Veritas Storage Foundation™ Release Notes

Linux

5.0 Maintenance Pack 2

symantec™

Veritas Storage Foundation Release Notes

Copyright © 2007 Symantec Corporation. All rights reserved.

Storage Foundation 5.0 Maintenance Pack 2

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 Veritas 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/283859

Linux is a registered trademark of Linus Torvalds.

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 phone or email support. 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

Veritas Storage Foundation Release Notes

Introduction

This document provides release information about the products in the Veritas Storage Foundation 5.0 Maintenance Pack 2 (MP2) product line for the Linux platform:

- Veritas Storage Foundation (Basic, Standard, Standard HA, Enterprise, and Enterprise HA)
- Veritas Storage Foundation for Oracle (Standard, Enterprise, and HA Editions)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)

Each of these Symantec products is activated by a single license key. You must obtain a license key before installing the product. For information on obtaining a license key, see the *Veritas Storage Foundation Installation Guide*.

The following products are not supported for the Oracle Enterprise Linux (OEL) platform in this release:

- Veritas Storage Foundation for DB2
- Veritas Storage Foundation for Oracle RAC
- Veritas Volume Replicator (VVR)

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

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

Review this entire document before installing your Veritas Storage Foundation product.

New features

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

Support for Oracle Enterprise Linux

Oracle Enterprise Linux (OEL) is a redistribution by Oracle of Red Hat Enterprise Linux Update 4 that has been customized for the Oracle product.

Veritas File System

The 5.0 MP1 and MP2 releases of Veritas File System includes the following new features and enhancements

Block device support in Veritas File System

VxFS 5.0 only supported the creation of a VxFS file system on a VxVM volume. Creating a VxFS file system on raw devices, such as SCSI or IDE disks, was not supported. VxFS in 5.0 MP1 and MP2 supports the creation of VxFS file systems directly on such devices, without requiring a VxVM volume.

Both partitioned and non-partitioned block devices may now be used directly with VxFS. For example:

```
# mkfs -t vxfs /dev/sdc1
# mount -t vxfs /dev/sdc1 /mnt
# fsck -t vxfs /dev/sdc1
```

Similarly, all local mount features of VxFS, such as clones and file system snapshots, do not require the creation of a VxVM volume.

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems

Storage Foundation operates on the following Linux operating systems and kernels distributed by Oracle, Red Hat and SUSE:

- Oracle Enterprise Linux, which is based on, and which has binary compatibility with, Red Hat Enterprise Linux 4 Update 4 (2.6.9-42 kernel) or higher on AMD Opteron or Intel Xeon EM64T (x86_64).
- Red Hat Enterprise Linux 4 (RHEL 4) with Update 3 (2.6.9-34 kernel) or higher on AMD Opteron or Intel Xeon EM64T (x86_64).
- SUSE Linux Enterprise Server 9 (SLES 9) with SP3 (2.6.5-7.244 kernels) or higher on AMD Opteron or Intel Xeon EM64T (x86_64).

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

Supported Linux operating system updates

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

Information about the latest supported Red Hat errata and updates and SUSE service packs is available in the following TechNote. Read this TechNote before installing any Veritas product.

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

Memory requirements

A minimum of 1 GB of memory is strongly recommended.

Supported Oracle versions

Oracle versions 9.2.0.6, 10g, and 10gR2 are all supported on the Linux operating systems listed above.

Mandatory patch required for Oracle Bug 4130116

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

Software and hardware requirements

The hardware compatibility list (HCL) contains the latest information about supported hardware and software and is updated regularly.

Note: Before installing or upgrading Veritas Volume Manager, review the current compatibility list to confirm the compatibility of your hardware and software.
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

If you do not find your hardware or software listed or if you have questions about the information in the compatibility list, contact Veritas Technical Services.

VxVM licenses

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

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

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

To see the license features that are enabled in VxVM, enter the following command:

```
# vxdctl license
```

Cross-Platform Data Sharing licensing

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

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

Component product release notes

Release notes for component products in all versions of 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 Storage Foundation Cluster File System Release Notes	`sfcfs_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 the Veritas Storage Foundation 5.0 MP2 software for the first time and not upgrading an existing system, read the *Veritas Storage Foundation Installation Guide* for pre-installation instructions. Also review the *Veritas Storage Foundation Release Notes* and all documents in the `release_notes` directory for important release information.

