

Veritas Storage Foundation™ Release Notes

HP-UX

5.0 Maintenance Pack 2



Veritas Storage Foundation Release Notes

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.0 MP2

Document version: 5.0MP2.0

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the Veritas Storage Foundation 5.0 Release Notes.

The Veritas Storage Foundation 5.0 Release Notes can be viewed at the following URL:
<http://entsupport.symantec.com/docs/283708>

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

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 as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

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

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

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

Contacting Technical Support

Customers with a current maintenance agreement may access Technical Support information at the following URL:

<http://www.symantec.com/business/support/index.jsp>

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:
<http://www.symantec.com/techsupp/>

Customer service

Customer service information is available at the following URL:
<http://www.symantec.com/techsupp/>

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization
 - Product registration updates, such as address or name changes
 - General product information (features, language availability, local dealers)
 - Latest information about product updates and upgrades
 - Information about upgrade assurance and maintenance contracts
 - Information about the Symantec Buying Programs
 - Advice about Symantec's technical support options
 - Nontechnical presales questions
 - Issues that are related to CD-ROMs or manuals

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance agreement administration team for your region as follows:

Asia-Pacific and Japan

contractsadmin@symantec.com

Europe, Middle-East, and Africa

semea@symantec.com

North America and Latin America

supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions

These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.

Managed Security Services

These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.

Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:
<http://www.symantec.com>
Select your country or language from the site index.

Contents

Introduction	11
System requirements	13
HP-UX operating system requirements	13
Hardware requirements	13
Supported Oracle and HP-UX combinations	13
Required HP-UX patches	14
Mandatory patch required for Oracle Bug 4130116	16
Mandatory patches for VEA	16
Hard disk array support	17
Component product release notes	17
About the common product installer	18
Installing Veritas Storage Foundation products	18
Installing Storage Foundation using the common product installer ..	18
Installing Storage Foundation and HA Solutions using the common product installer	20
Installing Veritas Volume Manager	22
Configuring Storage Foundation and HA products	23
Configuring Storage Foundation	23
Configuring Veritas Storage Foundation and HA	25
Configuring Veritas Storage Foundation on a cluster	26
Configuring Storage Foundation for Oracle on cluster	29
Database configuration requirements	33
Creating and configuring the repository database for Oracle	34
Setting administrative permissions for databases	36
Setting up Veritas extension for Oracle Disk Manager	36
Configuring Veritas Volume Manager	37
Upgrading Storage Foundation	38
Installation and upgrade requirements	39
Upgrading Veritas Storage Foundation or Veritas Storage Foundation Oracle from 5.0 or 5.0 MP1 to 5.0 MP2	39
Upgrading Veritas Storage Foundation HA from 5.0 or 5.0 MP1 to Storage Foundation HA 5.0 MP2	40
Upgrading Veritas Storage Foundation Oracle HA from 5.0 or 5.0 MP1 to 5.0 MP2	40
Preparing to upgrade the Veritas software	42
Upgrading Veritas Storage Foundation HA from 4.1MP2 to 5.0 MP2 ..	44
Upgrading Veritas Storage Foundation Oracle HA	

from 4.1 MP2 to 5.0 MP2	45
Configuring your system after the installation	46
Verifying the Storage Foundation installation	47
Verifying that the products were installed	47
Installation log files	47
Using the installation log file	48
Using the response log file	48
Using the summary file	48
Verifying Volume Manager Installation	48
Checking Volume File System installation	49
Command installation verification	49
Uninstalling Veritas Storage Foundation products	49
Summary of Veritas Storage Foundation uninstallation tasks	50
Dropping the repository database for Oracle	50
Shutting down cluster operations	51
Removing VxFS file systems and Storage Checkpoints	51
Removing the root disk from VxVM	52
Moving volumes to disk partitions	52
Shutting down Veritas Volume Manager	57
Uninstalling Veritas Storage Foundation packages	57
Uninstalling Veritas Volume Manager	58
Removing the Veritas Enterprise Administrator client	58
Uninstalling the SF 5.0 MP2 patches	59
Software limitations	61
Storage Foundation software limitations	61
Veritas Volume Manager software limitations	61
Veritas File System software limitations	62
Veritas Storage Foundation for Oracle software limitations	62
No longer supported	67
Fixed issues	67
Veritas Storage Foundation fixed issues	67
Veritas Volume Manager fixed issues	68
Veritas File System fixed issues	80
Veritas Storage Foundation for Oracle fixed issues	81
Known issues	84
Veritas Storage Foundation 5.0 MP2 known issues	84
Veritas Storage Foundation 5.0 MP1 known issues	91
Documentation errata	101
Veritas Storage Foundation Release Notes errata	101
Web GUI help errata	101
Manual pages errata	101
Veritas Volume Manager Administrator's Guide errata	102
Veritas Storage Foundation for Oracle Administrator's Guide errata	102

Veritas Storage Foundation for Oracle GUI Guide errata105

Veritas Storage Foundation Release Notes

Introduction

This document provides release information about the products in the Veritas Storage Foundation (SF) 5.0 Maintenance Pack 2 (MP2) for HP-UX 11iv2 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
- Veritas Volume Replicator (VVR)

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*.

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/319349>

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:

- This release of Veritas Storage Foundation and Veritas Storage Foundation for Oracle can only be installed on a system running the September 2004 HP-UX 11i version 2.0 release (or later) on PA-RISC or Itanium platforms. However, before you upgrade to SF 5.0 MP2, Symantec recommends you install the latest Fusion release (June, 2008 or later) of the HP-UX 11i version 2 operating system.

To verify the operating system version

- ◆ Enter the swlist command:

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

Before you install any Veritas software, JFS must be installed on your system.

To verify that JFS is installed

- ◆ Enter the swlist command:

```
# 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

Table 1 identifies the supported Oracle and HP-UX combinations if you plan to use Veritas Storage Foundation with an Oracle database:

Table 1 Supported Oracle and HP-UX combinations

Oracle release	September 2004 HP-UX 11i version 2.0 (or later)
9.2	Yes
10.1	Yes

Table 1 Supported Oracle and HP-UX combinations

Oracle release	September 2004 HP-UX 11i version 2.0 (or later)
10.2	Yes
11g R1	Yes

Required HP-UX patches

The SF 5.0 MP2 releases of Veritas Storage Foundation and Veritas Storage Foundation for Oracle require the following HP-UX patches.

Warning: Install all the latest required HP-UX patches before you install Veritas Storage Foundation, Veritas Storage Foundation High Availability Solutions, Storage Foundation Cluster File System, or Veritas Volume Replicator. You can use the `swlist` command to determine whether the correct update and patches are installed. If the correct patches are not found, the installation terminates. Make sure that Enable VXFS bundle has revision B.11.23.04 or later after installing the latest patches.

HP may release patches that supersede the ones in this list. To make sure that you have the latest HP-UX patches, go to the Symantec support website to view the relevant TechNote.

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

Also, you can get the patches from Patch Database of Hewlett-Packard offered under the Maintenance and Support section of the HP Services & Support - IT Resource Center. The Patch Database of HP provides fast, accurate searches for the latest recommended and superseded patches available for Veritas File System or Veritas Volume Manager.

Most of the above patches are available in the Feature11i bundle. The Feature11i bundle is available from the HP software download site at the following URL:

<http://h20293.www2.hp.com/>

In addition to the above patches the `EnableVXFS` bundle needs to be installed before installing the VxFS 5.0 MP2 file system. This bundle is an HP bundle and contains enhancements to various commands to understand the new disk layout Version 6 and later.

Table 2 Required HP-UX patches

HP-UX Patch ID	Description
PHCO_32385	Enables <code>fscat(1M)</code> .

Table 2 Required HP-UX patches

HP-UX Patch ID	Description
PHCO_32387	Enables <code>gettext(1M)</code> .
PHCO_32388	Enables <code>ext(1M)</code> .
PHCO_32389	Enables <code>vxdump(1M)</code> .
PHCO_32390	Enables <code>vxrestore(1M)</code> .
PHCO_32391	Enables <code>vxfssstat(1M)</code> .
PHCO_32392	Enables <code>vxtunefs(1M)</code> .
PHCO_32393	Enables <code>vxupgrade(1M)</code> .
PHCO_32488	Enables LIBC for VxFS 4.1 and later file systems.
PHCO_32523	Enhancement to <code>quota(1)</code> for supporting large uids.
PHCO_32524	Enhancement to <code>edquota</code> for supporting large uids.
PHCO_32551	Enhancement to <code>quotaon/quotaoff</code> for supporting large uids.
PHCO_32552	Enhancement to <code>repquota</code> for supporting large uids.
PHCO_32596	Enables <code>df(1M)</code> .
PHCO_32608	Enables <code>bdf(1M)</code> .
PHCO_32609	Enables <code>fstyp(1M)</code> .
PHCO_32610	Enables <code>mount(1M)</code> .
PHCO_32611	Fix <code>fs_wrapper</code> to accept "vxfs" from subtype.
PHCO_33238	<code>swapon(1M)</code> cumulative patch.
PHCO_34036	LVM commands patch.
PHCO_34208	SAM cumulative patch.
PHCO_34191	Cumulative <code>libc</code> patch.
PHSS_32228	LIBCL patch. If the patch is not installed, the VAILAgent process could terminate and LUN-to-disk mapping would fail.
PHSS_32231	FORTRAN I/O Library patch. If the patch is not installed, the VAILAgent process could terminate and LUN-to-disk mapping would fail.
PHSS_32674	Obam patch (backend for the SAM patch).
PHKL_31500	Sept04 Base patch.

Table 2 Required HP-UX patches

HP-UX Patch ID	Description
PHKL_32272	Changes to fix intermittent failures in getacl/setacl.
PHKL_32430	Changes to separate vxfs symbols from libdebug.a, so that symbols of VxFS 4.1 and later are easily available in q4/p4.
PHKL_32431	Changes to disallow mounting of a file system on a vnode having VNOMOUNT set. Enhancements for supporting quotas on large uids.
PHKL_33312	LVM Cumulative patch.
PHKL_34010	Cumulative VM patch.

[Table 3](#) lists the patches in the `EnableVXFS` bundle.

Table 3 EnableVXFS bundle

HP-UX Patch ID	Description
FsLibEnh	Enhancements to LIBC libraries to understand VxFS disk layout Version 6 and later.
DiskQuota-Enh	Enhancements to various quota-related commands to support large uids.
FsCmdsEnh	Enhancements to the <code>mount</code> command to support VxFS 4.1.

For VEA functionality to work reliably, you must have HP-UX patches PHSS_33903 and PHSS_38526 (Itanium systems only) installed prior to installing Veritas Storage Foundation Cluster File System 5.0 Maintenance Pack 2.

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.

Mandatory patches for VEA

For VEA functionality to work reliably, you must have HP-UX patches PHSS_33903 and PHSS_38526 (Itanium systems only) installed prior to installing Veritas Storage Foundation 5.0 Maintenance Pack 2.

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 the 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 SF 5.0 MP2 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 packages do not install any release notes, 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.

About the common product installer

The product installer is the recommended method to license and install the product. The installer also lets you configure the product, verify preinstallation requirements, and view the product's description.

If you obtained a standalone Veritas product from an electronic download site, the single product download files do not contain the general product installer. Use the product installation script to install the product.

At most points during an installation, you can type b (back) to return to a previous section of the installation procedure. The back feature of the installation scripts is context-sensitive, so it returns to the beginning of a grouped section of questions. If an installation procedure hangs, use Control-c to stop and exit the program. After a short delay, the script exits.

Default responses are in parentheses. Press Return to accept the defaults.

Additional options are available for the common product installer.

Installing Veritas Storage Foundation products

This section contains procedures for installing Veritas Storage Foundation:

[“Installing Storage Foundation using the common product installer” on page 18](#)

[“Installing Storage Foundation and HA Solutions using the common product installer” on page 20](#)

[“Installing Veritas Volume Manager” on page 22](#)

For instructions on installing Veritas Volume Replicator, see *Veritas Volume Replicator Installation Guide*.

Installing Storage Foundation using the common product installer

The Veritas product installer is the recommended method to license and install Storage Foundation. The following sample procedure is based on installing Storage Foundation on a single system.

The same procedure can also be used to install Veritas Storage Foundation for Oracle.

For an initial installation on a new system, you can use the following procedure.

To install Storage Foundation

- 1 To install Storage Foundation on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.
- 2 Load and mount the software disc.

- 3 Move to the top-level directory on the disc.
- 4 From this directory, type the following command to install on the local system only. Also use this command to install on remote systems using the secure shell (ssh) utilities:

```
# ./installer
```

If you use the remote shell utilities to install on remote systems, additionally specify the -rsh option:

```
# ./installer -rsh
```

The sample installation assumes that ssh is used.
- 5 Enter **I** to install and press Return.
- 6 When the list of available products is displayed, select Veritas Storage Foundation, enter the corresponding number, and press Return.
You can also install Veritas Storage Foundation for Oracle using this procedure. Select the number corresponding to this product, if desired.
- 7 You are prompted to enter the system names (in the following example, "host1") on which the software is to be installed. Enter the system name (or names), and press Return. Stop the VEA server:

```
Enter the system names separated by spaces on which to install SF: host1
```
- 8 Enter the product license information. Each system requires a product license before installation. License keys for additional product features must also be added at this time.

```
Enter a SF license key for host1:XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X  
XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered on host1  
SF license registered on host1
```
- 9 You are prompted to enter additional license information, until all licenses for all systems have been entered. Then reply that you have no additional licenses to enter.

```
Do you want to enter another license key for host1? [y,n,q] (n)  
n
```
- 10 You can choose to install required depots or all depots. For example, optional depots include man pages.
Each option displays the disk space that is required for installation. Select the option you want to install, and press Return.
You see output similar to the following:

```
SF can be installed without optional depots to conserve disk space.  
1) Install required Veritas Storage Foundation depots - 1132 MB required  
2) Install all Veritas Storage Foundation depots - 1132 MB required
```

```
Select the depots to be installed on all systems? [1-2,q,?] (2)
2
```

- 11 The list includes the items in the selected option. Press Return to continue.
- 12 The installation complete automatically view the log file, if needed, to confirm the installation.
*Installation log files, summary file, and response file are saved at: /opt/VRTS/install/logs/installer-*****
- 13 Reboot the systems.

Installing Storage Foundation and HA Solutions using the common product installer

The Veritas product installer is the recommended method to license and install Storage Foundation. Veritas Storage Foundation for Oracle can also be installed using this procedure.

To install Storage Foundation and High Availability products

- 1 To install on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.
- 2 Load and mount the software disc.
- 3 Move to the top-level directory on the disc.
- 4 From this directory, type the following command to install on the systems, if you use the `ssh` utilities:
`# ./installer`
If you use the remote shell utilities to install on remote systems, additionally specify the `-rsh` option:
`# ./installer -rsh`
The sample installation assumes that `ssh` is used.
- 5 Enter I to install and press Return.
- 6 When the list of available products is displayed, select Veritas Storage Foundation (SF), enter the corresponding number, and press Return.
With a Veritas Storage Foundation HA license, the high availability cluster components are also installed for this menu selection. Veritas Storage Foundation for Oracle can also be installed using this procedure. Select the number corresponding to this product, if desired.
- 7 You are prompted to enter the system names (in the following example, "host1" and "host2") on which the software is to be installed. Enter the system name or names, and press Return.

Enter the system names separated by spaces on which to install SF: host1
host2

- 8 During the initial system check, the installer verifies that communication between systems has been set up.

If the installer hangs or asks for a login password, stop the installer and set up ssh or rsh. Then run the installer again.

- 9 Enter the product license information.

```
Enter a SF license key for
host1: [?] XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered
on host1
```

```
Do you want to enter another license key for host1?
[y,n,q,?] (n) n
```

```
Enter a SF license key for
host2: [?] XXXX-XXXX-XXXX-XXXX-XXXX-X
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered
on host2
```

```
Do you want to enter another license key for host2? [y,n,q,?]
(n) n
```

Enter n if you have no further license keys to add for a system. You are then prompted to enter the keys for the next system.

Each system requires a product license before installation. License keys for additional product features must also be added at this time.

- 10 You can choose to either install only required depots or all depots. Optional depots include man pages, for example.

Each option displays the disk space that is required for installation. Select which option you want to install, and press Return.

For example, you see output similar to the following:

```
SF can be installed without optional depots to conserve disk space.
```

```
1) Required Veritas Storage Foundation depots - 929 MB required
```

```
2) All Veritas Storage Foundation depots - 930 MB required
```

```
Select the depots to be installed on all systems? [1-2,q,?] (2) 2
```

The list of optional depots may differ depending on the license key that you entered.

- 11 The installation complete automatically

View the log file, if required, to confirm the installation.

Installation log files, summary file, and response file are saved at:

```
/opt/VRTS/install/logs/installer-****
```

- 12 Reboot the systems.

- 13 If you installed Veritas Storage Foundation for Oracle, create a new repository database.

See “[Creating and configuring the repository database for Oracle](#)” on page 34

Installing Veritas Volume Manager

This section describes how to install the Veritas Volume Manager software and license key.

To install Veritas Volume Manager

- 1 Log in as superuser.
- 2 Run the installer command to install Veritas Volume Manager. For example:

```
# cd /dvdrom
# ./installer
```

From the product installer, choose the **I** option for Install, and select Veritas Volume Manager.

Note: If you have obtained a Veritas product from an electronic download site, the single product download files do not contain the installer installation script, so you must use the product installation script to install the product. For example, if you download Veritas Volume Manager, use the `installvm` script instead of the installer script.

- 3 Enter one or more system names on which Veritas Volume Manager is to be installed. For example:
Enter the system names separated by spaces on which to install VxVM:
system01
- 4 After the system check completes successfully, press Return to continue.
- 5 Enter a VxVM license key. For example: Enter a VxVM license key for system01:[?] XXXX-XXXX-XXXX-XXXX-XXXX-X
- 6 If required, you can enter another license key by typing **y** at the following prompt:
Do you want to enter another license key for system01? [y,n,q,?] (n)
- 7 After VxVM licensing completes successfully, press Return to continue.
- 8 You are asked which depots must be installed.
 - a Required Veritas Volume Manager packages - 698MB required
 - b All Veritas Volume Manager packages - 838 MB required
 - c Storage Foundation Enterprise HA packages - 1192 MB required.
- 9 Select the packages to be installed.

- 10 After the list of packages is displayed, you are told that you must reboot prior to reconfiguration.
- 11 After the Veritas Volume Manager installation completes successfully, you see a message identifying the location of log files. You must save these files for future reference.
- 12 Reboot the system.

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

Configuring Storage Foundation and HA products

This section contains procedures for configuring Veritas Storage Foundation and High Availability products.

Configuring Storage Foundation

This section describes how to configure Storage Foundation with the common product installer. The same procedure can also be used to install Veritas Storage Foundation for Oracle.

To configure Storage Foundation

- 1 To configure Storage Foundation, enter the following command:

```
# cd /opt/VRTS/install
# ./installsf -configure [-rsh]
```
- 2 Enter the names of the systems on which you want to configure the software.
Enter the system names separated by spaces on which to configure
SF: host1
- 3 The procedure checks system licensing, and you can enter additional licenses, if required.
Checking system licensing
SF license registered on host1
Do you want to enter another license key for host1? [y,n,q] (n)
n
- 4 The enclosure-based naming scheme (rather than disk device naming) is a feature of Veritas Volume Manager. You can reference disks using a symbolic name that is more meaningful than the operating system's normal device access name.

See the *Veritas Volume Manager Administrator's Guide*.

```
Do you want to set up the enclosure-based naming scheme?
[y,n,q,?] (n) n
```

- 5 You can specify the default name of a disk group that is to be assumed by Veritas Volume Manager commands if a disk group is not otherwise specified. Enter n if you do not want to specify the name of the default disk group at this time.

You can set the name of the default disk group after installation by running the `vxdctl defaultdg diskgroup` command on a system.

See the `vxdctl(1M)` manual page.

See the *Veritas Volume Manager Administrator's Guide*.

If you specify the name of a default disk group, this step does not create the disk group. After installation, you can use the `vxdiskadm` command to create the disk group.

Do you want to set up a default disk group for each system?

[y,n,q,?] (y) y

- 6 If you responded y, then enter the information for the default disk group name.

Will you specify one disk group name for all eligible systems?

[y,n,q,?] (y) y

Specify a default disk group name for all systems. [?] dg001

- 7 Verify the system's fully-qualified hostname.

Is the fully qualified hostname of system

"host1" = "host1.domain_name"? [y,n,q] (y) y

- 8 This product can be configured as a Storage Foundation Manager managed host or a stand-alone host.

Several prerequisites are necessary to configure the system as a Storage Foundation Manager managed host.

For more information, see the *Veritas Storage Foundation Manager Installation Guide*.

Enable Storage Foundation Management Server Management?

[y,n,q] (y) n

- 9 The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

Do you want to start Veritas Storage Foundation processes now?

[y,n,q]

(y) y

- 10 The configuration completes automatically.

Check the log file, if required, to confirm the configuration.

Configuration log files, summary file, and response file are saved at:

/opt/VRTS/install/logs/installer-****

Configuring Veritas Storage Foundation and HA

After installation, you must configure the product. To do this, run the Veritas product installer or the appropriate installation script using the -configure option.

Use the following procedures to configure Storage Foundation and High Availability Solutions and clusters using the common product installer.

Required information for configuring Storage Foundation and HA Solutions

To configure Storage Foundation and High Availability Solutions, the following information is required:

- A unique Cluster name
- A unique Cluster ID number between 0-65535
- Two or more NIC cards per system used for heartbeat links
 - One or more heartbeat links are configured as private links. One heartbeat link may be configured as a low-priority link.

Symantec Security Services

You can configure Veritas Storage Foundation to use Symantec Security Services.

Running Storage Foundation in Secure Mode guarantees that all inter-system communication is encrypted and that users are verified with security credentials.

When you use Storage Foundation in Secure Mode, NIS, and system user names and passwords are used to verify identity. Storage Foundation user names and passwords are no longer used when a cluster is running in Secure Mode.

Before you configure a cluster to operate using Symantec Security Services, another system must already have Symantec Security Services installed and be operating as a Root Broker.

Simple Mail Transfer Protocol (SMTP) notification

The following information is required to configure Simple Mail Transfer Protocol (SMTP) notification:

- The domain-based hostname of the SMTP server
 - The email address of each SMTP recipient
 - A minimum severity level of messages to be sent to each recipient
- The following information is required to configure SNMP notification:
- System names of SNMP consoles to receive VCS trap messages

- SNMP trap daemon port numbers for each console
- A minimum severity level of messages to be sent to each console

Configuring Veritas Storage Foundation on a cluster

The following sample procedure is based on the configuration of a Veritas Storage Foundation Enterprise HA cluster with two nodes: "host1" and "host2." Default responses are enclosed by parentheses. Press Return to accept defaults. The procedure in this section is only relevant if you are installing an HA version of the Storage Foundation software.

To configure Storage Foundation on a cluster

- 1 To invoke the common installer, run the installer command on the disc:

```
# cd /mnt/cdrom
# ./installer
```
- 2 At the product status page, enter C for Configure an Installed Product, and press Return. The product installer is displayed.
- 3 You are prompted to enter the system names (in the following example, "host1" and "host2") on which the software is to be installed. Enter the system name (or names) and press Return.
Enter the system names separated by spaces on which to install SF: **host1 host2**

- 4 At the prompt, enter y or press Return to configure the Storage Foundation product.

