In the French, Japanese or Simplified Chinese locales, the help for the Symantec Product Authentication Services is displayed in English. Type the following commands to display the correct help for the locale:

```
# cd /opt/VRTSat/bin/sparcv9
# mv VxSS_Help.vxa VxSS_Help_en.vxa
# mv VxSS_Help_lang.vxa VxSS_Help.vxa
```

where *lang* is *fr*, *ja* or *zh* as appropriate for the locale. [631206]

Veritas Enterprise Administrator issues

VRTSobc33 package not completely installed

The VRTSobc33 (Veritas Enterprise Administrator Core) package in this release may fail to install completely. The workaround is for the `root` user to use the `tar` command to extract the package, and then use the `patchadd` command to install it. [865949]

Volume tags not displayed

In the VEA client for Microsoft Windows systems, existing volume tags are not displayed when adding a new volume tag. [602953]

Search does not return any objects for non-Administrator users

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

The workaround is to add the user to the Administrator group. [840452]

Action Agent displays error V-39-53246-2

When **actionagent** is selected from the Control Panel on the VEA client, the following error message displays:

```
Error V-39-53246-2
  A plug-in Component may be corrupt.
  Download of plug-in component failed.
  Do you want to retry?
```

Workaround: Perform the following steps on the VEA server:

- 1 Execute `vxregctl` with the following parameters:


```
# /opt/VRTSobc/pal33/bin/vxregctl /etc/vx/isis/Registry \
setvalue Software/VERITAS/VRTSobc/pal33/Agents/actionagent/
ClientExtensions/{b0c6307c-d11e-46a8-8a33-82b599c20ee9}/
IsisSkin/1/1033 Version REG_SZ 3.3.74.0
```

Note that the value beginning with “Software/...” and ending with “.../1033” must be on a the same line with no spaces or breaks.
- 2 Restart the action agent:


```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c restart
```

Password field does not display on a localized Veritas Enterprise Administrator client

The password field in the Enter Password dialog of a localized Microsoft Windows VEA client does not display. This issue does not occur on a localized Solaris VEA client.

Workaround: Use one of the following workarounds:

- ◆ Change the password using a VEA client for Solaris.
- ◆ Temporarily use the English version of the VEA client for Windows.

To use the English Windows VEA client

- 1 On the Microsoft Windows system running the localized VEA client, remove the client.
- 2 Change the Regional Settings to **English** from the Control Panel.

Veritas Web GUI issues

Incorrect error message when importing a disk group

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

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

The workaround is to refresh the page. [607096]

Solaris x64 hosts cannot be managed

The Web GUI cannot be used to manage Solaris x64 for Opteron hosts that are running the Storage Foundation 4.1 software. [615554]

Error when creating a volume set

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.
```

The workaround is to refresh the page. [615960]

Veritas File System 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/283886>

The following are new known issues in this MP1 release of Veritas File System:

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

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

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

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

Veritas Volume Replicator known issues

Known 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/283918>

The following are new known issues in this MP1 release of Veritas Volume Replicator:

Synchronizing volumes and RVG with large volumes

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

Workaround: Instead of using the `vradmin syncrvg` command to synchronize the RVG, use the Automatic Synchronization feature when starting replication. To do this, use the `vradmin startrep -a` command. Or, reduce the size of the volume to below 1TB before running the `vradmin syncrvg` command.

For the `vradmin syncvol` command, the only workaround is to reduce the size of the volume to below 1TB. [840217]

Resynchronizing data between VVR 4.1MP2 (Solaris) and VVR 5.0MP1

You cannot use the `vradmin syncvol` command or the `vradmin syncrvg` command to resynchronize data if the Primary is using VVR 4.1MP2 on Solaris and the Secondary is using VVR 5.0MP1 or vice versa. Using these commands causes a core dump due to differences in the resynchronization code in these versions of VVR.

Workaround: Instead of using the `vradmin syncrvg` command to synchronize the RVG, use the Automatic Synchronization feature when starting replication. To do this, use the `vradmin startrep -a` command.

For the `vradmin syncvol` command, there is no workaround. We recommend upgrading both hosts to VVR 5.0MP1, if possible. [846685]

Issue with VVR VEA in the Japanese locale

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.

[616709]

Veritas Storage Foundation for Databases known issues

Veritas Storage Foundation for Oracle 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/283886>

The following are new known issues in this MP1 release of Veritas Storage Foundation for Oracle:

Cannot unmount single-host clone in HA environment after failover

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. [818522]

Workaround: Fix the issue that caused the failover to the other node, and then fall back to the fixed node.

File fragmentation check in the `qio_convertdbfiles` command may report errors

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. [819430]

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.

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 5 in the uninstallation procedure in “Uninstalling Veritas Storage Foundation” in the *Veritas Storage Foundation Installation Guide*.
- 3 Remove the Veritas Storage Foundation for Oracle packages using the `pkgrm` command.

