Veritas Storage Foundation™ Release Notes

AIX

5.0 Maintenance Pack 1

Veritas Storage Foundation Release Notes

Veritas Storage Foundation 5.0 MP1

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/283990

AIX is a registered trademark of IBM Corporation.

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

Chapter 1

Veritas Storage Foundation Release Notes

Introduction

This document provides release information on the following versions of the Veritas Storage Foundation Release 5.0 Maintenance Pack 1 (MP1) AIX product line:

- Storage Foundation Standard
- Storage Foundation Standard HA
- Storage Foundation Enterprise
- Storage Foundation Enterprise HA
- Storage Foundation for DB2 Standard
- Storage Foundation for DB2 Enterprise
- Storage Foundation for DB2 Enterprise HA
- Storage Foundation for Oracle Standard
- Storage Foundation for Oracle Enterprise
- Storage Foundation for Oracle Enterprise HA

All versions contain sets of Symantec products that can be activated by a single license key, or features installed with the product filesets and licensed separately.

Review this entire document before installing Veritas Storage Foundation components.

Note: For the latest information on updates, patches, and software issues regarding this release, refer to the following TechNote on the Symantec Technical Support website:
http://entsupport.symantec.com/docs/282024
The hardware compatibility list (HCL) is available at:
http://entsupport.symantec.com/docs/283161
The hardware TechNote is available at:
http://entsupport.symantec.com/docs/283282

Veritas Storage Foundation is a licensed product so you must obtain a license key prior to installation. For information on obtaining a license key, see the *Veritas Storage Foundation Installation Guide*.

Note: This document does not contain release Notes for Veritas Cluster Server. For release information on this product, see the *Veritas Cluster Server Release Notes*.

New features

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

Veritas Volume Manager

The following new features have been incorporated into Veritas Volume Manager.

Virtual SCSI device support

This release of Veritas Volume Manager includes support for virtual SCSI (vSCSI) devices.

Virtual SCSI devices are supported by default with DMP providing load balancing and failover. The `vxdmpvscsi` utility is provided to handle all operations that are specific to vSCSI devices, including enabling and disabling DMP on vSCSI devices, controlling the I/O policy, adding and removing supported arrays, and listing supported arrays.

The vxdmpvscsi command has the following usage:

```
# vxdmpvscsi addarraysupport VID offset length
# vxdmpvscsi disable
# vxdmpvscsi enable
# vxdmpvscsi iopolicy { A/P | A/P-LB }
# vxdmpvscsi listsupport
# vxdmpvscsi rmarraysupport VID offset length
```

The command keywords have the following effects:

addarraysupport	Adds support for using an array as a vSCSI device with DMP. By default, EMC, HITACHI and IBM arrays are supported as underlying devices. The required arguments are the vendor ID (VID) string for the array, the offset location in the inquiry page in hexadecimal, and the length of the VID string in hexadecimal. After using this command, run the rmdev and cfgmgr commands to remove and recreate the devices. The following are examples of using this command: # vxdmpvscsi addarraysupport **SHARK 08 04** # vxdmpvscsi addarraysupport **FUJITSU 08 07**
disable	Disables DMP for vSCSI devices. Once DMP support has been disabled, MPIO takes control of the devices, and implements multipathing features such as failover and load balancing; DMP acts in pass-through mode.
enable	Enables the use of DMP with vSCSI devices (this is the default behavior). Support for additional arrays can be added by using the addarraysupport operation.
iopolicy	Sets the I/O policy for vSCSI devices. The A/P policy makes the vSCSI devices behave like Active/Passive arrays. All the paths on one Virtual I/O (VIO) server are active, and the paths on the other VIO server acts in standby mode. At any point in time, I/O is allowed on only one path of the active VIO server. Load balancing of I/O on multiple paths at the same time is not supported. The A/P-LB policy configures LUN load-balancing for the vSCSI devices by dividing the LUN paths between two active VIO servers.
listarraysupport	Lists the arrays that are supported for use with vSCSI devices. By default, support for EMC, HITACHI and IBM arrays is enabled. Support for other arrays may be enabled by using the addarraysupport operation.

`rmarraysupport`	Removes support for using an array as a vSCSI device with DMP. See `addarraysupport` for a description of the arguments. After using this command, run the `rmdev` and `cfgmgr` commands to remove and recreate the devices.