```
Are you ready to configure SF? [y,n,q] (y) y
You have a VCS license and you are installing SF HA.
```

- 5 When prompted to configure VCS, enter y to configure VCS on these systems.
Do you want to configure VCS on these systems at this time?


```
[y,n,q] (y) y
```
- 6 Enter the unique cluster name and Cluster ID number.

```
Enter the unique cluster name: [?] vcs_cluster2
Enter the unique Cluster ID number between 0-65535: [b,?] 76
The installer discovers the NICs available on the first system and lists them:
Discovering NICs on host1 ...discovered lan0 lan1 lan2
```

- 7 Enter the private heartbeat NIC information for each host.
Enter the NIC for the first private heartbeat link on host1:


```
[b,?] lan1
```


Would you like to configure a second private heartbeat link?


```
[y,n,q,b,?] (y)
```


Enter the NIC for the second private heartbeat link on host1:


```
[b,?] lan2
```

Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n) **n**
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n)
Are you using the same NICs for private heartbeat links on
all systems? [y,n,q,b,?] (y)
When answering y, be sure that the same NICs are available on each system;
the installer does not verify this.
Notice that in this example, lan0 is not selected for use as a private
heartbeat NIC because it is already used as the public network interface.
The default responses are chosen.

- 8 A summary of the information you entered is given. When prompted,
confirm that the information is correct.
Is this information correct? [y,n,q] (y)
If the information is correct, press Return. If the information is not correct,
enter n. The installer prompts you to enter the information again.
- 9 When you are prompted if you want to configure the product to use Veritas
Security Services, enter y or n.
Before you configure a cluster to operate using Veritas Security Services,
another system must already have Veritas Security Services installed and
be operating as a Root Broker.
For more information on configuring a VxSS Root Broker, see the *Veritas
Cluster Server Installation Guide*.
Would you like to configure SF to use Veritas Security Services?
[y,n,q] (n) **n**
- 10 A message notifies you of the information required to add users. When
prompted, set the Administrator user name and/or password.
Do you want to set the username and/or password for the Admin
user (default username = 'admin', password='password')? [y,n,q] (n)
- 11 Enter n if you want to decline. If you enter y, you are prompted to change the
password. You are prompted to add another user to the cluster.
Do you want to add another user to the cluster? [y,n,q] (y)
- 12 Enter n if you want to decline, enter y if you want to add another user. You
are prompted to verify the user.
Is this information correct? [y,n,q] (y)
Enter y or n to verify if this information is correct.
- 13 You are prompted to configure the cluster management console. Enter y or n
to configure the cluster management console.
Do you want to configure the Cluster Management Console [y,n,q]
(n) **y**

- 14** Enter the NIC for the Cluster Manager (Web Console), then confirm whether the NIC is to be the public NIC used by all systems.

Enter the NIC for Cluster Manager (Web Console) to use on host1:

[b,?] (hme0)

Is hme0 to be the public NIC used by all systems [y,n,q,b,?] (y)

- 15** Enter the Virtual IP address for Cluster Manager.

- 16** You are prompted to verify the information.

Is this information correct? [y,n,q] (y)

Enter y or n to verify if this information is correct.

- 17** You are prompted to configure the cluster connector. Enter y or n to configure the cluster connector.

Do you want to configure the cluster connector [y,n,q] (n)

- 18** When prompted to configure SMTP notification, enter y or n to configure.

Do you want to configure SMTP notification? [y,n,q] (y)

- 19** When prompted to configure SNMP notification, enter y or n to configure.

Do you want to configure SNMP notification? [y,n,q] (y)

- 20** When prompted to set up the default disk group for each system, enter y to set up the disk group for each system.

Do you want to set up a default disk group for each system?

[y,n,q,?] (y)

- 21** When prompted to change any of the VVR ports on host1, enter y or n if you want to change any of the VVR ports on host1.

Do you want to change any of the VVR ports on host1? [y,n,q] (n)

- 22** When prompted to configure VVR statistics on host1, enter y or n to configure the VVR statistics.

Do you want to change the frequency of VVR statistics collection on host1? [y,n,q] (n)

Do you want to change the maximum number of days for retaining VVR statistics on host1? [y,n,q] (n)

- 23** When prompted to modify the VVR tunables on host1, enter y or n to modify VVR tunables host1.

Do you want to view or modify VVR tunables on host1? [y,n,q,?] (n)

- 24** When prompted to change any of the VVR ports on host2, enter y or n if you want to change any of the VVR ports on host2.

Do you want to change any of the VVR ports on host2? [y,n,q] (n)

- 25** When prompted to configure VVR statistics on host2, enter y or n to configure VVR statistics on host2.

Do you want to change the frequency of VVR statistics collection on host2 ? [y,n,q] (n)

Do you want to change the maximum number of days for retaining VVR statistics on host2 ? [y,n,q] (n)

- 26** When prompted to modify the VVR tunables on host2, enter y or n to modify the VVR tunables on host2.

Do you want to view or modify VVR tunables on host2 ? [y,n,q,?] (n)

- 27** You are prompted to enter the fully-qualified hostname of system host1. Enter y for the host1.domain_name.

Is the fully qualified hostname of system "host1" = "host1.domain_name"? [y,n,q] (y)

- 28** You are prompted to enter the fully-qualified hostname of system host1. Enter y for the host2.domain_name.

Is the fully qualified hostname of system "host2" = "host2.domain_name"? [y,n,q] (y)

- 29** You are prompted to enable Storage Foundation Management Server Management.

Enable Storage Foundation Management Server Management? [y,n,q] (y)

- 30** Enter y to enable Storage Foundation Management Server Management. You are prompted to start Storage Foundation.

Do you want to start Veritas Storage Foundation processes now? [y,n,q] (y)

...

Startup completed successfully on all systems.

Configuring Storage Foundation for Oracle on cluster

This section describes configuring Storage Foundation for Oracle using the Veritas Storage Foundation common product installer. However, a few additional configuration steps may also be required. Refer the following sections for more details.

To configure Storage Foundation product on a cluster

- 1 To invoke the common `installer`, run the `installer` command with the `configure` option:
`# ./installer -configure`
- 2 When the list of available products is displayed, select Veritas Storage Foundation for Oracle (SFORA), enter the corresponding number, and press Return.
`Select a product to configure:`

- 3 You are prompted to enter the system names (in the following example, "host1" and "host2") on which the software is to be installed. Enter the system name or names, and press Return.
Enter the system names separated by spaces on which to configure SF: **host1 host2**
- 4 During the initial system check, the installer checks that communication between systems has been set up.
The installer requires that ssh commands used between systems execute without prompting for passwords or confirmations. If the installer hangs or asks for a login password, stop the installer and run it again with ssh configured for password-free logins, or configure rsh and use the -rsh option.
- 5 The procedure checks system licensing, and you can enter additional licenses, if required.
Checking system licensing
SF license registered on host1
Do you want to enter another license key for host1? [y,n,q] (n)
n
- 6 When prompted to configure VCS, enter y to configure VCS on these systems.
Do you want to configure VCS on these systems at this time?
[y,n,q] (y) **y**
No configuration changes are made to the systems until all configuration questions are completed and confirmed. Press Return to continue.
All systems are configured to create one cluster.
- 7 Enter the unique cluster name and Cluster ID number.
Enter the unique cluster name: [?] **cluster2**
Enter the unique Cluster ID number between 0-65535: [b,?] **76**
- 8 The installer discovers the network interfaces (NICs) available on the first system and lists them:
Discovering NICs on host1 ... discovered lan0 lan1 lan2 lan3 lan4 lan5
- 9 Enter private heartbeat NIC information for each host.
Enter the NIC for the first private heartbeat link on host1: [b,?] **lan2**
Would you like to configure a second private heartbeat link? [y,n,q,b,?] (y) **y**
Enter the NIC for the second private heartbeat link on host1: [b,?] **lan3**
Would you like to configure a third private heartbeat link? [y,n,q,b,?] (n) **n**
Do you want to configure an additional low priority heartbeat link? [y,n,q,b,?] (n) **n**
Are you using the same NICs for private heartbeat links

on all systems? [y,n,q,b,?] (y) **y**

Warning: When answering y, make sure that the same NICs are available on each system; the installer may not verify this. The NICs must also be the same speed on both systems for the heartbeat links to function properly. Notice that in this example, lan0 selected for use as a private heartbeat NIC because it already used as the public network interface.

- 10 A summary of the information you entered is given. When prompted, confirm that the information is correct.
Is this information correct? [y,n,q]
If the information is correct, enter y. If the information is not correct, enter n. The installer prompts you to enter the information again.
- 11 When prompted to configure the product to use Veritas Security Services, enter n, unless a Root Broker has already been set up.

Warning: Before configuring a cluster to operate using Veritas Security Services, another system must already have Veritas Security Services installed and be operating as a Root Broker.

For more information on configuring a VxSS Root Broker, see the *Veritas Cluster Server Installation Guide*.

Would you like to configure SFORA to use Symantec Security Services? [y,n,q] (n) **y**

- 12 Security can be configured automatically, or it can be partially automatic. Select 1 to configure the security automatically.
Select the Security option you would like to perform [1-3,q,?] (1) **1**
- 13 At the prompt, enter the Root Broker name.
- 14 When you are prompted to configure SMTP notification, enter n or y to configure. To configure SNMP notification, enter the following information. You can then confirm that it is correct, or enter it again.

Do you want to configure SMTP notification? [y,n,q] (y) **y**

Active NIC devices discovered on host1: lan0

Enter the NIC for the SF Notifier to use on host1: [b,?] (lan0)
lan0

Is lan0 to be the public NIC used by all systems [y,n,q,b,?] (y)
y

Enter the domain-based hostname of the SMTP server (example: smtp.yourcompany.com): [b,?] **smtp.mycompany.com**

Enter the full email address of the SMTP recipient (example: user@yourcompany.com): [b,?] **user@mycompany.com**

```
Enter the minimum severity of events for which mail should be
sent
to user@163.com [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
```

- 15 When you are prompted to configure SNMP notification, enter n or y to configure. To configure SNMP notification enter the following information. You can then confirm that it is correct, or enter it again.