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 install the Storage Foundation software

1 Insert the Veritas Storage Foundation 5.0 MP2 software disc, and mount it on a suitable mount point. For example:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

2 Move to the top-level directory on the DVD:

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

3 To install the software, you can either run the product installation script or the generic `installer` script. For example, to install the Storage Foundation software on the local system by using the product installation script, enter the following command from the top-level directory of the mounted DVD-ROM:

```
# ./installsf
```

To install on more than one system, enter the following command:

```
# ./installmp system_name1 system_name2 ... [-rsh]
```

The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to install the software simultaneously on several systems.

Note: Provided that the remote shell (RSH) or secure shell (SSH) has been configured correctly, this command can be run on a single node of the cluster to install the software on all the cluster nodes.

For help on how to respond to the installation prompts, see the *Veritas Storage Foundation 5.0 Installation Guide.*

Preparing to upgrade to 5.0 MP2

If you are upgrading an installed Veritas Storage Foundation 5.0 or 5.0 RP1 system, read the *Veritas Storage Foundation Installation Guide* for instructions on how to preserve the existing configuration information.

In particular, perform the following actions:

- Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to recreate these entries in the `/etc/fstab` file on the freshly installed system.
- Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/boot/grub/menu.lst`, `/etc/grub.conf`, `/etc/elilo.conf`, or `/etc/lilo.conf` (as appropriate), and `/etc/fstab`. You should also run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.
- Use the `vxlicrep` command to make a record of the currently installed Veritas licenses.
- Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.

You should also review the *Veritas Storage Foundation Release Notes* and all documents in the `release_notes` directory for important release information.

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

- "Upgrading a cluster system to 5.0 MP2" on page 14.
- "Upgrading a standalone system to 5.0 MP2" on page 19.

Upgrading a cluster system to 5.0 MP2

An upgrade requires stopping cluster failover functionality during the entire procedure. The upgrade is performed in a number of stages depending on the type of upgrade you are performing. The supported upgrade paths are:

- SF 5.0 to 5.0 MP2
- SF 5.0 RP1 to 5.0 MP2

Caution: Phased upgrade procedure results in system PANIC on configurations where LLT is configured over UDP and this known issue is fixed in 5.0 MP2. This issue is specific to configurations where LLT is configured over UDP and not present in usual LLT Ethernet configurations.

Full upgrade procedure should be used for upgrading from SF 5.0 or SF 5.0 RP1 on configurations where LLT is configured over UDP.

Phased upgrade stages

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time. Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

Stages of a phased upgrade

1. Freeze service group operations and stop VCS on the cluster.
2. Select a group of one or more cluster nodes to upgrade, and leave a group of one or more nodes running.
3. Take the first group offline and install the software patches.
4. Take the second group offline **before** bringing the first group online.
5. Bring the first group (with the newly installed patches) online to restart cluster failover services.
6. Upgrade the remaining nodes in the second group and bring them online. The cluster is fully restored.

Proceed to “Performing a phased upgrade of a cluster to 5.0 MP2” on page 15.

Full upgrade stages

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

Stages of a full upgrade

1. Freeze service group operations and stop VCS on the cluster.
2. Take all nodes in the cluster offline and install the software patches.
3. Bring all the nodes (with the newly installed patches) online to restart cluster failover services. The cluster is fully restored.

Proceed to “Upgrading a standalone system to 5.0 MP2” on page 19.

Performing a phased upgrade of a cluster to 5.0 MP2

This section describes how to perform a phased upgrade of a cluster to 5.0 MP2.

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 a cluster

1. Log in as superuser.
2. From any node in the cluster, make the VCS configuration writable:
   ```
   # haconf -makerw
   ```
3. Enter the following command to freeze HA service group operations on each node:
   ```
   # hasys -freeze -persistent node_name
   ```
4. Make the configuration read-only:
   ```
   # haconf -dump -makero
   ```
5. Select the group of nodes that are to be upgraded first, and follow step 6 through step 32 for these nodes.
6. Stop VCS by entering the following command on each node in the group being upgraded:
   ```
   # hastop -local
   ```

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.

7 For each node, check if its root disk is under VxVM control:

```
# df -v /
```

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.

8 If required, upgrade the nodes, and patch them to a supported kernel version. For more information on supported operating systems and kernel versions, see "System requirements" on page 9.

9 On each node, check if any Veritas File Systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

10 Unmount all Storage Checkpoints and Veritas File Systems:

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

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 On each node, 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/fstab` file for Veritas File Systems, or for any file systems that are configured on VxVM volumes.