```
# pkgrm VRTSorgui VRTSdbed VRTSdbdoc VRTSdbcom
```
- 4 Verify the removal of the packages using the `pkginfo` command.

```
# pkginfo | grep VRTS
```

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. [840486]

dbed_vmclonedb -p failed to create clonedb with modified pfile

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 that must be either placed within quotes or escaped when used as pfile parameter values. In some cases, Oracle will process a pfile correctly at startup even if a parameter value contains unquoted special characters. However, the pfile parser uses by Veritas Storage Foundation for Oracle 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.[853792]

One-time scheduled tasks need Specific Date

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. [861274]

Cannot use Web GUI to view snapplan log

When trying to view a snapplan log with the Web GUI, the error message "Unable to load operation" is displayed.

Workaround: View snapplan logs with the VEA Java GUI or with the `dbed_vmchecksnap -o list` command. [861696]

Database FlashSnap archive log destinations

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. [862092, 862687]

Veritas Storage Foundation for DB2 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/283886>

The following are new known issues in this MP1 release of Veritas Storage Foundation for DB2:

Semaphore 0 conflict causes DB2 commands to hang after system reboot

If after system reboot you experience a hang in `db2ed_update` or other SFDB2 commands on a Solaris 10 machine with DB2 v8.1 software installed, it may be caused by semaphore 0 being locked by the DB2 commands. This problem usually happens after a system reboot followed by a first-time `db2start`. The SFDB2 scripts usually hang in the `db2ed_dbprocli` commands.

The `db2hmon` process will crash with a SEGV error on a Solaris 10 machine with Solaris patch 119963-04 (refer to IBM APAR number IY85577). If `db2hmon` crashes during `db2start`, semaphore 0 may be locked by the DB2 process and this will cause SFDB2 CLI to hang.

If you have verified that your machine has the Solaris 10 patch 119963-04, you can do one of the following to resolve this hanging issue:

- upgrade your DB2 v8.1 version to FixPak 13
- disable DB2 health monitor for the DB2 instance
- if the above is not applicable, run `ipcrm -s 0` to clear the semaphore 0 lock state [603674, 603722, 785803]

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.

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 5 in the uninstallation procedure in “Uninstalling Veritas Storage Foundation” in the *Veritas Storage Foundation Installation Guide*.
- 3 Remove the Veritas Storage Foundation for DB2 packages using the `pkgrm` command.

```
# pkgrm VRTsd2gui VRTSdb2ed VRTSdbdoc VRTSdbcom
```
- 4 Verify the removal of the packages using the `pkginfo` command.

```
# pkginfo | grep VRTS
```

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. [840486]

One-time scheduled tasks need Specific Date

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. [861274]

Cannot use Web GUI to view snapplan log

When trying to view a snapplan log with the Web GUI, the error message "Unable to load operation" is displayed.

Workaround: View snapplan logs with the VEA Java GUI or with the `db2ed_vmchecksnap -o list` command. [861696]

Veritas Storage Foundation for Sybase 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/283886>

The following are new known issues in this MP1 release of Veritas Storage Foundation for Sybase:

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

If Veritas Storage Foundation for Sybase 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 Sybase packages manually if you want to uninstall the product.

To uninstall the Veritas Storage Foundation for Sybase packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 Follow steps 1 through 5 in the uninstallation procedure in "Uninstalling Veritas Storage Foundation" in the *Veritas Storage Foundation Installation Guide*.
- 3 Remove the Veritas Storage Foundation for Sybase packages using the `pkgrm` command.

```
# pkgrm VRTSsybed VRTSdbdoc
```
- 4 Verify the removal of the packages using the `pkginfo` command.