```
Do you want to configure SNMP notification? [y,n,q] (y)
Active NIC devices discovered on host1: lan0
Enter the NIC for the SF Notifier to use on host1: [b,?] (lan0)
lan0
Is lan0 to be the public NIC used by all systems [y,n,q,b,?] (y)
y
Enter the SNMP trap daemon port: [b,?] (162) 162
Enter the SNMP console system name: [b,?] host1
Enter the minimum severity of events for which SNMP traps should
be sent to host1 [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
Would you like to add another SNMP console? [y,n,q,b] (n) n
```

- 16 You are prompted to configure permissions to allow database administrators (DBAs) access to the tools to support the Veritas Storage Foundation for Oracle. The default settings only allow access to the root user. Respond y to change permission for a DBA or a group of DBAs to access the support tools for Veritas Storage Foundation for Oracle. When prompted, enter the login account or group name.

```
Do you want to add single user access on host1? [y,n,q,?] (y) y
Enter login account name for DBA user: oracle
Do you want to add group access on host1? [y,n,q,?] (y) y
Enter group name for DBA users: oinstall
Are you using the same DBA user/group for all systems? [y,n,q,?] (y) y
```

- 17 The enclosure-based naming scheme (rather than disk device naming) is a feature of Veritas Volume Manager. You can reference disks using a symbolic name that is more meaningful than the operating system's normal device access name.

See the *Veritas Volume Manager Administrator's Guide*

```
Do you want to set up the enclosure-based naming scheme?
[y,n,q,?] (n) n
```

- 18 You are given the option of specifying the default name of a disk group that is to be assumed by Veritas Volume Manager commands if a disk group is not otherwise specified. Enter n if you do not want to specify the name of the default disk group at this time. You can set the name of the default disk group after installation. If you specify the name of a default disk group, this

step does not create the disk group. After installation, you can use the `vxdiskadm` command to create the disk group.

Do you want to set up a default disk group for each system?
[y,n,q,?] (y) **y**

- 19 If you responded y, then enter the information for the default disk group name.

Will you specify one disk group name for all eligible systems?
[y,n,q,?] (y) **y**
Specify a default disk group name for all systems. [?] **dg001**

- 20 Validate the default disk group information, and press Return.
- 21 You may be prompted to verify the systems' fully-qualified hostname. Press Return to continue.
- 22 This product can be configured as a Storage Foundation Manager managed host or a stand-alone host. Several prerequisites are necessary to configure the system as a Storage Foundation Manager managed host.
For more information, see the *Veritas Storage Foundation Manager Installation Guide*.
- Enable Storage Foundation Management Server Management? [y,n,q]
(y) **n**
- 23 The Veritas Storage Foundation for Oracle software is verified and configured. Start the Veritas Storage Foundation for Oracle processes.
Do you want to start Veritas Storage Foundation for Oracle processes now? [y,n,q] (y) **y**
- 24 The configuration and startup complete automatically. View the log file, if required, to confirm the configuration.
Configuration log files, summary file, and response file are saved at:
`/opt/VRTS/install/logs/installer-****`
- 25 If you installed Veritas Storage Foundation for Oracle, create a new repository database.

Database configuration requirements

Most relational database management system (RDBMS) software requires operating system parameters to be set prior to operation. The Oracle database requires modifications to kernel settings before the databases run correctly. The most critical settings are usually in the Shared Memory and Semaphore settings on HP-UX. For precise settings, see your current database installation and configuration documentation.

Creating and configuring the repository database for Oracle

After you install Veritas Storage Foundation for Oracle, you must create and configure the repository database using the `sfua_db_config` script.

The script detects whether your system is running in a stand-alone or HA configuration and then automatically configures the repository database.

Before you run the script, review the following requirements for a stand-alone configuration:

- You must have a mount point mounted on a VxVM volume with a VxFS file system. The mount point stores the repository database.

Before you run the script, review the following requirements for an HA configuration:

- Create a separate, non-shared disk group on shared storage. Create a VxVM volume and a VxFS file system and mount the volume.
- It is recommended that you have a separate disk group for the repository volume so that any failovers are independent of other service groups.
- The mount point stores the repository database.
- Obtain an unique virtual IP address for the public NIC interface.
- Obtain the device names for the public NIC interface for all systems in the cluster. For example, use the following names.
 - hme0
 - lan0
- Obtain a subnet mask for the public NIC interface.
- Make sure VCS is not in read-write (`-rw`) mode. To make sure VCS is in read-only mode, use the following command:
`# haconf -dump -makero`

[Table 4](#) indicates the options available for the `sfua_db_config` script.

Table 4 sfua_db_config options

Option	Description
<code>-ssh</code>	Use this option in a high availability (HA) configuration. The option indicates that <code>ssh</code> and <code>scp</code> is used for communication between systems. Either <code>ssh</code> or <code>rsh</code> must be preconfigured so that you can execute the commands without being prompted for passwords or confirmations.
<code>-o dropdb</code>	Drops the repository database.

Table 4 sfua_db_config options

Option	Description
-o unconfig_cluster	Use this option in a high availability (HA) configuration. Unconfigures the repository database from the VCS cluster.
-o dbstatus	Verifies the status of the database and database server.
-o stopserver	Stops the database server.
-o startserver	Starts the database server.
-o serverstatus	Reports the database server status.
-o stopdb	Detaches the repository database from the database server.
-o startdb	Attaches the repository database to the database server.

To create and configure the repository database

- 1 Run the sfua_db_config script:

```
# /opt/VRTSdbcom/bin/sfua_db_config
```
- 2 Confirm that you are ready to configure the Veritas Storage Foundation for Oracle repository:

```
Are you ready to configure SFORA repository (y/n/q) [y]?
```
- 3 The mount point is displayed.

```
filesystem mount point for SFORA repository: /sfua_rep
```
- 4 The network interfaces (NICs) are discovered, and you are prompted to enter the NIC for the repository configuration on each host:

```
Enter the NIC for system host1 for HA Repository configuration: [lan0]
Enter the NIC for system host2 for HA Repository configuration: [lan0]
```
- 5 Enter the Virtual IP address for repository failover.

```
Enter the Virtual IP address for repository failover:
xxx.xxx.xxx.xxx
Enter the netmask for public NIC interface: [xxx.xxx.xxx.xxx]
Following information will be used for SFORA HA configuration:
Public IP address: xxx.xxx.xxx.xxx
Subnet mask: xxx.xxx.xxx.xxx
Public interface: host1 -> lan0,host2 -> lan0
```
- 6 The mount point information is displayed, and the script asks for confirmation. Then the repository information is added.
- 7 Verify that the repository was configured. If you are installing in a high availability configuration, enter the following command:

```
# /opt/VRTS/bin/hagrp -state
```

```
Group Attribute System Value
Sfua_Base State guan |ONLINE|
Sfua_Base State plover |OFFLINE|
Note: Sfua_Base group should be online on one node in the
cluster.
```

- 8 If you are installing in a stand-alone configuration, enter the following command to verify that the repository was configured:

```
# /opt/VRTSdbcom/bin/sfua_db_config -o dbstatus
Database 'dbed_db' is alive and well on server
'VERITAS_DBMS3_host'.
```

Setting administrative permissions for databases

To allow database administrators to administer a database using Veritas Storage Foundation, you are required to change some permission settings. During the installation process, you have the opportunity to configure the product. Answering “y” allows you to provide database administrators access to various functionality. If you did not make the permission changes during installation, you can do so at a later time.

The default settings at installation time for the `/opt/VRTSdbed` directory allow only the root login to access the directory.

To allow the user “oracle” access to the `/opt/VRTSdbed` directory

- ◆ Use the `chown` and `chmod` commands:

```
# chown oracle /opt/VRTSdbed
# chmod 500 /opt/VRTSdbed
```

To allow users in the group “dba” access to the `/opt/VRTSdbed` directory

- ◆ Use the `chgrp` and `chmod` commands:

```
# chgrp dba /opt/VRTSdbed
# chmod 550 /opt/VRTSdbed
```

Setting up Veritas extension for Oracle Disk Manager

To link the Veritas extension for Oracle Disk Manager library into Oracle home for Oracle 11g

- ◆ Do one of the following:

- For HP-UX PA, enter:

```
# mv ${ORACLE_HOME}/lib/libodm11.sl \
${ORACLE_HOME}/lib/libodm11.sl.orig
# ln -s /opt/VRTSodm/lib/libodm.sl \
${ORACLE_HOME}/lib/libodm11.sl
```

- For HP-UX IA, enter:

```
# mv ${ORACLE_HOME}/lib/libodm11.so \
${ORACLE_HOME}/lib/libodm11.so.orig
# ln -s /opt/VRTSodm/lib/libodm.sl \
${ORACLE_HOME}/lib/libodm11.so
```

To link the Veritas extension for Oracle Disk Manager library into Oracle home for Oracle 10g

- ◆ Do one of the following:

- For HP-UX PA, enter:

```
# mv ${ORACLE_HOME}/lib/libodm10.sl \
${ORACLE_HOME}/lib/libodm10.sl.orig
# ln -s /opt/VRTSodm/lib/libodm.sl \
${ORACLE_HOME}/lib/libodm10.sl
```

- For HP-UX IA, enter:

```
# mv ${ORACLE_HOME}/lib/libodm10.so \
${ORACLE_HOME}/lib/libodm10.so.orig
# ln -s /opt/VRTSodm/lib/libodm.sl \
${ORACLE_HOME}/lib/libodm10.so
```

To link the Veritas extension for Oracle Disk Manager library into Oracle home for Oracle9i

- ◆ Do one of the following:

- For HP-UX PA, enter:

```
# mv ${ORACLE_HOME}/lib/libodm9.sl \
${ORACLE_HOME}/lib/libodm9.sl.orig
# ln -s /opt/VRTSodm/lib/libodm.sl \
${ORACLE_HOME}/lib/libodm9.sl
```

- For HP-UX IA, enter:

```
# mv ${ORACLE_HOME}/lib/libodm9.so \
${ORACLE_HOME}/lib/libodm9.so.orig
# ln -s /opt/VRTSodm/lib/libodm.sl \
${ORACLE_HOME}/lib/libodm9.so
```

When Oracle Disk Manager is enabled, a message similar to the following is sent to the alert log:

Oracle instance running with ODM: Veritas #.# ODM Library, Version #.#.

When the system and instance are configured correctly, the Oracle Disk Manager feature is used, by default, for accessing any database storage.

Configuring Veritas Volume Manager

If you deferred configuring VxVM during installation, you can configure it by running the `installvm` script with the `-configure` option.

To configure VxVM using the `installvm` script:

- 1 Enter the following commands.