14 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl 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 are complete, press **Return** to start the requirements checks.

19 After the requirements checks are complete, press **Return** to start upgrading the packages. If you are upgrading multiple nodes, you have the option of upgrading them simultaneously. You will be prompted after the upgrade is complete.

20 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.

21 Stop VCS on each of the second group of nodes:

```
# hastop -local
```

22 Reboot the upgraded nodes. Assuming the reboot was successful, application failover is now available for the first group of nodes.

23 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.

24 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.

25 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.

26 Make the VCS configuration writable again from any node in the upgraded group:

```
# haconf -makerw
```

27 Enter the following command on each node in the upgraded group to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent node_name
```

28 Make the configuration read-only:

```
# haconf -dump -makero
```

29 Bring the CVM service group online on each node in the upgraded group:

```
# hagrp -online cvm -sys node_name
```

30 Restart all the volumes by entering the following command for each disk group:

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

31 Remount all VxFS file systems and Storage Checkpoints on all nodes:

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

32 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

33 Repeat step 6 through step 32 for the second group of nodes.

Upgrading a standalone system to 5.0 MP2

The following procedure describes upgrading a standalone system to 5.0 MP2.

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 a standalone system

1 Log in as superuser.

2 If the system is part of a cluster, run the following command on any node in the cluster:

```
# hastop -all
```

This stops VCS on all nodes in the cluster.

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

```
# df -v /
```

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.

4 If required, upgrade your system, and patch it to a supported kernel version. For more information on supported operating systems and kernel versions, see “System requirements” on page 9.

5 Check if any Veritas File Systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

6 Unmount all Storage Checkpoints and Veritas File Systems:

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

7 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.

8 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
```

9 Comment out any mount points in the `/etc/fstab` file for Veritas File Systems, or for any file systems that are configured on VxVM volumes.

10 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

11 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
```

12 Move to the top-level directory on the DVD:

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

13 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
  ```

14 After the initial system checks have completed successfully, press **Enter** to start the requirements checks for the upgrade.

15 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.

16 After the upgrade of the packages is complete, use the following command to shut down the system:

```
# shutdown -r now
```

17 Reboot the system.

18 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.

19 Reinstate any missing mount points in the `/etc/fstab` file.

20 If the system is part of a cluster and any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.

21 Restart all the volumes by entering the following command for each disk group:

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

22 Remount all Veritas File Systems and Storage Checkpoints:

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

23 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

Changing permissions for Storage Foundation for Databases

After installing the Veritas Storage Foundation 5.0 MP2 patches, follow these post-installation steps to ensure Veritas Storage Foundation for Databases commands work correctly. [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:

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

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
   ```

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
      ```

Verifying software versions

To list the Veritas packages installed on your system, enter the following command:

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

Removing the 5.0 MP2 packages

Roll back of the 5.0 MP2 packages to a previous version of the packages is not supported. To restore release 5.0 on your system, you must completely remove 5.0 MP2, then reinstall Veritas Storage Foundation 5.0.

To uninstall the Veritas software

1. Log in as superuser.
2. Verify that /opt/VRTS/bin is in your PATH so you can execute all product commands.
3. If the systems are running as an HA cluster, use the following command to take all service groups offline, and shut down VCS:

   ```
   # /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 VCS, causing undesired results during uninstallation of the packages.

4. Check if the root disk is under VxVM control by running this command:

   ```
   # df -v /
   ```

   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. There must be at least one other disk in the rootdg disk group in addition to the root disk for vxunroot to succeed.

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

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

5 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
        # df -T | grep vxfs
```

6 Unmount all Storage Checkpoints and file systems:

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

7 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.

8 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
```

9 Check if the VEA service is running:

```
    # /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
    # /opt/VRTS/bin/vxsvcctrl stop
```

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

```
    # cd /opt/VRTS/install
    # ./uninstallsf [-rsh]
```

You can use this command to remove the packages from one or more systems. The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to uninstall the software simultaneously on several systems.

Note: Provided that the remote shell (RSH) or secure shell (SSH) has been configured correctly, this command can be run on a single node of the cluster to install the software on all the cluster nodes.

After uninstalling the Veritas software, reinstall the release 5.0 software as described in the *Veritas Storage Foundation Installation Guide*, *Veritas Storage Foundation Cluster File System Installation Guide*, or *Veritas Cluster Server (VCS) Installation Guide*.

