

# Veritas Storage Foundation™ Release Notes

HP-UX

5.0 Maintenance Pack 1



# Veritas Storage Foundation Release Notes

Copyright © 2007 Symantec Corporation. All rights reserved.

Storage Foundation 5.0 Maintenance Pack 1

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

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

THIS DOCUMENTATION IS PROVIDED “AS IS” AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID, SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation  
20330 Stevens Creek Blvd.  
Cupertino, CA 95014  
[www.symantec.com](http://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/283708>

HP-UX is a registered trademark of Hewlett-Packard Development Company, L.P.

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



# Contents

Introduction .....	7
System requirements .....	9
HP-UX operating system requirements .....	9
Hardware requirements .....	9
Supported Oracle and HP-UX combinations .....	9
Mandatory patch required for Oracle Bug 4130116 .....	9
Hard disk array support .....	10
Component product release notes .....	10
Installing for the first time .....	11
Installing Storage Foundation 5.0 MP1 .....	11
Changing permissions for Storage Foundation for Oracle .....	12
Removing patches .....	12
Software limitations .....	14
Storage Foundation software limitations .....	14
Veritas File System software limitations .....	14
Veritas Storage Foundation for Oracle software limitations .....	14
No longer supported .....	19
Fixed issues .....	20
Veritas Storage Foundation fixed issues .....	20
Veritas Volume Manager fixed issues .....	20
Veritas File System fixed issues .....	23
Veritas Volume Replicator fixed issues .....	25
Veritas Storage Foundation for Oracle fixed issues .....	27
Known issues .....	29
Veritas Storage Foundation known issues .....	29
Veritas Storage Foundation for Oracle known issues .....	31
Veritas Volume Manager known issues .....	34
Veritas File System known issues .....	39
Veritas Volume Replicator known issues .....	39
Documentation errata .....	41
Veritas Storage Foundation Release Notes errata .....	41
Web GUI help errata .....	41
Manual pages errata .....	41
Veritas Volume Manager Administrator's Guide errata .....	42
Veritas Storage Foundation for Oracle Administrator's Guide errata .....	42
Veritas Storage Foundation for Oracle Graphical User Interface Guide errata .....	45



# Veritas Storage Foundation Release Notes

## Introduction

This document provides release information about the products in the Veritas Storage Foundation 5.0 Maintenance Pack 1 (MP1) HP-UX product suite:

- Veritas Storage Foundation (Standard, Standard HA, Enterprise, and Enterprise HA)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)
- Veritas Storage Foundation for Oracle (Standard, Enterprise, and HA Editions), formerly known as Veritas Database Edition for Oracle.

---

**Note:** Veritas Storage Foundation QuickStart is not available in this release.

---

Each of these products is activated by a single license key. You must obtain a license key before installing the product.

See the *Veritas Storage Foundation Installation Guide*.

---

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

<http://entsupport.symantec.com/docs/281875>

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

---

This document does not contain release information for Veritas Cluster Server. See the *Veritas Cluster Server Release Notes*.



# System requirements

## HP-UX operating system requirements

The system requirements for this release are:

- September 2004 release of HP-UX 11i version 2.0 or later.

### To verify the operating system version

- ◆ Use the `swlist` command as follows:

```
# swlist | grep HPUXBaseAux
HPUXBaseAux                B.11.23.0409 HP-UX Base OS Auxiliary
```

JFS must be installed on your system prior to installing any Veritas software.