```
# pkginfo | grep VRTS
```

If Veritas Storage Foundation for Sybase 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 Sybase.

If the second upgrade fails, remove the Veritas Storage Foundation for Sybase packages as described above, then run the installer to upgrade Veritas Storage Foundation for Sybase. [840486]

Documentation errata

The following sections describe Veritas Storage Foundation documentation errata for the 5.0 release.

Veritas Storage Foundation Installation Guide errata

Please note the following errata in the *Veritas Storage Foundation Installation Guide*:

- For Solaris 10, patch 119254-06 or higher is a prerequisite.
- The `VRTSvxvm` patch is not required for Storage Foundation Basic, although it can optionally be installed.
- The `VRTSvxvm` patch is required if Veritas Volume Replicator (VVR) is licensed, or if Storage Foundation Basic is upgraded to Storage Foundation Enterprise. The patch must be applied by using the `patchadd` command.
- If the upgrade scripts are used to upgrade a system with an encapsulated root disk from VxVM 3.5 MP4 to VxVM 5.0, the upgrade completes successfully, but the root disk does not appear to be encapsulated. Use the following workaround:
 - 1 Copy the `/opt/VRTS/install/upgrade_start_system.SAV` file to `/etc/system`.
 - 2 Copy the `/opt/VRTS/install/upgrade_start_vfstab.SAV` file to `/etc/vfstab`.
 - 3 Edit the `/etc/vfstab` file, change all occurrences of "rootdg" in paths to read "bootdg", and then save the file.
 - 4 Reboot the system.
- Using the `installsf` or `installvm` installation scripts with the `-configure` option to upgrade VxVM 3.5 MP4 may fail with the following error:

```
CPI ERROR V-9-10-1052 Cannot upgrade SF to version 5.0 on
vm240v5 as 3.5MP4
```

Use the following workaround:

- 1 Remove the `.SF.upgrade` file.

```
# rm /opt/VRTS/install/.SF.upgrade
```
- 2 Recreate the `.SF.upgrade` file.

```
# touch /opt/VRTS/install/.SF.upgrade
```
- 3 Re-enter the `installsf` or `installvm` command, for example:

```
# installvm -configure hostname
```

- If you are uninstalling Veritas Storage Foundation using the `pkgrm` command, the packages must be removed in a specific order or else the uninstallation will fail.

Caution: Removing the packages out of order will result in some errors, including possible core dumps, although the packages will still be removed.

To uninstall Veritas Storage Foundation

- 1 Unmount all VxFS file systems and Storage Checkpoints, and close all VxVM volumes.
- 2 Remove the packages in the following order:

```
# pkgrm VRTSmapro VRTSgapms VRTSvxmsa VRTSfasdc \
VRTSfas VRTSvail VRTSfsmnd VRTSfssdk VRTSfsdoc \
VRTSfsman VRTSvxfs VRTSvrdoc VRTSvrw VRTSweb VRTSjre15 \
VRTSjre VRTSvcsvr VRTSvrpro VRTSddlpr VRTSvdid \
VRTSvsvc VRTSvmpro VRTSalloc VRTSdcli VRTSvmdoc \
VRTSvmman VRTSfspro VRTSdsa VRTSvxvm SYMCLma VRTSspt \
VRTSaa VRTSmh VRTSccg VRTSobgui VRTSob VRTSobc33 \
VRTSat VRTSsmf VRTSspb VRTSicsco VRTSperl
```

- (Page 59) Storage Foundation Basic can be installed on Solaris 8, 9, and 10 systems (both 32 and 64-bit), and not just on Solaris 10 (64-bit) systems.
- (Page 65) If a Storage Foundation Basic system is upgraded to Storage Foundation Enterprise with the Veritas Volume Replicator licensed feature, you must additionally use the product installer to add the VVR packages (available under item 4).
- (Page 85) In the procedure "Upgrading Veritas Storage Foundation on Solaris 8, 9, or 10," ignore step 7. You do not need to reboot the machine.
- (Pages 86-87) In the procedure "To upgrade on a system with an encapsulated root disk," there is no need to run the `upgrade_check`, `upgrade_start` and `upgrade_finish` scripts in steps 2, 3 and 13. These commands are run automatically by the packaging scripts.
- (Page 87) There should be an additional step between steps 14 and 15 to configure the software:
Run the `installer` script with the `-configure` option specified to configure the software.
- (Page 90) After step 2 in the procedure "To upgrade the operating system," there should be the following additional step:
Run the `installer` script with the `-configure` option specified to configure the software.

- (Page 91) After step 10 in the procedure "To upgrade the Veritas Storage Foundation packages after upgrading the operating system," there should be the following additional step:
Run the `installer` script with the `-configure` option specified to configure the software.
- (Page 98) Step 6 in the procedure "To upgrade from SUNWvxxvm if the root disk is unencapsulated" should read as follows:
6 Run the `installsf` program.
 - If the disc is mounted automatically, enter:

```
# cd /cdrom/cdrom0/storage_foundation
# installsf -installonly hostname
```
 - If the disc is mounted manually, enter:

```
# cd /mount_point/storage_foundation
# installsf -installonly hostname
```
- (Page 98) After step 10 in the procedure "To upgrade from SUNWvxxvm if the root disk is unencapsulated," there should be the following additional step:
Run the `installer` script with the `-configure` option specified to configure the software.
- (Page 136) The section "Uninstalling language packages using the `pkgadd` command" should be entitled "Uninstalling language packages using the `pkgrm` command."
- (Page 154) The procedure in the section "Upgrading VxVM and the Solaris OS" cannot be used to upgrade from VxVM 3.5 MP4. This is because the `upgrade_start` and `upgrade_finish` scripts are not supported for such an upgrade path.
- (Page 156) The procedure "Upgrading the Solaris OS only" should include an additional step after rebooting the system in step 14.
Rerun the `upgrade_finish` script from the `volume_manager/scripts` directory on the mounted software disc:

```
# cd /cdrom/cdrom0/volume_manager/scripts
# ./upgrade_finish
```

Note: This procedure only applies if you are upgrading the Solaris OS but not Veritas Volume Manager.

Web GUI help errata

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

Manual page errata

The `vxassist(1M)`, `vxddladm(1M)`, `vxdisk(1M)`, `vxdmpadm(1M)`, `vxdmpinq(1M)`, `vxpool(1M)`, `vxresize(1M)`, `vxtemplate(1M)`, and `vxvoladm(1M)` manual pages are updated in this Maintenance Pack to include corrections for several errors or omissions.

Veritas Volume Manager Administrator's Guide errata

The following errata apply to the *Veritas Volume Manager Administrator's Guide*:

Specifying storage for version 20 DCO plexes

The section “Specifying storage for version 20 DCO plexes” in the “Administering volumes” chapter of the *Veritas Volume Manager Administrator's Guide* includes the following example:

```
# vxsnap -g mydg prepare myvol ndcomirs=2 disk05 disk06
```

This should read:

```
# vxsnap -g mydg prepare myvol ndcomirs=2 alloc=disk05,disk06
```

The `vxsnap prepare` command requires that you use the `alloc` attribute when specifying the storage for DCO plexes.

DMP configuration values

The minimum value of the `dmp_path_age` tunable is documented as 1 second. The correct minimum configurable value of `dmp_path_age` is 0, which prevents DMP from detecting intermittently failing paths.

The default recovery option settings are stated to be `queuedepth=20` for throttling and `retrycount=30` for I/O error retrying. The correct default settings are `iotimeout=10` for throttling and `retrycount=5` for I/O error retrying.

Veritas Storage Foundation for Oracle Administrator's Guide errata

The following sections are missing from the *Veritas Storage Foundation for Oracle Administrator's Guide*:

Setting up Oracle 9i RAC objects with srvctl

When configured within an Oracle RAC environment, you must set up the Oracle `srvctl` service and register the name of the RAC database with `srvctl`, so that Veritas Storage Foundation for Oracle can learn the status of remote database instances. Otherwise, commands such as `dbed_ckptcreate -o offline` may fail.

To set up Oracle 9i RAC objects

- 1 Look in `/var/opt/oracle/srvConfig.loc` to learn the pathname to the SRVM configuration file as defined by the variable `srvconfig_loc`. For example:

```
srvconfig_loc=/db/srvm.ora
```

- 2 List the details of the SRVM configuration file with `ls -l`:

```
# ls -l /db/srvm.ora
```

- 3 If the configuration file does not exist, create and initialize the file:

```
# touch /db/srvm.ora
# srvconfig -init
```

- 4 If the configuration file exists, note the size of the file shown by the output of `ls -l`.