Note the following limitations of support for vSCSI devices:

- This feature is only supported for AIX 5.3 TL5 with SP1. It is not supported for AIX 5.2 TL8.
- Cross-platform Data Sharing (CDS) (also known as Portable Data Containers (PDC)) is not supported for vSCSI devices. Such devices must be initialized as `aixdisk` format disks, and be added to non-CDS disk groups.
  By default, VxVM tries to initialize disks with the `cdsdisk` format. To avoid this, create the file `/etc/defaults/vxdisk`, and add the following line to this file:
  ```
  format=aixdisk
  ```
  To turn off CDS as the default for disk groups, create the file `/etc/defaults/vxdg`, and add the following line:
  ```
  cds=off
  ```
  With the defaults files set up in this way, VxVM will initialize disks with the `aixdisk` format in non-CDS disk groups by default.
- vSCSI devices do not support the SCSI-3 Persistent Group Reservation feature. This prevents SFRAC functionality from being used.

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.

AIX operating system requirements

The minimum system requirements for this release are:

- AIX 5.2 TL8
- AIX 5.3 TL5 with SP1

Install or migrate AIX 5.2 or 5.3, and then install the appropriate TL and SP before installing Veritas Storage Foundation 5.0 MP1.

Note: The installation script for Veritas Storage Foundation 5.0 MP1 verifies that the installation of AIX is at the proper TL before the upgrade to MP1 begins. Additional requirements for SP or other fixes are not verified by the install script. Refer to the following TechNote for any additional requirements that may exist:
http://entsupport.symantec.com/docs/282024

The script will abort without upgrading any component of Veritas Storage Foundation 5.0 MP1 if the target system fails to meet the AIX TL requirements.

Component product release notes

In addition to reading these Release Notes, review all component product release notes before installing the product. The following component product release notes are included as PDF files on the Veritas software disc:

Veritas Storage Foundation Release Notes (`sf_notes.pdf`)

Veritas Cluster Server Release Notes (`vcs_notes.pdf`)

Because release notes are not installed by any packages and are not on the documentation disc, it is recommended that you copy them from the software disc to the `/opt/VRTS`*`product_name`*`/docs` directory so they are available on your system for reference.

Migrating from AIX 5.2 to AIX 5.3

Veritas Storage Foundation 5.0 MP1 will not operate properly on AIX 5.3 unless AIX 5.3 is at the required TL or the applicable SPs are installed.

See "System requirements" on page 11 for the minimum required TL and SP of AIX 5.3 for Storage Foundation.

Installing for the first time

For new installations, first install Veritas Storage Foundation 5.0, then upgrade to the MP1 level. Review and follow the installation procedure for Storage Foundation 5.0 described in the *Veritas Storage Foundation 5.0 Installation Guide*. Follow the procedures in this document to upgrade Storage Foundation to 5.0 MP1.

Installing Storage Foundation 5.0 MP1

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

- "Installing Storage Foundation on a cluster" on page 12.
- "Installing 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 install Storage Foundation 5.0 MP1 on a cluster

1. Log in as superuser.
2. Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/filesystems` file. You will need to recreate these entries in the `/etc/filesystems` file on the freshly installed system.
3. Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in such as `/etc/filesystems`. 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.

4 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the installation.

5 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 systems.

6 If you are installing 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 installation procedure.

7 Perform any necessary preinstallation checks and configuration.
See the *Veritas Storage Foundation 5.0 Installation Guide*.

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

```
# mount | grep vxfs
```

9 Unmount all Storage Checkpoints and Veritas File Systems:

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

10 Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | /opt/VRTSvxfs/sbin/fsdb -V 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 -V vxfs filesystem
# mount -V 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 10 to verify that the unclean file system is now clean.

11 If there are still diskgroups that are imported at this time then proceed to step 12, otherwise skip to step 15.

12 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps. Otherwise, proceed to step 13 on page 14.

a Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

b Use the `vxrvg stop` command to stop each RVG individually:

```
# vxrvg -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 13.