### To verify that JFS is installed

- ◆ Use the `swlist` command as follows:

```
# swlist -l product JFS
JFS B.11.23                The Base VxFS File System
```

## Hardware requirements

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>

## Supported Oracle and HP-UX combinations

The following list identifies the supported Oracle and HP-UX combinations if you plan to use Veritas Storage Foundation with an Oracle database:

Oracle Release	September 2004 HP-UX 11i version 2.0 or later
9.2	Yes
10.1	Yes
10.2	Yes

## Mandatory patch required for Oracle Bug 4130116

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

## Hard disk array support

### Required patches

To enable hard disk array support on Veritas Storage Foundation for Oracle, you must install the following patches to September 2004 HP-UX 11i version 2.0 (update 2):

- PHSS\_32228
- PHSS\_32231

---

**Note:** These patches apply only to the IA architecture.

---

### Disk array support for Storage Mapping

The EMC 8000 disk array, which supported the Storage Mapping feature in the previous release, is not yet confirmed for operation with the 5.0 MP1 release. The hardware compatibility list (HCL) will be updated when support becomes available.

## Component product release notes

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

*Veritas Storage Foundation Release Notes (sf\_notes.pdf)*

*Veritas Cluster Server Release Notes (vcs\_notes.pdf)*

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

## Installing for the first time

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

## Installing Storage Foundation 5.0 MP1

Install the patches using the `installmp` script. You must have superuser (`root`) privileges.

### To install the patches using `installmp`

- 1 Verify the status of the Veritas Enterprise Administrator (VEA) Service:  

```
# /opt/VRTS/bin/vxsvcctl status  
Current state of server : RUNNING
```
- 2 Stop the VEA server:  

```
# /opt/VRTS/bin/vxsvcctl stop  
DBED: Successfully unloaded the Storage Foundation Provider  
4.0 for Oracle
```
- 3 Verify the status of the VEA Service again:  

```
# /opt/VRTS/bin/vxsvcctl status  
Current state of server : NOT RUNNING
```
- 4 Insert the patch disc into the DVD-ROM drive.
- 5 Determine the block device file for the DVD-ROM drive by entering:  

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.
- 6 Mount the software disc. For example, to mount the patch disc at the mount point `/dvdrom`, enter:  

```
# /usr/sbin/mount -F cdfs /dev/dsk/c#t#d# /dvdrom
```

where `/dev/dsk/c#t#d#` is the location of the DVD-ROM drive.
- 7 Go to the DVD-ROM file system:  

```
# cd /dvdrom
```
- 8 Install the patches using the `installmp` command.  

```
# ./installmp
```
- 9 Reboot the system:  

```
# /usr/sbin/shutdown -r now
```

## Changing permissions for Storage Foundation for Oracle

After installing the Veritas Storage Foundation 5.0 MP1 patches, follow these post-installation steps to ensure Veritas Storage Foundation for Oracle commands work correctly. [772592]

---

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

---

### To change permissions

- 1 Change permissions for the `VRTSdbed` directory:  

```
# 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
```

## Removing patches

The following procedure removes the patches if you want to uninstall MP1.

### To remove the patches

- 1 Log in as superuser (`root`).
- 2 If you are running a Storage Foundation HA cluster, stop the cluster:  

```
# hastop -all
```
- 3 If you are running a Storage Foundation HA cluster, stop the VxFEN processes on all nodes:  

```
# /sbin/init.d/vxfen stop
```

- 4 On all nodes, stop the Storage Agents:  

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c stop
```
- 5 Check the status of the Storage Agents:  

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c status
```
- 6 On all nodes, remove all of the patches using the `swremove` command:  

```
# swremove -x autoreboot=true patch_name1, patch_name2 ...
```

Symantec recommends that all the patches installed for 5.0 MP1 be removed through a single command line. The system automatically reboots after removing the patches.

---

**Caution:** The patches `PHKL_35178`, `PHCO_35179`, and `PVCO_03696` should always be installed or removed with a single command line. Not installing or removing them together can result in data loss.

---

## Software limitations

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

### 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/283708>

### 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/283708>

### Veritas Storage Foundation for Oracle software limitations

#### **DBDST limitations with non-English filenames and placement class names**

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

#### **Differing locales produces unintelligible characters in GUI**

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

#### **Some features stop working after a GCO failover**

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

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

### Avoid using UNIX VEA via PC-based UNIX emulators

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

### No support for Intelligent Storage Provisioning

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

### Disk layouts Version 5 and earlier do not display Storage Checkpoint quotas in the GUI

You can click the **Quota** tab for the Storage Checkpoint available via the GUI. If the file system and the Storage Checkpoint are on a version prior to Disk Layout Version 6, you will see the error 4646.