Upgrading a High Availability cluster

If you are upgrading an HA cluster, see the *Veritas Cluster Server Installation Guide* for information on preserving your VCS configuration across the upgrade procedure. In particular, back up configuration files, such as `main.cf` and `types.cf`, in the `/etc/VRTSvcs/conf/config` directory. If you installed any VCS agents, back up any additional configuration files in this directory, such as `Oracletypes.cf`.

Fixed issues

Note: There are no software issues fixed in the 5.0 Maintenance Pack 2 release relative to the 5.0 Maintenance Pack 1 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 table contains information about fixed issues in the 5.0 MP1 and MP2 releases 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 `vxmake` command could not be used to recreate site records.
601274	In a CVM cluster, DMP did not fail over to a secondary path when the primary paths were disconnected.
605743	If a disk group were split from a source disk group, volumes in the split-off disk group did not retain their volume tags.
609199	When the `vxdmpadm disable` command was applied to a primary path on one node in a CVM cluster, the other nodes did not fail over to the secondary path.
611333	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	I/O was not restored on a path that was re-enabled after a failback or a non-disruptive upgrade (NDU) operation.
618068	A system panic could occur when EMC PowerPath was configured to coexist with DMP as a third-party driver.
618317	A system crash could occur while bringing up cluster if I/O were performed on a unopened path.
619958	After encapsulating a SAN-connected root disk on an EMC CLARiiON CX700 array, a system panic occurred during booting when multiple paths were enabled. Systems with an encapsulated root disk on an Active/Passive array in explicit failover mode (A/PF array) failed to boot if more than one path was enabled. Systems with an encapsulated root disk on an Active/Passive array (A/P array) took a longer time than usual to boot.
621832	Immediately after installation, the `vxesd` daemon had the DVD mount point as its current working directory, which prevented the DVD from being unmounted.
625877	The error "/etc/vx/vxvm-startup: line 241: /usr/sbin/vxddladm: No such file or directory" was seen at boot time.
643089	Relayout from `mirror-stripe` to `concat-mirror` did not work for site-consistent volumes.
645749	Growing a volume by a specified amount did not work for a site-consistent volume with more than 2 disks per site.
771691	The `vxrootmir` command would not copy the Master Boot Record (MBR) to a root mirror disk. This behavior was generic to all arrays.
778352	Mirroring the root disk to a SAN boot disk could cause data corruption. This behavior was generic to all arrays.
793159	Automatic reattachment of a remote site did not work correctly.
801445	The DMP feature to detect and respond to intermittently failing paths was turned off by default in the 5.0 release, and the values of the `dmp_health_time` and `dmp_path_age` tunables were both set to 0. This feature is now enabled by default in 5.0 MP1 and MP2. The default values of `dmp_health_time` and `dmp_path_age` are 60 and 300 seconds respectively.

Veritas Enterprise Administrator fixed issues

The following issues have been fixed in the 5.0 MP1 and MP2 releases 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 were actually healthy, and ‘‘Healthy’’ when they were actually not healthy.
614761	The volume set creation wizard showed cache volumes in the ‘‘Available Volumes’’ list.
616661	When connecting to the central host, an ‘‘OutOfBoundException’’ error could occur.
618146	A Java exception error occurred in the Statistics View.

Veritas Web GUI fixed issues

The following issues have been fixed in the 5.0 MP1 and MP2 releases 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.

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

Veritas File System fixed issues

The following table contains information about fixed issues in the 5.0 MP1 and MP2 releases 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 does not function properly for cluster file systems. For the **Unmount Storage Checkpoint** operation, the **More info** link on the second wizard page does not function properly for cluster file systems.
770917	Inode ownership issues detected in large directory related code paths have been fixed.