```
# vxvoltune vol_vvr_use_host_byte_order 1
```

13 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.

14 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
```

15 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
```

16 Insert the disc containing the Veritas software into the DVD-ROM drive, and mount the disc on a suitable mount point. For example:

```
# mount -V cdrfs -o ro /dev/cd0 /mnt/cdrom
```

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

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

18 To install 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
  ```

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

20 After the requirement checks have completed successfully, press **Enter** to begin installing the packages.

Note: As you are installing multiple systems in a cluster, select to install the systems simultaneously.

---

21 After the installation of the packages is complete, use the following command to shut down the cluster nodes:

```
# shutdown -r now
```

22 Reboot the cluster nodes, and reinstate any missing mount points in the `/etc/filesystems` file.

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

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

24 Restart each RVG (if any were stopped in step 12):

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

25 Remount all Veritas File Systems and Storage Checkpoints:

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

26 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
```

27 If you set the value of the `vol_vvr_use_host_byte_order` tunable to 1 in step d of step 12, reboot the system.

Note: When the installer installs software, some software may be applied rather than committed. It is the responsibility of the system administrator to commit the software, which can be performed later using the `-c` option of the `installmp` command.

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.
- Restore any VCS configuration files as described in the *Veritas Cluster Server 5.0 Installation Guide* and *Veritas Cluster Server 5.0 Release Notes*.

Installing Storage Foundation on a standalone system

Caution: Existing data could be destroyed on any disks that are touched by upgrading the operating system. While installing, 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 installation.

To install Storage Foundation 5.0 MP1 on a standalone system

1 Log in as superuser.

2 Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/filesystems` file. You will need to recreate these entries in the `/etc/filesystems` file on the freshly installed system.

3 Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in such as `/etc/filesystems`. 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.

4 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the installation.

5 Perform any necessary preinstallation checks and configuration. See the *Veritas Storage Foundation 5.0 Installation Guide* for more information.

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

```
# mount | grep vxfs
```

7 Unmount all Storage Checkpoints and Veritas File Systems:

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

8 Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | /opt/VRTSvxfs/sbin/fsdb -V 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 -V vxfs filesystem
# mount -V 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 8 to verify that the unclean file system is now clean.

9 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps. Otherwise, proceed to step 10 on page 18.

a Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

b Use the `vxrvg stop` command to stop each RVG individually:

```
# vxrvg -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 10.

```
# vxvoltune vol_vvr_use_host_byte_order 1
```

10 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.

11 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
```

12 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
```

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 -V cdrfs -o ro /dev/cd0 /mnt/cdrom
```

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

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

15 To install 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 installation.

17 After the requirement checks have completed successfully, press **Enter** to begin installing the packages.

Note: If you are installing on multiple standalone systems, you can choose to install on the systems simultaneously.

---

18 After the installation of the packages is complete, use the following command to shut down the system:

```
# shutdown -r now
```

19 Reboot the system, and reinstate any missing mount points in the `/etc/filesystems` 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 9):

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

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
```

24 If you set the value of the `vol_vvr_use_host_byte_order` tunable to 1 in step d of step 9, reboot the system.

Note: When the installer installs software, some software may be applied rather than committed. It is the responsibility of the system administrator to commit the software, which can be performed later using the `-c` option of the `installmp` command.

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.

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.[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
```

Migrating an internal root disk to a SAN root disk under DMP control

If the system has been booted from an internal disk (such as `hdisk0`), an alternate root disk can be configured on the attached SAN storage before putting it under DMP control.

To migrate a root disk from LVM to DMP control

1. Use the `chdev` command to clear the PVID of the disk in the SAN storage that is to be the SAN root disk:
   ```
   # chdev -l hdisk373 -a pv=clear
   ```
   If the SAN disk is already multipathed by DMP, this command needs to be run for all of the paths to the disk.
2. Use the `alt_disk_install` command to create the SAN root disk:
   ```
   # alt_disk_install -C hdisk373
   ```
3. Reboot the system from the SAN root disk.
4. To configure DMP, follow the steps in "Configuring DMP support for booting over a SAN" on page 23.

See "Example of creating a SAN boot disk under DMP control" on page 24.

Migrating a SAN root disk to DMP control

If the system has been booted from a SAN disk under MPIO control, MPIO must be disabled before DMP control can be enabled.

To migrate a SAN root disk from MPIO to DMP control

1. If the disk is not already under MPIO control, skip to step 3.
   Disable MPIO by installing a device-specific ODM definition package as described in the following TechNote:
   http://library.veritas.com/docs/263558
2. Reboot the system. The system is booted without any multipathing support.
3. To configure DMP, follow the steps in "Configuring DMP support for booting over a SAN" on page 23.

Configuring DMP support for booting over a SAN

Configuring DMP to work with a LVM root disk over a SAN requires that the system is correctly configured to use the boot device over all possible paths.

To configure DMP support for booting over a SAN

1. If the LVM root disk is multipathed with 4 paths, the output from the `lspv` command output for the root volume group (`rootvg`) is similar to the following:

```
# lspv | grep rootvg
hdisk374   00cbf5ce56def54d rootvg   active
hdisk375   00cbf5ce56def54d rootvg
hdisk376   00cbf5ce56def54d rootvg
hdisk377   00cbf5ce56def54d rootvg
```

   The PVID and volume group entries in the second and third columns should be identical for all the paths.

2. If the PVID and volume group entries are not set correctly on one or more of the paths, use the `chdev` command to set the correct value, as shown in this example:

```
# chdev -l hdisk377 -a pv=clear
hdisk377 changed
# chdev -l hdisk377 -a pv=yes
hdisk377 changed
```

3. Check that the output from the `bootlist` command is correct:

```
# bootlist -m normal -o
hdisk373 blv=hd5
hdisk372 blv=hd5
hdisk371 blv=hd5
hdisk370 blv=hd5
```

   In this example, the default boot volume, `hd5`, should be shown for each path.

4. If the `blv` option is not set for a path to the disk, use the `bootlist` command to set it:

```
# bosboot -ad
# bootlist -m normal hdisk373 blv=hd5
```

5. Run the `vxdmproot` command to configure DMP on the root disk:

```
# vxdmproot install
```

6. Reboot the system. DMP takes control of the SAN boot device in order to perform load balancing and failover.

7. To verify whether DMP controls the root disk, after reboot, the major number of the root disk (for example, hdisk373) should be the same as the

major number of DMP configured device (in this example, /dev/vx/dmpconfig.

```
# ls -l /dev/`bootinfo -b` /dev/vx/dmpconfig
crw--- 1    root    system      42,65535    Nov 24      14:25
/dev/vx/dmpconfig
brw--- 1    root    system      42,    1    Sep 07      16:16
/dev/hdisk373
```

Mirroring the DMP controlled rootdisk

If the rootdisk is mirrored, then run the following command to add the new mirror rootdisk to the DMP. This command enables the DMP on mirror rootdisk when the mirror is booted on mirror rootdisk.

```
# vxdmproot add <mirror_rootdisk>
```

Example of creating a SAN boot disk under DMP control

In this example, a SAN boot disk with multiple paths is created by cloning the existing root disk, and then enabling multipath support by DMP.

1. The existing root disk is hdisk125.

```
# lspv | grep rootvg
hdisk125        00cdee4fd0e3b3da                    rootvg
active
# lsvg
rootvg
```

2. The PVIDs of all the paths to the SAN boot disk are cleared:

```
# chdev -l hdisk542 -a pv=clear
hdisk542 changed
# chdev -l hdisk557 -a pv=clear
hdisk557 changed
# chdev -l hdisk558 -a pv=clear
hdisk558 changed
# chdev -l hdisk559 -a pv=clear
hdisk559 changed
```

3. The `alt_disk_install` command is used to create the SAN boot disk:

```
# alt_disk_install -C -P all hdisk542
+-------------------------------------------------------------+
ATTENTION: calling new module /usr/sbin/alt_disk_copy. Please
see the
alt_disk_copy man page and documentation for more details.
Executing command: /usr/sbin/alt_disk_copy -P "all" -d
"hdisk542"
+-------------------------------------------------------------+
Calling mkszfile to create new /image.data file.
```