For Disk Layout Version 5 and earlier to display Storage Checkpoints Quotas via the GUI, perform the following steps:

- 1 Using Veritas File System 3.5, create a file system.
- 2 Upgrade to Veritas File System 5.0 but *do not* upgrade the file system.
- 3 Create a checkpoint and then click the **Quota** tab accessible via the GUI.

### Storage Checkpoint limitations

- You cannot create a clone database using a mounted Storage Checkpoint. [32726]
- You must run the `dbed_update` command after upgrading to Veritas Storage Foundation 4.1 for Oracle from a previous release. This will allow you to roll back to a Storage Checkpoint that was created prior to this release. [86431]
- If you create an Oracle instance using the `spfile` option, you must run the `dbed_update` command before you can successfully perform any Storage Checkpoint or Database FlashSnap functions.

## VEA limitations

- Veritas Enterprise Administrator (VEA) does not display tablespace information when the `v$table` column names are changed using the SQL\*Plus profile facility. Normally this happens when SQL\*Plus settings are used in `login.sql` to change column names in reports. [34446]
- VEA may display system fonts incorrectly. On a Japanese desktop, VEA may incorrectly display system fonts. Japanese characters may not be properly displayed when you select the non-default font for the VEA GUI.

## Database FlashSnap limitations

- The Database FlashSnap feature does not support RAID-5 volumes. [34570]
- When cloning a database using Database FlashSnap, the Oracle database must have at least one mandatory archive destination, otherwise `dbed_vmchecksnap` results in this error message:  
SFORA dbed\_vmchecksnap ERROR V-81-5677 Could not find a mandatory, primary and valid archive destination for database PROD.  
Please review the `LOG_ARCHIVE_DEST_n` parameters and check `v$archive_dest`.  
This example shows how to establish a mandatory archive destination using SQL\*Plus:  

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

  
For more information about Oracle parameters for archiving redo logs, see your Oracle documentation. [270905]
- Existing snapshot plexes created by the `vxassist` command are not supported. A combination of snapshot plexes created by `vxassist` and `vxsnap` is also not supported.



## Oracle Disk Manager limitations

- Because Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O, do not turn off the Quick I/O mount option. The default option is the correct option to use.
- Using Oracle Disk Manager with Cached Quick I/O enabled is not supported and could cause your system to panic [34281]. To avoid a system panic, ensure the following:
  - If you are using Oracle Disk Manager, do not enable Cached Quick I/O on your file system.
  - If you are converting from Quick I/O to Oracle Disk Manager, make sure you disable Cached Quick I/O.

## Clone database ORACLE\_SID character limit

When cloning an Oracle instance using the `dbed_clonedb` or `dbed_vmclonedb` command, the clone database's `ORACLE_SID` can be only eight characters or less. You will receive an error (ERROR V-81-5713) if the `ORACLE_SID` is more the eight characters.

## Renaming columns in `login.sql`

Renaming columns in `login.sql` can sometimes cause scripts to fail or produce incorrect results. To prevent this, make the following changes in the user environment to generally avoid loading `login.sql`:

- 1 Move `login.sql` to another directory, for example, to `~oracle/login.sql`.
- 2 Make sure this new directory is included in `SQLPATH`, for example:  

```
export SQLPATH=~oracle/sql:$SQLPATH
```
- 3 Do not make `SQLPATH` read-only, so that Storage Foundation for Oracle scripts can unset it at runtime.
- 4 Avoid starting Storage Foundation for Oracle scripts from the directory where `login.sql` resides, unless you are sure that `login.sql` does not contain any settings or commands that change the default output for queries against the data dictionary or increase the startup time for SQL\*Plus.