Incident	Description
770935	Prevented the system from panicking when setting access time (*atime*) or modification time (*mtime*) of named data streams by calling `vxfs_nattr_utimes`() API on 32-bit kernel.
770953	`fsck` used to create the `lost+found` directory with the `rwxrwxrwx` permissions if it decided to create one. Now, it creates the directory with the `rwxr-xr-x` permissions, which is consistent with the behavior of `mkfs`.
770964	`fsck` has been enhanced to replay file systems created with earlier log versions on volume sets.
771086	Fixed an `fsck` problem in which users could end up creating multiple `lost+found` directories when running the `fsck -o full` command and answering `fsck` questions interactively. Now, `fsck` creates only one. It also checks for multiple `lost+found` entries and removes duplicate directory entries.
771996	Enhanced VxFS to use less CPU when doing administrative tasks on the devices of multi-volume file systems.
772013	Enhanced the `fsck` command to enforce the `lost+found` file name in the root directory of the file system to be a directory file type.
777012	If the system crashed or there was a metadata I/O error, after the `fsadm` command reorganized the `lost+found` directory, running the `fsck -o full` command may not have been able to clean the file system with regard to names that needed to be added to the `lost+found` directory. The problem happened on single-volume and multi-volume file systems.
785649	A situation where `vxfsconvert` of a dusty file system loops forever in user-level code when an inode with pending truncation operation is encountered has been fixed.
793022	The `vxfs_nattr_open`() API interface has been fixed to shrink files, as appropriate, when invoked with `O_TRUNC` flag.
793030	`vxfsutil.h` uses `struct timeval` in one of the function declarations, but does not include `time.h`. This causes user applications to report warnings during compilation. This issue has been fixed by including `time.h` in `vxfsutil.h`.
795073	The increased CPU utilization when writing to a file system that is almost full due to more background processing threads than are actually required being enqueued has been fixed.

Veritas Storage Foundation for Oracle fixed issues

The following issues have been fixed in the 5.0 MP1 and MP2 releases of the 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 executed with a different locale than the locale in use when the SFDB server was started no longer fail with the following message: `([Sybase][ODBC Driver][Adaptive Server Anywhere]Syntax error ).`
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`
600431	Storage Checkpoint operations are now supported for databases cloned with Database FlashSnap.
607001	Repository changes resulting from executing SFDB Storage Checkpoint CLIs are no longer delayed in the SFDB GUI.
608697	You can now refresh the View Statistics wizard in the Firefox browser. The Web GUI statistic scheduler no longer skips the first statistic collection.
609682	Specifying a snapplan that does not exist to `dbed_vmsnap` no longer produces the following error: `SFORA dbed_vmsnap ERROR V-81-6518 Could not find snapplan 'snap_plan' in repository.`
786989	The qio_getdbfiles_ora script now detects when an Oracle instance is in Standby mode.

Incident	Description

853363	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:
	1 Determine the agent name under which the Symmetrix provider is configured:
	`# /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh -l`
	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.
	2 Set the discovery mode to discover host-visible devices only:
	`# /opt/VRTSvail/bin/vail_symm_discovery_cfg.sh \`
	`-a agent_name -s 0`
	where *`agent_name`* is the agent name out put from the `-l` option in the previous step.

Known issues

Veritas Storage Foundation known issues

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

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

Upgrade issues

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

Because of a bug in the DDL provider package (`VRTSddlpr`) in the 5.0 release, the entries for the DDL provider are removed from the registry if the `rpm` command is used to upgrade to 5.0 MP2. This results in the Web GUI not being able to display any LUNs that are attached to the system.

The workaround is to specify the `--nopreun` option to the `rpm -U` command when upgrading the `VRTSddlpr` package.

The error does not occur if you use the `installmp` script to upgrade to 5.0 MP2.

[833516]

Upgrading the DCLI package

The `--nopreun` option should be specified if you use the `rpm -U` command to upgrade the distributed command line interface (`VRTSdcli`) package.

This issue does not apply if you use the `installmp` script to upgrade to 5.0 MP2.

Veritas Volume Manager known issues

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

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

The following sections contain information about known problems and issues in this release of VxVM.

Rootability issues

Encapsulated root disk cannot be mirrored

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

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

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

You can use the following workaround:

1. Unencapsulate the root disk:

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

   The system will reboot.

2. Remove the `/etc/vx/disk.info` file:

   ```
   # rm /etc/vx/disk.info
   ```

3. Shut down and reboot the system.

4. Use the `vxdiskadm` or `vxencap` command to re-encapsulate the root disk.

5. Use the `vxdiskadm` or `vxrootmir` command to remirror the root disk to the correct device.

You can avoid potentially misleading mapping between DMP node names and OS-based device names by turning off persistent device naming:

```
# vxddladm set namingscheme={ebn|osn} persistence=no
```

where `namingscheme=ebn` would select enclosure-based naming, and `namingscheme=osn` would select OS-based naming. [797829]