```
Checking disk sizes.
Creating cloned rootvg volume group and associated logical
volumes.
Creating logical volume alt_hd5.
Creating logical volume alt_hd6.
Creating logical volume alt_hd8.
Creating logical volume alt_hd4.
Creating logical volume alt_hd2.
Creating logical volume alt_hd9var.
Creating logical volume alt_hd3.
Creating logical volume alt_hd1.
Creating logical volume alt_hd10opt.
Creating logical volume alt_lg_dumplv.
Creating /alt_inst/ file system.
Creating /alt_inst/home file system.
Creating /alt_inst/opt file system.
Creating /alt_inst/tmp file system.
Creating /alt_inst/usr file system.
Creating /alt_inst/var file system.
Generating a list of files
for backup and restore into the alternate file system...
Backing-up the rootvg files and restoring them to the alternate
file system...
Modifying ODM on cloned disk.
Building boot image on cloned disk.
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst
forced unmount of /alt_inst
Changing logical volume names in volume group descriptor area.
Fixing LV control blocks...
Fixing file system superblocks...
Bootlist is set to the boot disk: hdisk542
```

4 The `lspv` command is used to confirm that `altinst_rootvg` has been created for one of the paths to the SAN disk:

```
# lspv | grep rootvg
hdisk125   00cdee4fd0e3b3da   rootvg     active
hdisk542   00cdee4f5b103e98   altinst_rootvg
```

5 The remaining paths to the SAN disk are updated to include the correct `altinst_rootvg` information:

```
# chdev -l hdisk557 -a pv=clear
hdisk557 changed
# chdev -l hdisk557 -a pv=yes
hdisk557 changed
# chdev -l hdisk558 -a pv=clear
hdisk558 changed
# chdev -l hdisk558 -a pv=yes
```

```
hdisk558 changed
# chdev -l hdisk559 -a pv=clear
hdisk559 changed
# chdev -l hdisk559 -a pv=yes
hdisk559 changed
# lspv | grep rootvg
hdisk125   00cdee4fd0e3b3da   rootvg     active
hdisk542   00cdee4f5b103e98   altinst_rootvg
hdisk557   00cdee4f5b103e98   altinst_rootvg
hdisk558   00cdee4f5b103e98   altinst_rootvg
hdisk559   00cdee4f5b103e98   altinst_rootvg
```

6 The `bootlist` command verifies that the boot device has been updated for only one of the paths to the SAN disk:

```
# bootlist -m normal -o
hdisk542 blv=hd5
```

7 The `bootlist` command is used to include the other paths to the new boot device:

```
# bootlist -m normal hdisk557 blv=hd5
# bootlist -m normal hdisk558 blv=hd5
# bootlist -m normal hdisk559 blv=hd5
# bootlist -m normal -o
hdisk542 blv=hd5
hdisk557 blv=hd5
hdisk558 blv=hd5
hdisk559 blv=hd5
```

8 Run the following command to make sure all the above disks are bootable.

```
# pl_varyon -i | egrep -w "hdisk542|hdisk557|hdisk559"

hdisk542 YES   oocdee4fd0e3b3da00000000000000 00cbf5ce00004c00
hdisk557 YES   oocdee4f5b103e9800000000000000 00cbf5ce00004c00
hdisk558 YES   oocdee4f5b103e9800000000000000 00cbf5ce00004c00
hdisk559 YES   oocdee4f5b103e9800000000000000 00cbf5ce00004c00
```

9 DMP can now be enabled:

```
# vxdmproot install
```

10 After rebooting the system to enable DMP support for LVM bootablity, the following commands confirm that the system is booted from the new multipathed SAN disk:

```
# bootinfo -b
hdisk542
# bootlist -m normal -o
hdisk542 blv=hd5
hdisk557 blv=hd5
hdisk558 blv=hd5
hdisk559 blv=hd5
# lspv | grep rootvg
hdisk125   00cdee4fd0e3b3da   old_rootvg
hdisk542   00cdee4f5b103e98   rootvg     active
hdisk557   00cdee4f5b103e98   rootvg     active
hdisk558   00cdee4f5b103e98   rootvg     active
hdisk559   00cdee4f5b103e98   rootvg     active
```

Software limitations

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

Veritas Storage Foundation software limitations

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

http://entsupport.symantec.com/docs/283990

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

Advanced POWER Virtualization (APV)

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

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

Veritas File System software limitations

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

http://entsupport.symantec.com/docs/283990

The following are new software limitations in this MP1 release of Veritas File System:

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

The VX_CONCURRENT advisory cannot be set on a file that is actively open by Quick I/O or ODM, nor can that file be concurrently opened with the O_CIO flag. A file opened with the O_CIO flag or that has the VX_CONCURRENT advisory set may not be concurrently opened by Quick I/O or ODM. Quick I/O and ODM access are not allowed for any files on a file system that is mounted with the -o cio mount option.

Veritas Storage Foundation for Oracle software limitations

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

http://entsupport.symantec.com/docs/283990

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

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

Since DNS hostname lookup queries are, by definition, case insensitive, make sure the SFDB repository is running on a host with a name that is truly unique -- regardless of case -- within the local subnet. Errors may occur if the repository hostname differs from another hostname only by case. [851129]

Veritas Storage Foundation for DB2 software limitations

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

http://entsupport.symantec.com/docs/283990

The following are new software limitations in this MP1 release of Veritas Storage Foundation for DB2:

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]

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]

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

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

[643964]

Repository hostnames are case insensitive

Since DNS hostname lookup queries are, by definition, case insensitive, make sure the SFDB repository is running on a host with a name that is truly unique -- regardless of case -- within the local subnet. Errors may occur if the repository hostname differs from another hostname only by case. [851129]

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 table contains information about fixed issues in the Veritas Storage Foundation 5.0 MP1 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 `vxmake` command could not be used to recreate site records. This is now supported if the `-d` option is used to read from a description file.
601274	In a CVM cluster, DMP did not fail over to a secondary path when the primary paths were disconnected.
603522	If a VPATH metadevice was opened with the O_NDELAY flag, ioctls or I/O to the device failed with ENXIO.
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.

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

Veritas Enterprise Administrator fixed issues

The following issues have been fixed in the Veritas Storage Foundation 5.0 MP1 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 the Veritas Storage Foundation 5.0 MP1 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.

Veritas File System fixed issues

The following table contains information about fixed issues in the Veritas Storage Foundation 5.0 MP1 release of VxFS.

Incident	Description
596307	Dramatically decreased the time required to freeze file systems during a dynamic reconfiguration. This is a performance enhancement.
616323	For WebGUI online help, the following issues have been fixed: For the **Remount Storage Checkpoint** operation, the **More info** link on the second wizard page does not function properly for cluster file systems. For the **Unmount Storage Checkpoint** operation, the **More info** link on the second wizard page does not function properly for cluster file systems.
770916	Fixed the problem in which `vxfsconvert` failed to convert the file system if it was invoked with `-s` *`estimated_value`* option, and the estimated value happened to be an exactly the required space for conversion and not a multiple of 8KB.
770917	Inode ownership issues detected in large directory related code paths have been fixed.
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`.

Incident	Description
770964	`fsck` has been enhanced to replay file systems created with earlier log versions on volume sets.
771085	An uninitialized minor device number in `vxuncfg` command could cause a system panic. This was fixed by initializing the variable before making system calls.
771086	Fixed an `fsck` problem in which users could end up creating multiple `lost+found` directories when running the `fsck -o full` command and answering `fsck` questions interactively. Now, `fsck` creates only one. It also checks for multiple `lost+found` entries and removes duplicate directory entries.
771996	Enhanced VxFS to use less CPU when doing administrative tasks on the devices of multi-volume file systems.
772013	Enhanced the `fsck` command to enforce the `lost+found` file name in the root directory of the file system to be a directory file type.
777012	If the system crashed or there was a metadata I/O error, after the `fsadm` command reorganized the `lost+found` directory, running the `fsck -o full` command may not have been able to clean the file system with regard to names that needed to be added to the `lost+found` directory. The problem happened on single-volume and multi-volume file systems.
785649	A situation where `vxfsconvert` of a dusty file system loops forever in user-level code when an inode with pending truncation operation is encountered has been fixed.
793022	`vxfs_nattr_open`() API interface has been fixed to shrink files, as appropriate, when invoked with `O_TRUNC` flag.
793030	`vxfsutil.h` uses `struct timeval` in one of the function declarations, but does not include `time.h`. This causes user applications to report warnings during compilation. This issue has been fixed by including `time.h` in `vxfsutil.h`.
795073	The increased CPU utilization when writing to a file system that is almost full due to more background processing threads than are actually required being enqueued has been fixed.

Veritas 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.
859597	A volume could be resized incorrectly and data could be lost if the requested size contained a decimal point.

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.

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

Veritas Storage Foundation for Databases fixed issues

Veritas Storage Foundation for Oracle fixed issues

The following table contains a list of issues fixed in the Veritas Storage Foundation 5.0 MP1 release for Storage Foundation for Oracle:

Incident	Description
567342	An unmounted checkpoint clone database reappears in the Java GUI tree after rescanning.
582416	Clicking the **Help** button on a GUI wizard produces the following error message: `Error V-39-53246-8 Get EntryPoint failed. Please check the manifest related information`
582069	If SFDB commands are executed with a locale that differs from the locale in use when the SFDB server was started, the commands may fail with the following message: `([Sybase][ODBC Driver][Adaptive Server Anywhere]Syntax error ).`
600431	Storage Checkpoint operations are currently not supported for databases cloned with Database FlashSnap.
608697	The Web GUI statistic scheduler uses the client timestamp instead of the server timestamp which causes the first statistic collection to be skipped. When using the Web GUI with the Mozilla Firefox browser, the Refresh command in the View Statistics wizard does not work.
613681	The `sfua_db_config -o startdb` command does not print a message to indicate whether the SFDB repository database was successfully started or not.

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 output from the `-l` option in the previous step.

Veritas Storage Foundation for DB2 fixed issues

The following table contains a list of issues fixed in the Veritas Storage Foundation 5.0 MP1 release for Storage Foundation for DB2:

Incident	Description
567342	An unmounted checkpoint clone database reappears in the Java GUI tree after rescanning.
582416	Clicking the **Help** button on a GUI wizard produces the following error message: `Error V-39-53246-8 Get EntryPoint failed. Please check the manifest related information`
582069	If SFDB commands are executed with a locale that differs from the locale in use when the SFDB server was started, the commands may fail with the following message: `([Sybase][ODBC Driver][Adaptive Server Anywhere]Syntax error ).`
608697	The Web GUI statistic scheduler uses the client timestamp instead of the server timestamp which causes the first statistic collection to be skipped. When using the Web GUI with the Mozilla Firefox browser, the Refresh command in the View Statistics wizard does not work.

Incident	Description
600490, 600492	The `db2ed_clonedb` and `db2ed_vmclonedb` commands do not support automatic storage databases. Using either command on an automatic storage database results in the following message: DBT1037N: CONT_PATH entries cannot be specified for an automatic storage table space
613681	The `sfua_db_config -o startdb` command does not print a message to indicate whether the SFDB repository database was successfully started or not.

Known issues

Storage Foundation known issues

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

http://entsupport.symantec.com/docs/283990

Veritas Volume Manager known issues

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

http://entsupport.symantec.com/docs/283990

The following incidents were reported for this release of VxVM.

Utility issues

Conversion of an LVM volume fails

Using the `vxconvert` command to convert an LVM volume to a VxVM volume can fail under the following circumstances:

- The disk access name, as shown by the `vxdisk list` command, does not indicate that the device belongs to LVM. The workaround is to run the `vxdisk scandisks` command before proceeding with the conversion.
- A previous conversion failed, and an LVM disk has been left with a VxVM disk label. The workaround is to run the `vxdiskunsetup` command on the disk before proceeding with the conversion.

[898339]

DMP issues

Limitations of DMP support for LVM root disks

DMP support for LVM root disks has the following limitations:

- DMP can be enabled only on the boot disk. The root volumes on disks other than boot disks cannot be controlled by DMP. In addition, the mirror disk cannot be controlled by DMP. However, when you reboot on the mirrored disk, it will come up with DMP enabled.
- The `vxdmproot add` and `vxdmproot remove` commands cannot be applied on non-bootable boot disks. In other words, these commands can only be applied to the mirror rootdisks.

- DMP may not operate correctly if Dynamic Reconfiguration (DR) or similar operations are performed that shuffle device numbers or device names. If such a reconfiguration occurs, use the following commands to reconfigure DMP before the system is next rebooted.