Also avoid using any settings or commands in the `glogin.sql` file that change the default output for queries against the data dictionary, or that may increase the startup time for SQL\*Plus.

## 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 Oracle 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:

- `/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]

## 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]

## 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 Storage Foundation fixed issues

The following table contains information about fixed issues in this release of Veritas Storage Foundation.

Incident	Description
	The Veritas product installer exited when a required patch was not detected, even if a superseding patch was available.

### Veritas Volume Manager fixed issues

The following table contains information about fixed issues in this release of VxVM.

Incident	Description
528677	Volume relayout is now supported for site-confined volumes and for site-consistent volumes.
540351	Reattaching a site when the disks were in the serial-split brain condition gave an error message.
540523	Under some circumstances, DMP nodes could be incorrectly enabled.
543803	The VxVM configuration daemon, <code>vxconfigd</code> , could dump core under rare conditions if the <code>vxctl enable</code> command was run on a system with an HDS array.
563524	Split, join and move operations failed on a source disk group that had any site-confined volumes.
584200	The <code>vxmake</code> command could not be used to recreate site records. This is now supported if the <code>-d</code> option is used to read from a description file.
601274	In a CVM cluster, DMP did not fail over to a secondary path when the primary paths were disconnected.

Incident	Description
605743	If a disk group were split from a source disk group, volumes in the split-off disk group did not retain their volume tags.
609199	When the <code>vxddmpadm disable</code> command was applied to a primary path on one node in a CVM cluster, the other nodes did not fail over to the secondary path.
611333	DMP could not obtain the correct serial number for a device if its LUN serial number contained a comma. This problem was seen on EMC Symmetrix arrays with more than 8096 LUNs.
614061, 614787	Adding cache volumes (used by space-optimized instant snapshots) to volume sets could cause data corruption and system panics.
617331	I/O was not restored on a path that was re-enabled after a failback or a non-disruptive upgrade (NDU) operation.
618317	A system crash could occur while bringing up cluster if I/O were performed on a unopened path.
621832	Immediately after installation, the <code>vxesd</code> daemon had the DVD mount point as its current working directory, which prevented the DVD from being unmounted.
625877	The error <code>"/etc/vx/vxvm-startup: line 241: /usr/sbin/vxddladm: No such file or directory"</code> was seen at boot time.
643089	Relayout from <code>mirror-stripe</code> to <code>concat-mirror</code> did not work for site-consistent volumes.
645749	Growing a volume by a specified amount did not work for a site-consistent volume with more than 2 disks per site.
774434	The <code>-e</code> option to the <code>vxdisk list</code> command had been omitted.
793159	Automatic reattachment of a remote site did not work correctly.
794228	Using option 4 of the <code>vxdiskadm</code> command to replace a failed or removed disk gave error V-5-1-582.
801445	The DMP feature to detect and respond to intermittently failing paths was turned off by default in the 5.0 release, and the values of the <code>dmp_health_time</code> and <code>dmp_path_age</code> tunables were both set to 0. This feature is now enabled by default in 5.0 MP1. The default values of <code>dmp_health_time</code> and <code>dmp_path_age</code> are 60 and 300 seconds respectively.

## Veritas Enterprise Administrator fixed issues

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

Incident	Description
578688	The maximum size of the Alert and Task logs has been documented as 2MB.
596284	An Action pull-down menu item did not exist for the Layout View, the Disk View or the Volume View.
599060	Controller states were reported as “Not Healthy” when they are actually healthy, and “Healthy” when they were actually not healthy.
614761	The volume set creation wizard showed cache volumes in the “Available Volumes” list.
616661	When connecting to the central host, an “OutOfBoundException” error could occur.
618146	A Java exception error occurred in the Statistics View.

## Veritas Web GUI fixed issues

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

Incident	Description
564455	Removing a volume from a volume set returned a Java exception.
565072	Creating a file system on a disabled volume returned both success and failure messages.
566619	The Scan Disks By Controller View did not list the available controllers.
574410	Attempting to create a volume without an existing disk group produced a misleading error.
575262	Disabling a path to a SENA storage array produced an erroneous message.
576794	Ghost entries for disconnected disks in the All Disks View could not be removed by using the GUI.