```
-rw-r--r-- 1 oracle dba 10569216 Jan 20 14:29 /db/srvm.ora
```

- 5 If the configuration file size is greater than zero (as shown in the example above), the file is initialized. If the file size is zero, initialize it:

```
# srvconfig -init
```

- 6 Start the Oracle RAC Manageability daemon on each system:

```
$ gsdctl start
```

- 7 Confirm the GSD daemon status:

```
$ gsdctl stat
GSD is running on the local node
```

- 8 Add the database to the `srvctl` configuration:

```
$ srvctl add database -d KPRDADV1 \
-o /apps/oracle/product/920rac
$ srvctl config database
KPRDADV1
```

- 9 Add each instance to the configuration. For example, in a two-instance configuration, add the first instance:

```
$ srvctl add instance -d KPRDADV1 -i KPADV1R1 -n node1
$ srvctl config database -d KPRDADV1
node1 KPADV1R1 /apps/oracle/product/920rac
```

Then add the second instance:

```
$ srvctl add instance -d KPRDADV1 -i KPADV1R2 -n node2
$ srvctl config database -d KPRDADV1
node1 KPADV1R1 /apps/oracle/product/920rac
node2 KPADV1R2 /apps/oracle/product/920rac
```

- 10 Check the status of the instances to confirm they are running:

```
$ srvctl status database -d KPRDADV1
Instance KPADV1R1 is running on node node1
Instance KPADV1R2 is running on node node2
```

Reconfigure virtual IP address for repository configuration

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

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

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

In a cluster environment, do the following:

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

a Unconfigure the SFUA repository:

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

b Import the repository disk group.

c Then, start then repository disk volume.

d Mount the repository file system.

e Then, run the command:

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

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

The following information is incorrect in the *Veritas Storage Foundation for Oracle Administrator's Guide*:

- (Page 289) In step 3 of the procedure "To remove a snapplan and snapshot volume," the correct command to remove a snapplan is:

```
# /opt/VRTS/bin/dbed_vmchecksnap -s db -f snapplan -o remove
```
- (Pages 179 and 343) In the table describing `dbed_clonedb` command options, the description of the `-d` option is potentially misleading. The description should read as follows:
Used with the `-o umount` option. If the `-d` option is specified, the read-write Storage Checkpoint mounted by `dbed_clonedb` is deleted along with the clone database.
Note that this does not delete the read-only Storage Checkpoint first created by `dbed_ckptcreate`, which is subsequently used by `dbed_clonedb` to create a read-write checkpoint.

Veritas Storage Foundation for Oracle Graphical User Interface Guide errata

The following information is incorrect in the *Veritas Storage Foundation for Oracle Graphical User Interface Guide*:

- (Page 23) In the procedure "To start the DBED agent," the command in step 2 should read as follows:

```
/etc/rc2.d/S75vxpal.DBEDAgent start
```
- (Page 23) In the procedure "To stop the DBED agent," the command in the single step should read as follows:

```
/etc/rc2.d/S75vxpal.DBEDAgent stop
```

Veritas Storage Foundation for DB2 Administrator's Guide errata

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

Reconfigure virtual IP address for repository configuration

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

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

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

In a cluster environment, do the following:

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

```
# /opt/VRTSdbcom/bin/sfua_db_config -o unconfig_cluster
```
 - b Import the repository disk group.
 - c Then, start then repository disk volume.
 - d Mount the repository file system.
 - e Then, run the command:

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

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

The following information is incorrect in the *Veritas Storage Foundation for DB2 Administrator's Guide*:

- (Page 243) In the procedure "To abort reverse resynchronization," the single step should read as follows:

Use the `-o reverse_resync_abort` option of the `db2ed_vmsnap` command:

Note that the command following this text is correct.
- (Pages 150 and 299) In the table describing `db2ed_clonedb` command options, the description of the `-d` option is potentially misleading. The description should read as follows:

Used with the `-o umount` option. If the `-d` option is specified, the read-write Storage Checkpoint mounted by `db2ed_clonedb` is deleted along with the clone database.

Note that this does not delete the read-only Storage Checkpoint first created by `db2ed_ckptcreate`, which is subsequently used by `db2ed_clonedb` to create a read-write checkpoint.

Veritas Storage Foundation for DB2 Graphical User Interface Guide errata

The following information is incorrect in the *Veritas Storage Foundation for DB2 Graphical User Interface Guide*:

- (Page 23) In the procedure "To start the DBED agent," the command in step 2 should read as follows:

```
/etc/rc2.d/S75vxpal.DBEDAgent start
```

- (Page 24) In the procedure "To stop the DBED agent," the command in the single step should read as follows:

```
/etc/rc2.d/S75vxpal.DBEDAgent stop
```