  ```
  # cfgmgr
  # vxdmproot install
  ```

PowerPath disks and vxconvert

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

The problem of having an entry in the `/etc/vx/darecs` file is because the PowerPath disks were configured as foreign disks in Storage Foundation 4.0, and this was not changed after subsequent upgrades. [857504]

Workaround:

1. Remove the PowerPath disk configured as simple using the following command:

   ```
   # vxdisk rm daname
   ```

2. Rediscover the PowerPath disk as an aixdisk instead of as a simple/foreign disk by running the following command:

   ```
   # vxdisk scandisks
   ```

3. Rerun the `vxconvert` command.

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 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 Veritas File System 5.0 release are listed in the Veritas Storage Foundation 5.0 Release Notes which is available at the following URL:

http://entsupport.symantec.com/docs/283990

The following are new known issues in this MP1 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 Volume Replicator known issues

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

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

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

Synchronizing volumes and RVG with large volumes

The `vradmin syncrvg` and the `vradmin syncvol` commands do not work correctly for volumes larger than 1TB. When either of these two commands is used to synchronize large volumes, the command still runs, but it reports wrong total size of the volumes being synchonized 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 coredump 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 Veritas Storage Foundation for Oracle 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

http://entsupport.symantec.com/docs/283990

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

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

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

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

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

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]

Error message displayed after successfully creating a new repository

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

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

You can safely ignore the message. [830629]

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

Note: The primary instance's pfile is saved at the time of snapshot creation. If you attempt to clone the database using that snapshot you will be using the saved pfile, not the current pfile. Therefore you must create a new snapshot in order to ensure that the clone will use an updated pfile.[852188]

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.

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 Veritas Storage Foundation for DB2 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

http://entsupport.symantec.com/docs/283990

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

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

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

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

To uninstall the Veritas Storage Foundation for DB2 packages

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

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

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]

`db2ed_vmsnap` gives error: unable to duplicate owner and mode after move

This move error message can be ignored. `/tmp` has gid bit set so the snapshot completed successfully. We will suppress or remove this warning in a later release.[841766]

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]

Documentation errata

Web GUI help errata

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

Manual pages errata

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 Storage Foundation Installation Guide errata

The following errata apply to the V*eritas Storage Foundation Installation Guide*:

- (Page 76) The `uninstall` command for the Veritas Storage Foundation for Oracle package should read as:

  `# /opt/VRTS/install/uninstallsfora`
- (Page 85) In Table B-1, which lists the Storage Foundation packages and their details, the `VRTScmccc` package is an optional package.

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 errata apply to the V*eritas Storage Foundation for Oracle Administrator's Guide*:

- (Page 291) 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 181 and 345) 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.

The following information is 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 ninja
$ srvctl config database -d KPRDADV1
ninja KPADV1R1 /apps/oracle/product/920rac
```