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 this release of VxFS.

Incident	Description
616323	For Web GUI online help, the following issues have been fixed: For the <b>Remount Storage Checkpoint</b> operation, the <b>More info</b> link on the second wizard page does not function properly for cluster file systems. For the <b>Unmount Storage Checkpoint</b> operation, the <b>More info</b> 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 ( <i>atime</i> ) or modification time ( <i>mtime</i> ) of named data streams by calling <code>vxfs_nattr_utimes()</code> API on 32-bit kernel.
770953	<code>fsck</code> used to create the <code>lost+found</code> directory with the <code>rwrxrwxrwx</code> permissions if it decided to create one. Now, it creates the directory with the <code>rwxr-xr-x</code> permissions, which is consistent with the behavior of <code>mkfs</code> .
770964	<code>fsck</code> has been enhanced to replay file systems created with earlier log versions on volume sets.
771086	Fixed an <code>fsck</code> problem in which users could end up creating multiple <code>lost+found</code> directories when running the <code>fsck -o full</code> command and answering <code>fsck</code> questions interactively. Now, <code>fsck</code> creates only one. It also checks for multiple <code>lost+found</code> entries and removes duplicate directory entries.
771996	Enhanced VxFS to use less CPU when doing administrative tasks on the devices of multi-volume file systems.
772013	Enhanced the <code>fsck</code> command to enforce the <code>lost+found</code> file name in the root directory of the file system to be a directory file type.
777012	If the system crashed or there was a metadata I/O error, after the <code>fsadm</code> command reorganized the <code>lost+found</code> directory, running the <code>fsck -o full</code> command may not have been able to clean the file system with regard to names that needed to be added to the <code>lost+found</code> directory. The problem happened on single-volume and multi-volume file systems.
785649	A situation where <code>vxfsconvert</code> of a dusty file system loops forever in user-level code when an inode with pending truncation operation is encountered has been fixed.
793022	<code>vxfs_nattr_open()</code> API interface has been fixed to shrink files, as appropriate, when invoked with <code>O_TRUNC</code> flag.
793030	<code>vxfsutil.h</code> uses <code>struct timeval</code> in one of the function declarations, but does not include <code>time.h</code> . This causes user applications to report warnings during compilation. This issue has been fixed by including <code>time.h</code> in <code>vxfsutil.h</code> .
795073	The increased CPU utilization when writing to a file system that is almost full due to more background processing threads than are actually required being enqueued has been fixed.



## 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 <b>Next</b> 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 Oracle fixed issues

The following table contains information about fixed issues in this release of Veritas Storage Foundation for Oracle.

Incident	Description
567342	An unmounted checkpoint clone database no longer reappears in the Java GUI tree after rescanning.
582069	SFDB commands 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 <b>Help</b> 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	The Web GUI statistic scheduler no longer skips the first statistic collection.
608697	You can now refresh the View Statistics wizard in the Firefox browser.
609682	Specifying with <code>dbed_vmsnap</code> a snapplan that does not exist no longer produces the error <code>SFORA dbed_vmsnap ERROR V-81-6518 Could not find snapplan 'snap_plan' in repository.</code>

Incident	Description
786989	The <code>qio_getdbfiles_ora</code> script now detects when an Oracle instance is in Standby mode.
853363	<p>The I/O performance of EMC Symmetrix arrays has been improved in this release. To enable these changes, after upgrading to this release, set the discovery mode of the VAIL provider to discover only those Symmetrix devices that are visible to the host:</p> <ol style="list-style-type: none"><li>Determine the agent name under which the Symmetrix provider is configured: # <code>/opt/VRTSvail/bin/vail_symm_discovery_cfg.sh -1</code> The agent name will be "VAILAgent" for installations of Veritas Storage Foundation, Veritas Storage Foundation for Databases, or Veritas Storage Foundation for RAC. It will be "StorageAgent" if VxFAS is configured.</li><li>Set the discovery mode to discover host-visible devices only: # <code>/opt/VRTSvail/bin/vail_symm_discovery_cfg.sh \</code> <code>-a agent_name -s 0</code> where <code>agent_name</code> is the agent name out put from the <code>-1</code> option in the previous step.</li></ol>