Utility issues

NFS cannot handle minor numbers greater than 255

The NFS implementation in Linux does not support minor numbers greater than 255 unless a patch is applied (see the description of Red HAT Bugzilla Bug 143897 or SUSE Bugzilla Bug 64552 for details). Without the patch, volume devices with large minor numbers cannot be remotely mounted via NFS. The workaround is to use the `vxdg` command to change the base minor number of the disk group that contains the volumes, as shown here:

```
# vxdg -g diskgroup reminor 2
```

DMP issues

Identification of ATA and SATA disks

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

To disable multipathing for ATA and SATA disks

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

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

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

[862137]

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 MP2 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 vxdmpadm 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, vxsited, has been implemented to provide automatic reattachment of sites. vxsited 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, vxsited 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, `vxsited` reattaches a site only when all the disks at that site become accessible. After reattachment succeeds, `vxsited` 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: `vxsited` does not try to reattach a site that you have explicitly detached by using the `vxdg detachsite` command.

The automatic site reattachment feature is enabled by default. The `vxsited` 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 `vxsited` in the `/etc/init.d/vxvm-recover` startup script, and reboot the system.

If you do not want a site to be recovered automatically, kill the `vxsited` 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 `vxsited`, and kill it by specifying its process ID:

```
    # kill -9 PID
```

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

To prevent the automatic site reattachment feature from being restarted, comment out the line that starts `vxsited` in the `/etc/init.d/vxvm-recover` startup script.

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 to 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]

Veritas Enterprise Administrator issues

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]

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 a includes a newly created volume:

```
Error: 0xcfff0021 Facility: 0xfff 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/283859

The following are new known issues in this MP2 release of Veritas File Sytem:

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

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

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

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

Veritas Storage Foundation for Oracle known issues

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

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

The following are new known issues in this MP2 release of Veritas Storage Foundation for Oracle:

Unsupported features

The following features of Veritas Storage Foundation for Oracle are not supported on the Oracle Enterprise Linux platform:

- mapping
- statistics
- GUI
- Web GUI

Commands fail on the OEL platform

The dbdst_file_move, dbdst_makelbfs and dbed_vmsnap commands fail on the OEL platform because the grep, mkfs and sed utilities are located in different directories than on the RHEL platform.

The workaround is to use the following commands to create symbolic links for the grep, mkfs and sed utilities:

```
# ln -s /bin/grep /usr/bin/grep
# ln -s /sbin/mkfs /usr/sbin/mkfs
# ln -s /bin/sed /usr/bin/sed
```

[1055246, 1055249, 1065248]

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.

Workaround: Fix the issue that caused the failover to the other node, and then fallback to the fixed node. [818522]

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. Stop the repository database and unmount the repository volume.

In a stand-alone configuration:

Stop the database repository:

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

Unmount the database repository:

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

In an HA configuration:	Stop VCS processes on either the local system or all systems.
	To stop VCS processes on the local system:
	`# hastop -local`
	To stop VCS processes on all systems:
	`# hastop -all`

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

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

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

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

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]

Software limitations

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

Security-enhanced Linux

Security-enhanced Linux (SELinux) support is provided for evaluation purposes only. Security policy files are not currently available for the Veritas product stack.

To disable SELinux at boot time on both SLES9 and RHEL4, set the kernel boot parameter `selinux` to `0` (`selinux=0`) and reboot the machine.

Assuming the system has been configured for booting from the machine *`machine_name`*, edit the file `/boot/`*`machine_name`*`/menu.lst` to include `selinux=0` on the kernel line. Then reboot the machine to ensure the setting takes effect.

Veritas File System software limitations

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

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

Veritas Storage Foundation for Oracle software limitations

The following are the software limitations for Veritas Storage Foundation for Oracle.

Some features stop working after a GCO failover

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

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

Differing locales produces unintelligible characters in GUI

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

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]

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]

Deep mapping unsupported

Deep mapping on EMC SYMCLI is currently not supported.

Use of buffered I/O

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

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

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

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]

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.

Documentation errata

Web GUI help errata

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

Veritas Storage Foundation Installation Guide errata

In several places in the *Veritas Storage Foundation Installation Guide*, the wrong path is given for the DVD-ROM device special file. The correct path is `/dev/cdrom`. [625481]

Veritas Storage Foundation Release Notes errata

In the *Veritas Storage Foundation 5.0 Release Notes*, it is stated that EFI disks are not supported. In fact, EFI disks are supported except for use as CDS disks or as encapsulated root disks. [914736]

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 V*eritas Storage Foundation for Oracle Administrator's Guide*:

- (Page 223) 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 127 and 266) 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
  ```