```
# cd /dvrom/volume_manager  
# ./installvm -configure
```

- 2 The script runs an initial system check, and informs you that you cannot configure already configured features such as enclosure-based naming and default disk groups.

- 3 Decide whether you want to set up the enclosure-based naming scheme:

```
Do you want to set up the enclosure-based naming scheme? [y, n, q]
```

- 4 You are then asked if you want to set up a default disk group for each system:

- 5 Do you want to set up a default disk group for each system? [y, n, q]

- 6 If you have a VVR license installed, the next phase concerns configuration of VVR:

```
Do you want to change any of the VVR ports ... [y, n, q]
```

- 7 You are asked questions regarding the frequency of VVR statistics collection.

- 8 The next phase of the configuration procedure consists of setting up a centrally managed host:

```
Enable Centralized Management? [y, n, q]
```

- 9 If you select centralized management, you are asked a series of questions relating to hostnames.

- 10 And finally, you see the message:

```
Startup completed successfully on all systems
```

After the installation and configuration of VxVM is complete, you can use the vxdiskadm command and the VEA GUI to create disk groups, and to populate these with disks.

See the *Veritas Volume Manager Administrator's Guide* and the VEA online help for details.

Upgrading Storage Foundation

This section contains procedures for upgrading Veritas Storage Foundation:

- “[Installation and upgrade requirements](#)”
- “[Upgrading Veritas Storage Foundation or Veritas Storage Foundation Oracle from 5.0 or 5.0 MP1 to 5.0 MP2](#)” on page 39

Installation and upgrade requirements

Installing the SF 5.0 MP2 Veritas software overwrites Veritas File System 4.1. After installing the SF 5.0 MP2 version of Veritas File System, to revert back to the 4.1 version you must uninstall the SF 5.0 MP2 version and then re-install the 4.1 version. The SF 4.1 and 5.0 MP2 versions of Veritas File System cannot coexist.

Upgrading Veritas Storage Foundation or Veritas Storage Foundation Oracle from 5.0 or 5.0 MP1 to 5.0 MP2

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/vxsvctrl status
Current state of server : RUNNING
```
- 2 Stop the VEA server.

```
# /opt/VRTS/bin/vxsvctrl stop
```
- 3 Verify the status of the VEA Service again.

```
# /opt/VRTS/bin/vxsvctrl 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.

```
# 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 the following command:

```
# /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
```

Upgrading Veritas Storage Foundation HA from 5.0 or 5.0 MP1 to Storage Foundation HA 5.0 MP2

To verify Storage Foundation HA 5.0 or 5.0 MP1 to Storage Foundation HA 5.0 MP2

- 1 Log in as superuser on one of the systems for installation and stop the cluster using `hastop -all`.

- 2 Stop the VEA service on both the nodes

```
# /opt/VRTS/bin/vxsvctrl status  
# /opt/VRTS/bin/vxsvctrl stop  
# /opt/VRTS/bin/vxsvctrl status
```

You must backup all configuration files like `main.cf`, `llhosts`, `lltab`, `gabtab`, `types.cf`, and any other agent configuration files.

- 3 Insert the disc containing the SF 5.0 MP2 software into the disc drive of one of the cluster nodes.

- 4 Mount the disc on a suitable mount point.

- 5 Navigate to the folder containing the `installmp` script.

- 6 Install SFHA 5.0 MP2 using the `installmp` script.

```
./installmp [-rsh]
```

- 7 After the initial system checks and the requirements checks are complete, press Return to start installing the patches.

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

- 9 Execute the following command to restart your systems:

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

- 10 Verify if VEA process has started.

```
# /opt/VRTS/bin/vxsvctrl status
```

- 11 If the VEA process has not started, then start the process.

```
# /opt/VRTS/bin/vxsvctrl start
```

Upgrading Veritas Storage Foundation Oracle HA from 5.0 or 5.0 MP1 to 5.0 MP2

To upgrade Veritas Storage Foundation Oracle HA from 5.0 or 5.0 MP1 to 5.0 MP2

- 1 Log in as superuser on one of the systems for installation and stop the cluster using `hastop -all`.

- 2 Stop the VEA service on both the nodes

```
# /opt/VRTS/bin/vxsvctrl status
# /opt/VRTS/bin/vxsvctrl stop
# /opt/VRTS/bin/vxsvctrl status
```

You must backup all configuration files like main.cf, llthosts, llttab, gabtab, types.cf, and any other agent configuration files.

- 3 Insert the disc containing the SF 5.0 MP2 software into the disc drive of one of the cluster nodes.

- 4 Mount the disc on a suitable mount point.

- 5 Navigate to the folder containing the installmp script.

- 6 Install SFORAHA 5.0 MP2 using the installmp script.

```
./installmp [-rsh]
```

- 7 After the initial system checks and the requirements checks are complete, press Return to start installing the patches.

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

- 9 Execute the following command to restart your systems:

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

- 10 Perform any regular post install steps related to VRTSdbd permissions and repository upgrade. For example:

```
# chown oracle:dba /opt/VRTSdbd
# chmod 750 /opt/VRTSdbd
```

- 11 To upgrade the repository in a cluster environment, do the following in the order presented:

- Save the repository disk group name, mount point, and other attributes.

- Unconfigure the SFUA repository from the VCS configuration.

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

The preceding command does the following:

- Removes the repository information from the VCS configuration

- Unmounts the repository

- Deports the disk group and takes the volumes offline

- Reconfigure the repository.

```
# vxdg import sfua_rep
```

```
# vxvol -g sfua_rep startall
```

- Mount the repository file system manually.

```
# mount -F vxfs /dev/vx/dsk/sfua_rep/dbed_rep /sfua_rep
```

- Run the repository upgrade command again without any options.

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

- 12 Check if VEA process has started.

```
# /opt/VRTS/bin/vxsvcctrl status
```
- 13 If the VEA process has not started, then start the process.

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

Preparing to upgrade the Veritas software

Ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/etc/fstab`. You must 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.

If you are upgrading an HA cluster, follow the guidelines given in the *Veritas Cluster Server (VCS) Installation Guide* for information on preserving your VCS configuration across the upgrade procedure. In particular, you must backup configuration files, such as `main.cf` and `types.cf`, in the `/etc/VRTSvcs/conf/config` directory. Additional configuration files, such as `OracleTypes.cf`, may also be present in this directory if you have installed any VCS agents. You must also back up these files.

For VxVM 3.5, disk group version 90 is the default disk group version.

For VxVM 4.1, disk group version 120 is the default disk group version.

VxVM 5.0 MP2 supports 90 and 120 disk group versions.

All the disk groups created by VxVM 3.5 or VxVM 4.1 can be imported after upgrading to VxVM 5.0 MP2. However, certain new features and tasks work only with the current disk group versions. To use new features, the disk group versions must be upgraded.

To prepare for the Veritas software upgrade

- 1 Log in as superuser.
- 2 Perform any necessary preinstallation checks and configuration.
- 3 If you are upgrading Veritas Storage Foundation for Oracle, resynchronize all existing snapshots before upgrading.

```
# /opt/VRTSvcs/bin/hagrp -offline service_group \
-sys system_name
```
- 4 Use the `vxlicrep` command to make a record of the currently installed Veritas licenses. Print the output or save it on a different system.
- 5 If you are upgrading a high availability (HA) product, take all service groups offline.
List all service groups.

```
# /opt/VRTSvcs/bin/hagrp -list
```

For each service group listed, take it offline.

```
# /opt/VRTSvcs/bin/hagrp -offline service_group \
-sys system_name
```

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

```
# df -F vxfs
```

- 7 Unmount all Storage Checkpoints and VxFS file systems.

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

- 8 Verify that all file systems have been cleanly unmounted

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

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

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

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

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

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

```
file system device busy
```

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.
```

- 10 (Optional) If an extended operation is pending, you must leave the file system mounted longer so the operation can finish. Removing a very large fileset clone can take several hours.

- 11 (Optional) Repeat [step 8](#) on page 43 to verify that the unclean file system is clean.

- 12 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.

- 13 Stop all the 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
```
- 14 If any external Array Policy Modules (APMs) are installed, uninstall the APMs. Most of the APMs were part of the base product in SF 5.0 and may not be required after the upgrade. Use following command to get the list of APMs installed:
- ```
# vxdmadm listapm all
```
- 15 Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the /etc/fstab file. You need to recreate these entries in the /etc/fstab file on the freshly installed system.

Upgrading Veritas Storage Foundation HA from 4.1MP2 to 5.0 MP2

- 1 Stop the VEA service on both the nodes.

```
# /opt/VRTS/bin/vxsvctrl status
# /opt/VRTS/bin/vxsvctrl stop
# /opt/VRTS/bin/vxsvctrl status
```

You must backup all configuration files like main.cf, llhosts, lltab, gabtab, types.cf, and any other agent configuration files.
- 2 Offline all application service groups.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
```
- 3 Insert the disc containing the SFHA 5.0 MP2 software into the disc drive of one of the cluster nodes.
- 4 Mount the disc on a suitable mount point.
- 5 Navigate to the folder containing the installer script.
- 6 Install SFHA 5.0 MP2 using the installer script.

```
./installer [-rsh]
```
- 7 Select I to upgrade to SF 5.0 MP2.
- 8 Select the appropriate number for the SF product.
- 9 During the installation process enter y for the question similar to the following:
Are you sure you want to upgrade using current configuration
- 10 Restart all nodes.
- 11 Use ./installsf [-rsh] -configure to configure the upgraded stack.
- 12 Check if VEA process has started.

```
# /opt/VRTS/bin/vxsvctrl status
```
- 13 If VEA process has not started, then start the process using

```
# /opt/VRTS/bin/vxsvctrl start
```

Upgrading Veritas Storage Foundation Oracle HA from 4.1 MP2 to 5.0 MP2

To upgrade Veritas Storage Foundation Oracle High Availability from 4.1 MP2 to 5.0 MP2

- 1 Log in as superuser.
- 2 Resynchronize all existing snapshots before upgrading (if any).

```
# /opt/VRTS/bin/dbed_vmsnap -s $ORACLE_SID -f SNAPPLAN -o \
resync
```
- 3 Use the vxlicrep command to make a record of the currently installed Veritas licenses.
- 4 Take all service groups offline.
List all service groups.

```
# /opt/VRTSvcs/bin/hagrp -list
```

For each service group listed, take it offline.

```
# /opt/VRTSvcs/bin/hagrp -offline service_group -sys system_name
```
- 5 Upgrade the cluster to SF 5.0 MP2 using the installer script.

```
./installer [-rsh]
```
- 6 Reboot the upgraded systems.
- 7 Reinstate the mount points in the /etc/fstab file that you recorded in step 9.
- 8 If Oracle DB is under VCS control and if the path for OracleTypes.cf is not /etc/VRTSvcs/conf/config, then systems will be in STALE ADMIN WAIT state after the upgrade.
- 9 To bring the cluster systems in RUNNING state, do the following in the order presented:
 - Stop the cluster.

```
# hastop -all
```
 - Edit the main.cf to have the proper path for including the OracleTypes.cf
 - Start the cluster on all nodes.

```
# hastart
```
- 10 Configure the systems.

```
# cd /opt/VRTS/install
```

To configure the product, use the installation script with the -configure option.

```
# ./installsfora [-rsh] -configure
```

- 11 Perform the other post-install operations, such as dbed repository creation and so on.

```
# chown oracle:dba /opt/VRTSdbed  
# chmod 750 /opt/VRTSdbed
```

- 12 To configure the repository in a cluster environment, run the repository config command.

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

- 13 To configure the repository for SF Oracle, make sure you have the following:

- A mount point of already mounted Veritas Volume on a shared storage, with Veritas File system
- A public NIC used by each system in the cluster
- A Virtual IP address and netmask

- 14 Check if VEA process has started.

```
# /opt/VRTS/bin/vxsvctrl status
```

- 15 If the VEA process has not started, then start the process.