# Known issues

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

## Veritas Storage Foundation known issues

### Software disc cannot be ejected during installation

During installation, if any of the products were configured and started, the software disc cannot be ejected. This may prevent installation from continuing in following circumstances:

- If the language pack disc needs be loaded so that the associated packages can be installed.
- A product was installed that did not require a system reboot to complete the installation.

This problem is not an issue if a product was installed or upgraded that required a system reboot to complete the installation.

#### To avoid this problem at install time

- 1 Specify the `-installonly` option to the `installer` script in addition to any other options.
- 2 Eject the software disc.
- 3 Run the `installer` script with the `-configure` option specified.

#### If a software disc cannot be ejected

- 1 Stop the event source daemon:

```
# /usr/sbin/vxddladm stop eventsource
```
- 2 Kill the `vxcached`, `vxrelocd` and `vxnotify` processes by using the `kill -9` command with their process IDs as reported by the `ps` command.
- 3 Eject the software disc.
- 4 Restart the VxVM daemon processes:

```
# /usr/sbin/vxddladm start eventsource  
# /etc/vx/bin/vxcached  
# /etc/vx/bin/vxrelocd
```

[622442]

## Upgrading the software

To install a Veritas Storage Foundation product, configure Veritas Volume Manager as described in Chapter 2, "Installing the Veritas Software," of the *Veritas Storage Foundation Installation Guide*. Then, follow the instructions in Chapter 3, "Upgrading the Veritas Software," of the *Veritas Storage Foundation Installation Guide*.

Because Veritas Volume Manager 3.5 and Veritas File System 3.5 are bundled in the September 2004 HP-UX 11i version 2.0 release, a fresh installation is *not* possible. You must perform an upgrade to move to the 5.0 versions of the Veritas products.

## Uninstalling the VRTSmapro package

Uninstalling the VRTSmapro (mapping provider) package does not remove the corresponding entry from the VEA registry. The recommended workaround is to run the following command before uninstalling the VRTSmapro package:

```
# /opt/VRTSmapro/bin/vxmapping_prov.config -r  
[617740]
```

## DBMS security issue

The Symantec Shared DBMS feature creates the following configuration files:

- /etc/vxdbms/VERITAS\_DBMS3\_hostname/conf/databases.conf
- /etc/vxdbms/VERITAS\_DBMS3\_hostname/conf/databases1.conf
- /etc/vxdbms/VERITAS\_DBMS3\_hostname/conf/registration.dat

These configuration files are created or modified by `vxdbms_start_db.pl`, `vxdbms_start-server.pl`, and `vxdbms_register.pl` respectively.

The files are writable by everyone if the file mode creation mask is insufficiently restrictive. Symantec strongly recommends that you restrict the default file mode creation permissions (using the `umask` command) for root and administrator accounts to avoid a potential security issue. Specifically, change the group/world write and execute permissions in the default `umask`. At its least restrictive, the default `umask` for root should be 022. Symantec recommends setting to 077.

## Host name may need to be entered manually on clustered host

When installing SF Management Server on a clustered host on which the Domain Name Service (DNS) is not configured, the you will be prompted to enter the fully-qualified host name manually to proceed with the installation.

When installing on a clustered host on which the DNS is running slowly, the you may occasionally be prompted to enter the fully-qualified host name manually to proceed with the installation.