Then add the second instance:

```
$ srvctl add instance -d KPRDADV1 -i KPADV1R2 -n tyrolia
$ srvctl config database -d KPRDADV1
ninja KPADV1R1 /apps/oracle/product/920rac
tyrolia 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 ninja
Instance KPADV1R2 is running on node tyrolia
```

Reconfiguring virtual IP address for repository configuration

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

For a standalone instance:

1 Change the IP address.

2 Update the new IP information for SFUA repository access by running the `sfua_db_config` command.

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

For a cluster environment:

1 Change the IP address for the cluster.

2 Update the IP address for the repository configuration in HA environment by running the following set of commands:

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.

Veritas Storage Foundation for DB2 Administrator's Guide errata

The following errata apply to the V*eritas Storage Foundation for DB2 Administrator's Guide*:

- (Page 241) In the procedure "To abort reverse resynchronisation", the text in the first bullet should read as follows:
  Use the `-o reverse_resync_abort` option of the `db2ed_vmsnap` command.
- (Pages 149 and 295) 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.

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

Reconfiguring virtual IP address for repository configuration

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

For a standalone instance:

1. Change the IP address.
2. Update the new IP information for SFUA repository access by running the `sfua_db_config` command.

   `# /opt/VRTSdbcom/bin/sfua_db_config`

For a cluster environment:

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

   a. Unconfigure the SFUA repository from the VCS configuration:

   `# /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.

Veritas Storage Foundation for Oracle Graphical User Interface Guide errata

The following errata apply to the V*eritas Storage Foundation for Oracle Graphical User Interface Guide*:

- (Page 23) In step 2 of the procedure “To Start a DBEDAgent”, the correct command to start a DBEDAgent is:
  `/etc/rc2.d/S75vxpal.DBEDAgent start`
- (Page 23) In the procedure, “To Stop a DBEDAgent”, the correct command to stop a DBEDAgent is:
  `/etc/rc2.d/S75vxpal.DBEDAgent stop`