```
# /opt/VRTS/bin/vxsvctrl start
```

Configuring your system after the installation

Use the following procedure to configure your system after installation.

To configure your system after the software upgrade

- 1 Reinstate the mount points in the `/etc/fstab` file that you recorded in the preparation steps.
- 2 Reboot the upgraded systems.
- 3 Restart all the volumes by entering the following command for each disk group:

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

Optional configuration steps

Perform the following optional configuration steps:

- If you want to use features of Veritas Storage Foundation 5.0 MP2 or Veritas Storage Foundation 5.0 MP2 for Oracle 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.
- Stop the cluster, restore the VCS configuration files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.
- To create root volumes that are under VxVM control after installation, use the `vxcp_lvmroot` command.

- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions. See the *Veritas Volume Manager Administrator's Guide* and *Veritas File System Administrator's Guide*
- 4 After you complete the installation procedure, proceed to initializing (where required), setting up, and using Veritas Storage Foundation.

Verifying the Storage Foundation installation

This section includes the following topics:

- [“Verifying that the products were installed” on page 47](#)
- [“Installation log files” on page 47](#)
- [“Using the installation log file” on page 48](#)
- [“Using the response log file” on page 48](#)
- [“Using the summary file” on page 48](#)
- [“Verifying Volume Manager Installation” on page 48](#)
- [“Checking Volume File System installation” on page 49](#)
- [“Command installation verification” on page 49](#)

Verifying that the products were installed

Verify that the Veritas Storage Foundation products are installed.

You can use the `swlist` command to check which packages have been installed:

```
# swlist -l product | grep VRTS
```

Use the following sections to further verify the product installation.

Installation log files

After every product installation, the installer creates three text files:

- Installation log file
- Response file
- Summary file

The name and location of each file is displayed at the end of a product installation, and are always located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

Using the installation log file

The installation log file contains all commands executed during the procedure, their output, and errors generated by the commands. This file is for debugging installation problems and can be used for analysis by Veritas Support.

Using the response log file

The response file contains the configuration information that you entered during the procedure. You can use the response file for future installation procedures by invoking an installation script with the `responsefile` option. The response file passes arguments to the script to automate that product's installation. You can edit the file to automate installation and configuration of additional systems.

Using the summary file

The summary file contains the results of the installation by the common product installer or product installation scripts. The summary includes the list of the packages and patches, and the status (success or failure) of each package or patch. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

Verifying Volume Manager Installation

Use the following procedure to verify that Volume Manager processes are running.

To verify that key Volume Manager processes are running

- 1 Type the following command:

```
# ps -e | grep vx
```

Entries for the `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxsmf`, `vxpal`, `vxcached`, and `vxconfigbackupd` processes appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

- 2 Verify availability of APMs installed on the system by the following command:

```
# vxdmnpadm listapm all
```

It is recommended that you install Array Policy Modules (APMs) which were used earlier but do not appear in the list.

Checking Volume File System installation

The Veritas File System package consists of a kernel component and administrative commands.

Command installation verification

The Veritas File System commands are installed in the following directories:

- `/sbin/fs/vxfs5.0`
Contains the Veritas mount command required to mount file systems.
- `/opt/VRTS/bin`
Contains symbolic links to all Veritas-specific commands installed in the directories listed above.

Determine whether the following subdirectories are present.

```
# ls /sbin/fs/vxfs5.0
# ls /opt/VRTS/bin
```

Make sure you have adjusted your environment variables accordingly.

Uninstalling Veritas Storage Foundation products

This section includes the following topics:

- “[Summary of Veritas Storage Foundation uninstallation tasks](#)” on page 50
- “[Dropping the repository database for Oracle](#)” on page 50
- “[Shutting down cluster operations](#)” on page 51
- “[Removing VxFs file systems and Storage Checkpoints](#)” on page 51
- “[Removing the root disk from VxVM](#)” on page 52
- “[Moving volumes to disk partitions](#)” on page 52
- “[Shutting down Veritas Volume Manager](#)” on page 57
- “[Uninstalling Veritas Storage Foundation packages](#)” on page 57
- “[Uninstalling Veritas Volume Manager](#)” on page 58
- “[Removing the Veritas Enterprise Administrator client](#)” on page 58
- “[Uninstalling the SF 5.0 MP2 patches](#)” on page 59

For instructions on uninstalling Veritas Volume Replicator, see *Veritas Volume Replicator Installation Guide*.

Summary of Veritas Storage Foundation uninstallation tasks

Complete the following preparations to uninstall a Veritas Storage Foundation product.

Warning: Failure to follow the preparations that are outlined in this chapter can result in loss of data.

Uninstallation of Veritas Storage Foundation products consists of the following tasks:

- If you are uninstalling Veritas Storage Foundation for Oracle, dropping the repository database
- Shutting down cluster operations
- Removing the root disk from VxVM control
- Removing VxFS file systems and Storage Checkpoints
- Moving volumes to disk partitions
- Removing the Veritas Storage Foundation
- Removing the license files (optional)

Warning: Failure to follow the instructions in the following sections may result in unexpected behavior.

After you uninstall Veritas Volume Manager, your machine will not have volume management software.

To uninstall Veritas Storage Foundation and Veritas Storage Foundation for Oracle, refer to the appropriate procedure.

Dropping the repository database for Oracle

When you uninstall Veritas Storage Foundation for Oracle, drop the repository database. If you want to recreate the repository database, you can drop the existing repository database using these steps.

To drop the repository database in a stand-alone configuration

- 1 Make sure the repository database volume is mounted using the `df` command.

If the repository database volume is not mounted, run the `sfua_rep_mount` command to mount the volume.

```
# /sbin/init.d/sfua_rep_mount
```

- 2 Use the sfua_db_config command with the -o dropdb option to remove the database.

```
# /opt/VRTS/bin/sfua_db_config -o dropdb
```

To drop the repository database in an Oracle cluster or Oracle RAC configuration

- 1 Drop the repository database from the VCS configuration and deport the repository disk group.

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

- 2 Import the repository database disk group.

```
# /opt/VRTS/bin/vxdg import repository_diskgroup_name
```

- 3 Run the sfua_rep_mount command to mount the repository database volume.

```
# /sbin/init.d/sfua_rep_mount start
```

- 4 Use the sfua_db_config command with the -o dropdb option to remove the database.

```
# /opt/VRTS/bin/sfua_db_config -o dropdb
```

Shutting down cluster operations

If the systems are running as an HA cluster, you have to take all service groups offline and shutdown VCS.

To take all service groups offline and shutdown VCS

- Use the hastop command:

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

Warning: Do not use the -force option when executing hastop. This leaves all service groups online and shuts down VCS, causing undesired results when you uninstall the packages.

Removing VxFS file systems and Storage Checkpoints

It is advisable to unmount any user VxFS file systems before you uninstall VxFS to help smooth uninstallation of VxVM package if the VxFS file system is mounted on VxVM volumes. System partitions need not be unmounted as part of this operation. After you remove the VRTSvxfs package, VxFS file system versions greater than those supported by OnlineJFS bundled with HP-UX OS are not mountable or accessible until another VRTSvxfs package supporting them is installed.

To unmount a file system

- 1 Check if any VxFS file systems are mounted.

```
# cat /etc/mnttab | grep vxfs
```

- 2 Unmount any file systems that are not system partitions.

```
# umount special | mount_point
```

Specify the file system to be unmounted as a `mount_point` or `special` (the device on which the file system resides).

See the `umount_vxfs(1M)` manual page.

If you are using the VxFS file system, system partitions need not be unmounted.

To unmount a Storage Checkpoint

- 1 Check if any Storage Checkpoints are mounted.

```
# cat /etc/mnttab | grep vxfs
```

- 2 Unmount any Storage Checkpoints.

```
# umount /checkpoint_name
```

Removing the root disk from VxVM

If the system's root disk is under VxVM control, use the following command to copy its contents to a new LVM root disk:

```
# /etc/vx/bin/vxres_lvmroot -v -b [-p c#t#d#2,c#t#d#3,...]  
c#t#d#
```

where `c#t#d#` is the access name of the new LVMroot disk. If the root disk volumes are distributed over several disks, use the `-p` option to specify a comma-separated list of additional disks that are to be used to set up the LVM root volume group. The operation to clone a new LVM root volume group can take some time, so the `-v` (verbose) option is specified to show how far this has progressed.

Moving volumes to disk partitions

You must move all volumes to disk partitions. You can do this using one of the following procedures:

- Back up the system fully onto tape and then recover from it.
- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Use VxVM to move volumes incrementally onto disk partitions as described in the following section.

Moving volumes onto disk partitions for HP-UX

Use the following procedure to move volumes to disk partitions.

To move volumes to disk partitions

- 1 Evacuate disks using `vxdiskadm`, the GUI, or the `vxevac` script. Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.
- 2 Remove the evacuated disks from VxVM control by entering the following:

```
# vxdg rmdisk diskname
# vxdisk rm devname
```
- 3 Decide which volume to move first, and if the volume is mounted, unmount it.
- 4 If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that you have applied the `sync` command to the data on the volume.
- 5 Create a partition on free disk space of the same size as the volume using the `format` command. If there is not enough free space for the partition, add a new disk to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this first volume.
- 6 Copy the data on the volume onto the newly created disk partition using a command such as `dd`.

```
# dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/dsk/c2t2d2
```

where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly created partition.
- 7 Replace the entry for that volume (if present) in `/etc/fstab` with an entry for the newly created partition.
- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Remove the volume from VxVM using the following command:

```
# vxedit -rf rm volume_name
```
- 10 Remove any free disks (those having no subdisks defined on them) by removing the volumes from VxVM control.

To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.

```
# vxprint -F '%sdnum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VxVM control.

```
# vxdg rmdisk diskname
```

```
# vxdisk rm devname
```

Use the free space created for adding the data from the next volume you want to remove.

- 11 After you successfully convert all volumes into disk partitions, reboot the system.
- 12 After the reboot, make sure none of the volumes are open.

```
# vxprint -Aht -e v_open
```

If any volumes remain open, repeat the steps listed above.

Example of moving volumes to disk partitions on HP-UX

This example shows how to move the data on a volume to a disk partition. In the example, there are three disks: disk1 and disk2 are subdisks on volume vol01 and disk3 is a free disk. The data on vol01 is copied to disk3 using vxevac.