## Unconfigure VEA Action Agent after a MANAGED host installation

When installing in MANAGED mode you must execute the following script at any time after the Veritas packages have been installed:

```
/opt/VRTSaa/config/remove_vxaa.sh
```

This script unconfigures the VEA Action Agent and prevents it from starting on a MANAGED host installation. [616057]

## Veritas Storage Foundation for Oracle known issues

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

<http://entsupport.symantec.com/docs/283708>

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

### Cannot unmount single-host clone in HA environment after failover

In an HA environment, after successfully taking a snapshot and cloning the database on the same host where primary is running, if a node failover happens then `dbed_vmclonedb -o umount` does not work. [818522]

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

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

The file fragmentation check in `qio_convertdbfiles` may report errors when run on multi-volume file systems. These errors are harmless and may be safely ignored. This issue also causes the `dbed_checkconfig` command to fail with an error if run on a database which uses one or more multi-volume file systems.

The method used to determine fragmentation in `qio_convertdbfiles` has been deprecated. The preferred way to check and resolve file or file system fragmentation is through the use of the `fsadm` tool. Refer to the *Veritas File System Administrator's Guide* for more information on using `fsadm` to display and resolve file system fragmentation. [819430]

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

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

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

### To uninstall the Veritas Storage Foundation for Oracle packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 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 `swremove` command.

```
# swremove VRTSorgui VRTSdbed VRTSdbcom 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]

### **dbed\_vmclonedb -p failed to create clonedb with modified pfile**

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

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

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

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

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

---

**Note:** The primary instance's pfile is saved at the time of snapshot creation. If you attempt to clone the database using that snapshot you will be using the saved pfile, not the current pfile. Therefore you must create a new snapshot in order to ensure that the clone will use an updated pfile. [853792]

---

### **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 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/283708>

The following issues were reported for this release of VxVM.

### Device issues

#### Importing EMC BCV devices

The following procedure can be used to import a cloned disk (BCV device) from an EMC Symmetrix array.

#### To import an EMC BCV device

- 1 Verify that that the cloned disk, EMC0\_27, is in the error udid\_mismatch state:

```
# vxdisk -o alldgs list
DEVICE          TYPE          DISK          GROUP  STATUS
EMC0_1          auto:cdsdisk EMC0_1        mydg   online
EMC0_27         auto          -             -      error udid_mismatch
```

In this example, the device EMC0\_27 is a clone of EMC0\_1.

- 2 Split the BCV device that corresponds to EMC0\_27 from the disk group mydg:

```
# /usr/symcli/bin/symmir -g mydg split DEV001
```

In this example, the corresponding BCV device to EMC0\_27 is DEV001.

- 3 Update the information that VxVM holds about the device:

```
# vxdisk scandisks
```

- 4 Check that the cloned disk is now in the online `udid_mismatch` state:

```
# vxdisk -o alldgs list
DEVICE          TYPE          DISK          GROUP  STATUS
EMC0_1          auto:cdsdisk EMC0_1        mydg   online
EMC0_27         auto:cdsdisk -            -      online udid_mismatch
```

- 5 Import the cloned disk into the new disk group `newdg`, and update the disk's UDID:

```
# vxdbg -n newdg -o useclonedev=on -o updateid import mydg
```

- 6 Check that the state of the cloned disk is now shown as `online`

`clone_disk`:

```
# vxdisk -o alldgs list
DEVICE          TYPE          DISK          GROUP  STATUS
EMC0_1          auto:cdsdisk EMC0_1        mydg   online
EMC0_27         auto:cdsdisk EMC0_1        newdg  online clone_disk
```

## DMP issues

### Fabric Monitoring

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

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

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

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