To move volumes to disk partitions on HP-UX

- 1 Display voldg content before the data on vol01 is copied to disk3.

```
# vxprint -g voldg -ht
DG NAME NCONFIG NLOG MINORS GROUP-ID
DM NAME DEVICE TYPE PRIVLEN PUBLEN STATE
RV NAME RLINK_CNT KSTATE STATE PRIMARY DATAVOLS SRL
RL NAME RVG KSTATE STATE REM_HOST REM_DG REM_RLNK
V NAME RVG KSTATE STATE LENGTH READPOL PREFPLEX
UTYPE
PL NAME VOLUME KSTATE STATE LENGTH LAYOUT NCOL/WID
MODE
SD NAME PLEX DISK DISKOFFS LENGTH [COL/]OFF DEVICE
MODE
SV NAME PLEX VOLNAME NVOLLAYR LENGTH [COL/]OFF AM/NM
MODE
DC NAME PARENTVOL LOGVOL
SP NAME SNAPVOL DCO
dg voldg default default 115000
1017856044.1141.hostname.veritas.com

dg voldg default default 115000
1017856044.1141.hostname.veritas.com
dm disk1 c1t12d0 auto:hpdisk 2591 17900352 -
dm disk2 c1t14d0 auto:hpdisk 2591 17899056 -
dm disk3 c1t3d0 auto:hpdisk 2591 17899056 -
v vol1 - ENABLED ACTIVE 4196448 ROUND -
fsgen
pl pl1 vol1 ENABLED ACTIVE 4196448 CONCAT -

RW
sd sd1 pl1 disk1 0 2098224 0 c1t12d0
ENA
sd sd2 pl1 disk2 0 2098224 2098224 c1t14d0
```

ENA

2 Evacuate disk1 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk1 disk3
# vxprint -g voldg -ht
DG NAME NCONFIG NLOG MINORS GROUP-ID
DM NAME DEVICE TYPE PRIVLEN PUBLEN STATE
RV NAME RLINK_CNT KSTATE STATE PRIMARY DATAVOLS SRL
RL NAME RVG KSTATE STATE REM_HOST REM_DG REM_RLINK
V NAME RVG KSTATE STATE LENGTH READPOL PREFPLEX
UTYPE
PL NAME VOLUME KSTATE STATE LENGTH LAYOUT NCOL/WID
MODE
SD NAME PLEX DISK DISKOFFS LENGTH [COL/]OFF DEVICE
MODE
SV NAME PLEX VOLNAME NVOLLAYR LENGTH [COL/]OFF AM/NM
MODE
DC NAME PARENTVOL LOGVOL
SP NAME SNAPVOL DCO
dg voldg default default 115000
1017856044.1141.hostname.veritas.com

dm disk1 c1t12d0 auto:hpdisk 2591 17900352 -
dm disk2 c1t14d0 auto:hpdisk 2591 17899056 -
dm disk3 c1t3d0 auto:hpdisk 2591 17899056 -
v vol1 - ENABLED ACTIVE 4196448 ROUND -
fsgen
pl pl1 vol1 ENABLED ACTIVE 4196448 CONCAT -
RW
sd disk3-01 pl1 disk3 0 2098224 0 c1t3d0
ENA
sd sd2 pl1 disk2 0 2098224 2098224 c1t14d0
ENA
```

3 Evacuate disk2 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk2 disk3
# vxprint -g voldg -ht
DG NAME NCONFIG NLOG MINORS GROUP-ID
DM NAME DEVICE TYPE PRIVLEN PUBLEN STATE
RV NAME RLINK_CNT KSTATE STATE PRIMARY DATAVOLS SRL
RL NAME RVG KSTATE STATE REM_HOST REM_DG REM_RLINK
V NAME RVG KSTATE STATE LENGTH READPOL PREFPLEX
UTYPE
PL NAME VOLUME KSTATE STATE LENGTH LAYOUT NCOL/WID
MODE
SD NAME PLEX DISK DISKOFFS LENGTH [COL/]OFF DEVICE
MODE
SV NAME PLEX VOLNAME NVOLLAYR LENGTH [COL/]OFF AM/NM
MODE
DC NAME PARENTVOL LOGVOL
SP NAME SNAPVOL DCO
dg voldg default default 115000
```

```
1017856044.1141.hostname.veritas.com

dm disk1 c1t12d0 auto:hpdisk 2591 17900352 -
dm disk2 c1t14d0 auto:hpdisk 2591 17899056 -
dm disk3 c1t3d0 auto:hpdisk 2591 17899056 -
v vol1 - ENABLED ACTIVE 4196448 ROUND -
fsgen
pl p11 vol1 ENABLED ACTIVE 4196448 CONCAT -
RW
sd disk3-01 p11 disk3 0 2098224 0 c1t3d0
ENA
sd disk3-02 p11 disk3 2098224 2098224 2098224 c1t3d0 ENA
```

4 Remove the evacuated disks from VxVM control.

```
# vxdisk -g voldg list
DEVICE TYPE DISK GROUP STATUS
c1t3d0 auto:hpdisk disk3 voldg
online
c1t12d0 auto:hpdisk disk1 voldg
online
c1t14d0 auto:hpdisk disk2 voldg
# vxdg rmdisk disk1
# vxdg rmdisk disk2
# vxdisk rm c1t12d0
# vxdisk rm c1t14d0
```

5 Verify that the evacuated disks have been removed from VxVM control.

```
# vxdisk -g voldg list
DEVICE TYPE DISK GROUP STATUS
c1t3d0 auto:hpdisk disk3 voldg online
```

6 Check to see whether the volume you want to move first is mounted.

```
# mount | grep vol1
/vol1 on /dev/vx/dsk/voldg/vol1
read/write/setuid/log/nolargefiles/dev=12dc138 on Wed Apr 3
10:13:11 2002
```

7 Copy the data on vol01 to the newly created disk partition.

```
# dd if=/dev/vx/dsk/voldg/vol01 of=/dev/dsk/c1t12d0
```

8 In the /etc/fstab file, remove the following entry.

```
/dev/vx/dsk/voldg/vol1 /dev/vx/rdsck/voldg/vol1 /vol1 vxfs 4
yes rw
```

9 Replace it with an entry for the newly created partition.

```
/dev/dsk/c1t12d0 /dev/rdsck/c1t12d0 /vol01 vxfs 4 yes rw
```

10 Mount the disk partition.

```
# mount -F vxfs /dev/dsk/c1t12d0 /vol01
```

11 Remove vol01 from VxVM.

```
# vxedit -rf rm /dev/vx/dsk/voldg/vol01
```

To complete the procedure, perform the remaining steps.

Shutting down Veritas Volume Manager

Use the following procedure to shut down Veritas Volume Manager.

To shut down Veritas Volume Manager

- ◆ Enter the vxldctl and vxiod commands:

```
# vxldctl stop
# vxiod -f set 0
```

Uninstalling Veritas Storage Foundation packages

Use the following procedure to shut down and remove the installed Veritas Storage Foundation packages.

To shut down and remove the installed Veritas Storage Foundation packages

- 1 In a stand-alone configuration, if you are uninstalling Veritas Storage Foundation for Oracle, stop the repository database and unmount the database repository volume.

```
# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb
# /sbin/init.d/sfua_rep_mount stop
```
- 2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system, enter the following:

```
# hastop -local
```

To stop VCS processes on all systems, enter the following:

```
# hastop -all
```

- 3 Move to the /opt/VRTS/install directory and run the uninstall script.

```
# cd /opt/VRTS/install
```

For Veritas Storage Foundation, enter the following:

```
# ./uninstallsf
```

For Veritas Storage Foundation for Oracle, enter the following:

```
# ./uninstallsfora
```

You can use these commands to remove the packages from one or more systems.

To remove packages from remote systems, configure `ssh` or `rsh`.

Depending on the choices that you made when you installed VxVM, not all these packages may be installed on your system.

If you have obtained a Veritas product from an electronic download site, the single product download files do not contain the `uninstallsf` installation script, so you must use the product uninstallation script to uninstall the product. For example, if you download Veritas Volume Manager, use the `uninstallvm` script instead `uninstallsf`.

Uninstalling Veritas Volume Manager

This section describes how to uninstall Veritas Volume Manager and the product license.

To uninstall Veritas Storage Foundation or Veritas Storage Foundation for Oracle, use the following procedures.

See “[Uninstalling Veritas Storage Foundation packages](#)” on page 57.

To uninstall Veritas Volume Manager

- 1 Log in as superuser.
- 2 Run the `installer` command to uninstall Veritas Volume Manager. For example, enter the following:

```
# cd /dvdrrom  
# ./installer
```

- 3 From the product installer, choose the **U** option for Uninstall, and select Veritas Volume Manager.
- 4 Enter one or more system names from which Veritas Volume Manager is to be uninstalled. For example:

Enter the system names separated by spaces on which to uninstall VxVM: **system01**

- 5 After the system check completes successfully, press Return to continue.
- 6 Enter **y** to uninstall the VxVM. For example:
Are you sure you want to uninstall VxVM? [y,n,q] (y)
- 7 After VxVM successfully stops, the system tells you the location of the log files. You must save these files for future reference.
- 8 After uninstallation completes, reboot the system.

```
# /usr/sbin/shutdown
```

Removing the Veritas Enterprise Administrator client

You must also remove the client software from any of the machines you have used to access the Veritas software.

To remove the VEA client from an HP-UX system other than the server

- 1 Stop the VEA Service.
`# /opt/VRTS/bin/vxsvctrl stop`
- 2 Use the `swremove` command to remove the VEA client packages.
`# swremove VRTSobgui VRTSat VRTSpbx VRTSicsco`

To remove the VEA client from a Windows system

- 1 Log in as the database administrator.
- 2 Select **Start > Settings > Control Panel**.
- 3 Double-click **Add/Remove Programs** to list the installed products.
- 4 Select **Veritas Enterprise Administrator** from the list, and click the **Remove** button.
- 5 When a dialog box appears asking you to confirm the removal, click **Yes**.

Uninstalling the SF 5.0 MP2 patches

The following procedure removes the patches if you want to uninstall SF 5.0 MP2. It is recommended to refer the Installation Summary Files for the list of patches that get installed during installation.

To remove the 5.0 MP2 patches

- 1 Log in as superuser.
- 2 Stop the cluster:

```
# hastop -all
```
- 3 On all the nodes, stop VxFEN:

```
# /sbin/init.d/vxfen stop
```
- 4 On all the nodes, stop the currently running VxPAL agents. See `vxpalcctrl` (1M):
 - Stop the storage agent:

```
# /opt/VRTSobc/pal33/bin/vxpalcctrl -a StorageAgent -c stop
```
 - Check the status of the storage agent:

```
# /opt/VRTSobc/pal33/bin/vxpalcctrl -a StorageAgent -c status
```
 - Stop the action agent:

```
# /opt/VRTSobc/pal33/bin/vxpalcctrl -a actionagent -c stop
```
 - Check the status of the action agent:

```
# /opt/VRTSobc/pal33/bin/vxpalcctrl -a actionagent -c status
```
 - Stop the gridnode agent:

```
# /opt/VRTSobc/pal33/bin/vxpalcctrl -a gridnode -c stop
```
 - Check the status of the gridnode agent:

```
# /opt/VRTSobc/pal33/bin/vxpalcctrl -a gridnode -c status
```
- 5 On all the nodes, remove all the patches using the `swremove` command:

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

Symantec recommends that all the patches installed during SF 5.0 MP2 installation be removed through a single command line. The system automatically reboots after removing the patches.

Software limitations

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

Storage Foundation software limitations

Software limitations in the SF 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 Volume Manager software limitations

Common Interswitch Linking (ISL) for heartbeat and storage traffic in Campus Cluster and EDC configurations

It is recommended to have common Interswitch Linking (ISL) for both heartbeat and storage traffic in Campus Cluster and Extended Distance Cluster (EDC) configurations.

In case of configurations involving separate links for heartbeat and storage traffic, if the storage link fails and the heartbeat link continues between the two sites, each node loses connectivity to storage on the remote site. In such cases, the volumes could be left in variety of states depending on the I/Os at the time of failure. Therefore, a manual intervention is required as a generic recovery procedure may not work.

Cluster Volume Manager (CVM) fail back behavior for non-Active/Active arrays

This describes the fail back behavior for non-Active/Active arrays in a CVM cluster. This behavior applies to A/P, A/PF, APG, A/A-A, and ALUA arrays.