```
# vxddladm gettune monitor_fabric
```

The default setting of `monitor_fabric` is `on`. [784343]

### Handling intermittently failing paths

The `dmp_health_time` and `dmp_path_age` tunables control how DMP handles intermittently failing paths. The default values in VxVM 5.0 MP1 of `dmp_health_time` and `dmp_path_age` are 60 and 300 seconds respectively. The value of `dmp_health_time` represents the minimum time in seconds for which a path must stay healthy. If a path changes state between enabled and disabled on a shorter time scale than this, DMP marks the path as intermittently failing and disables I/O on the path. I/O is not re-enabled on an intermittently failing path until `dmp_path_age` seconds have elapsed without further outage.

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

## Cluster issues

### Handling intermittently failing paths in a Campus Cluster

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

### Automatic site reattachment

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

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

If `vxrelocd` is not running, `vxstied` reattaches a site only when all the disks at that site become accessible. After reattachment succeeds, `vxstied` 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:** `vxstied` 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 `vxstied` 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 `vxstied` in the `/sbin/init.d/vxvm-recover` startup script, and reboot the system.

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

```
# kill -9 PID
```

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

To prevent the automatic site reattachment feature from being restarted, comment out the line that starts `vxstied` in the `/sbin/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 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]

### Error when creating a volume set

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

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

The workaround is to refresh the page. [615960]

## Veritas File System known issues

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

<http://entsupport.symantec.com/docs/283708>

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

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

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

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

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

## Veritas Volume Replicator known issues

Known issues in the 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/283745>

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

### Synchronizing volumes and RVG with large volumes

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

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

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

## 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]



## Documentation errata

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

### Veritas Storage Foundation Release Notes errata

The following errata apply to the *Veritas Storage Foundation Release Notes*:

#### **Applications using statvfsdev(3C) interface**

The description for this Veritas File System known issue should read as follows:

The applications that use the `statvfsdev(3C)` interface on a device name need to relink with the new LIBC to be able to understand disk layout Version 6 or 7. Disk Layout Version 7 is available with this release. If the applications are not relinked, they cannot recognize a file system with disk layout Version 6 or 7. However, they would continue to recognize file systems with disk layout Version 4 and 5.

### 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)`, `vxdumpadm(1M)`, `vxdumping(1M)`, `vxpool(1M)`, `vxresize(1M)`, `vxtemplate(1M)`, and `vxvoladm(1M)` manual pages are updated in this Maintenance Pack to include corrections for several errors or omissions.

## Veritas Volume Manager Administrator's Guide errata

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

### Specifying storage for version 20 DCO plexes

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

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

This should read:

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

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

### DMP configuration values

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

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

## Veritas Storage Foundation for Oracle Administrator's Guide errata

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

### Setting up Oracle 9i RAC objects with `srvctl`

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

#### To set up Oracle 9i RAC objects

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

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

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

```
# ls -l /db/srvm.ora
```
- 3 If the configuration file does not exist, create and initialize the file:  

```
# touch /db/srvm.ora  
# srvconfig -init
```
- 4 If the configuration file exists, note the size of the file shown by the output of `ls -l`.  

```
-rw-r--r-- 1 oracle dba 10569216 Jan 20 14:29 /db/srvm.ora
```
- 5 If the configuration file size is greater than zero (as shown in the example above), the file is initialized. If the file size is zero, initialize it:  

```
# srvconfig -init
```
- 6 Start the Oracle RAC Manageability daemon on each system:  

```
$ gsdctl start
```
- 7 Confirm the GSD daemon status:  

```
$ gsdctl stat  
GSD is running on the local node
```
- 8 Add the database to the `srvctl` configuration:  

```
$ srvctl add database -d KPRDADV1 -o /apps/oracle/product/920rac  
$ srvctl config database  
KPRDADV1
```
- 9 Add each instance to the configuration. For example, in a two-instance configuration, add the first instance:  

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

Then add the second instance:  

```
$ srvctl add instance -d KPRDADV1 -i KPADV1R2 -n node2  
$ srvctl config database -d KPRDADV1  
node1 KPADV1R1 /apps/oracle/product/920rac  
node2 KPADV1R2 /apps/oracle/product/920rac
```
- 10 Check the status of the instances to confirm they are running:  

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

## Reconfigure virtual IP address for repository configuration

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

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

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

In a cluster environment, do the following:

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

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

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

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

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

- (Page 285) 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 175 and 337) 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
```