When all of the Primary paths fail or are disabled in a non-Active/Active array in a CVM cluster, the cluster-wide failover is triggered. All hosts in the cluster start using the Secondary path to the array. When the Primary path is enabled, the hosts fail back to the Primary path.

However, suppose that one of the hosts in the cluster is shut down or disabled while the Primary path is disabled. If the Primary path is then enabled, it does not trigger failback. The remaining hosts in the cluster continue to use the Secondary path. When the disabled host is rebooted and rejoins the cluster, all of the hosts in the cluster continue using the Secondary path. This is expected behavior.

If the disabled host is rebooted and rejoins the cluster before the Primary path is enabled, enabling the path does trigger the failback. In this case, all of the hosts in the cluster fail back to the Primary path. [e1441769]

Node and host name expansion

As part of the Software Pack (SPK) for HP-UX 11i v2 of May 2005, HP provided an optional Node and Host Name Expansion (NodeHostNameXpnd) product bundle, revision B.11.23.01. This product bundle enables the capability to create node and host names of up to 255 bytes. The current limitation is 8 bytes for node names and 64 bytes for host names. Installation of this product bundle does not automatically activate 255 byte length support for node and host names. The default OS configuration is still 8 bytes for node names and 64 bytes for host names. You must enable a dynamic kernel tunable parameter, `expanded_node_host_names`, using the `kctune` command to allow the use of larger names on the system. Veritas Storage Foundation 5.0 products for HP-UX were not tested with the Host Name Expansion bundle and therefore do not support node and host name expansion.

Mixed PA-RISC and Itanium clusters are not supported

Within a cluster, all systems must run on the same processor type and use the same operating system version and patch level.

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 [599164]

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.

Differing locales produces unintelligible characters in GUI [605487]

The GUI does not support Oracle users having a different locale than the superuser's locale. The GUI displays unintelligible characters if the SFDB repository server starts with a locale that differs from the Oracle user locale (client).

Some features stop working after a GCO failover [563603]

Some features for Veritas Storage Foundation for Oracle do not work correctly after a Global Cluster (GCO) Failover. In SF 5.0, the Veritas 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) do not function normally after a failover. However, features such as Oracle Disk Manager (ODM), Quick I/O, and Concurrent I/O (CIO) continue to work after a failover. This issue will be fixed after the next release.

Do not use UNIX VEA via PC-based UNIX emulators

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

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 see the error 4646.

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

- 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 [32726]

- You cannot create a clone database using a mounted Storage Checkpoint.

VEA limitations [34446]

- 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.
- 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 the following 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 panic your system. [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_vmclosedb` command, the clone database's ORACLE_SID can be only eight characters or less. You get an error (ERROR V-81-5713) if the ORACLE_SID is more than 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 Veritas Storage Foundation for Oracle scripts can unset it at runtime.
- 4 Avoid starting Veritas 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 SQL*Plus startup time.
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 SQL*Plus startup time.

DBDST class names limited to 29 characters [601746]

The `dbdst_admin -o rmclass` command fails when attempting to remove a class name of 30 characters or more. The maximum class name length is 29 characters.

Selected utilities require setuid [643964]

Some Veritas Storage Foundation for 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 queries. For these reasons, the following programs are setuid-enabled:

- /opt/VRTSdbd/.dba/dbed_analyzer
- /opt/VRTSdbd/.dba/vxckptplan
- /opt/VRTSdbd/.dba/vxstorage_stats
- /opt/VRTSdbcom/.dba/vxdbd_start

Multiple archive log destinations with RAC [795617]

Multiple archive log locations are not supported in RAC configurations.

Repository hostnames are case-insensitive [859863]

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, in the local subnet. Errors may occur if the repository host name differs from another host name only by case.

One-time scheduled tasks need Specific Date [861274]

When scheduling a one-time task from the GUI, the task may not be executed if a Specific Date (Include Date) is not set for it.

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 is not supported in the next major release of Veritas Storage Foundation.

Fixed issues

The following sections describe some of the critical Veritas Storage Foundation issues that were fixed in this release.

Veritas Storage Foundation fixed issues

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

Incident	Description
1449870	During the installation of Veritas Storage Foundation Manager VRTSvsvc package is removed from common product installer.

Veritas Volume Manager fixed issues

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

Incident	Description
1089876	<p>The fix for JAGaf60923 introduced the problem importing the disk groups with vol_maxspecialio and/or vol_maxio tunable value being smaller than 1024 or if an underlying disk driver does not support as large as 1MB I/O size.</p> <p>The typical error observed during the disk group import is:</p> <pre>VxVM vxconfigd ERROR V-5-1-569 Disk group rootdg, Disk c0t0d0s2: Cannot auto-import group: Disk group has no valid configuration copies</pre> <p>The correct fix does not have any regression if the tunable parameters reduced by the user from their default values of 1024 or if an underlying disk driver maximum supported I/O size is less than 1MB.</p>
1274157	<p>While running VxVM test cases, vxassist command breaks its connection with vxconfigd. Upon failure, following error messages are seen:</p> <pre>VxVM vxassist ERROR V-5-1-684 IPC failure: Configuration daemon is not accessible VxVM vxassist ERROR V-5-1-10128 Configuration daemon is not accessible</pre> <ul style="list-style-type: none"> ■ Modify the backlog by setting it to a bigger value so more clients can listen to the socket. ■ Set VXV_RETRY_SET to true value, even if the backlog of the listen socket is full.
1274173	<p>After a master CVM node leaves the cluster, a slave node takes over as a master and has to recover shared volumes. Recovery of the volumes with DRL logs takes a significant amount of time and prevents the returning node from joining the cluster.</p> <p>The DRL maps recovery is first completed on all the shared volumes. Resynchronization is started after DRL recovery is complete.</p>

Incident	Description
1274218	<p>The <code>vxrdg free</code> command displays incorrect information about the offset and the length of the free space on a disk media.</p> <p>A check has been added for record IDs if the DA names do not match.</p>
1274220	<p>On a CVM cluster with two or more nodes, any configuration change transaction on an online imported shared disk group, like adding a new disk into a disk group or starting/stopping a volume, fails with the following error:</p> <pre>VxVM vxrdg ERROR V-5-1-10128 Cannot assign minor number</pre> <p>The secondary node recognizes that the minor numbers for shared disk group volumes have been temporarily changed.</p>
1275019	<p>Full resync is performed after force import of a dg where fast resync (FMR) must have been performed instead.</p> <p>Rather than leaving the determination of whether a disk has changed to kernel/vxio, we leave this to vxconfigd as it has the on-disk disked information. Then it flags kernel to use FMR if it determined that it is the same disk.</p>
1275036	<p>Under certain HA configurations, <code>vxassist(1M)</code> command after a relayout and a convert operations could lead to a non-HA friendly configuration. This could potentially cause the mirrored plexes of the volume to get disabled after a single controller failure.</p> <p>The <code>mirror=ctlr mirror=target</code> functionality has been fixed by reordering the way the volume components are laid out during relayout.</p>

Incident	Description
1296031	<p>In a CVM cluster, which has multiple mirrored volumes with Data Change Object (DCO) volume association, a node crash or a reboot prevents another node from joining the cluster. This issue has the following indicators:</p> <ol style="list-style-type: none"> 1) A vxio warning is displayed: <pre>VxVM vxio V-5-3-0 volcvm_needsync: object testvol2 in dg testdg is still marked as needing noderecover by node 1. Join will fail.</pre> <ol style="list-style-type: none"> 2) The vxconfigd(1M) daemon displays the error: <pre>'volume recovery in progress'</pre> <p>The correct diskgroup is set for lookup operations before searching the DCO record by name. Additionally, for faster recovery, the vxrecover utility uses the non-verbose mode to avoid numerous console messages when there are a lot of shared volumes.</p>
1296040	<p>In a CVM environment, the vxrecover(1M) command does not execute the vxvol noderecover operation for mirrored volumes with the Fast Mirror Resynchronization- 3 (FMR3) Dirty Region Logging (DRL). This can render subsequent node joining operations to fail with the following error:</p> <pre>"recovery in progress"</pre> <p>The code has been modified to identify the FMR DRL correctly, so that the vxvol noderecover operation is executed for mirrored volumes with the FMR3 DRL.</p>
1314428	<p>In case of a large storage configuration (2K LUNs with 2 paths each) the vxconfigd(1M) daemon takes a long time to complete (1-2 hrs) its startup.</p> <p>The DMP device open is cached till vxconfigd(1M) completely starts, that is, till all the devices are discovered.</p>
1379576	<p>While importing diskgroups with the -o useclosedev option, the vxconfigd(1M) daemon dumps core.</p> <p>During a refresh operation, all the existing references to objects are validated prior to use.</p>

Incident	Description
1177904	<p>After Unique Disk Identifier (UDID) update through <code>vxdisk updateudid</code> command on an imported diskgroup disk, the next VxVM command involving a transaction fails with an error.</p> <p>The code has been modified, so that the existence of the dmrec record is verified for the disk, and if available, it is used for reloading the disk in the kernel.</p>
1203522	<p>In the CVM environment, the auto fail over of a DMP path to secondary path does not happen the on primary path failure.</p> <p>The following errors can be seen in syslog:</p> <pre>kernel: SCSI error : <3 0 0 52> return code = 0x20000 kernel: end_request: I/O error, dev sdbb, sector 164530 kernel: end_request: I/O error, dev sdbb, sector 164538 vxvm:vxconfigd: V-5-1-0 kernel_read_disk: K_IOCTL (VOL_VOLDIO_READ) Failed errno 0x5</pre> <p>Now, the event notifications are sent to all the impacted nodes so that failovers occur when required.</p>
901010	<p>After adding a new path to a CLARiiON array online, <code>vxdisk scandisks</code> command may take more than an hour to complete if the CLARiiON array is in the Active/Passive Explicit Failover (A/PF) mode.</p> <p>DDL is modified to do a SCSI pass-through open instead of a regular open for discovery before issuing SCSI inquiry.</p>
1019507	<p>The <code>vxdisk list</code> command shows the capacity of the disk incorrectly.</p> <p>The correct disk geometry is calculated after comparing both the mode sense data and data obtained from the explicit DIOC_IOCTL ioctl(2).</p>

Incident	Description
1106031	<p>Consider a volume with two or more mirrors enabled with FMR2 (Fast Mirror Resync). If its DCO plex is detached due to a missing disk and an attempt is made to re-attach the DCO plex while one of its underlying disks is still missing, data written to the volume could be corrupted after volume recovery when the DCO plex is finally re-attached with the missing disk returns.</p> <p>If the DCO plex is in the detached state, the DCO map is not cleared even when I/O transactions are completed. It means that there will be resynchronization during volume recovery but it avoids data corruption.</p>
1106038	<p>The -k adddisk option of <code>vxldg(1M)</code> command gives the following error:</p> <pre>VxVM vxldg ERROR V-5-1-10127 associating disk-media 'disk1' with 'disk2': Serial Split Brain detected. Run vxsplittlines</pre> <p>In case of a transaction failure perform the rollback of the private header updates on all the disks in the disk group to make the header information consistent.</p>
1106040	<p>An enhancement to allow mixed Naming Schemes, namely, OSN (Operating System Name)/EBN (Enclosure based name) for third party discovered disks.</p> <p>An additional check is added for third party devices during device discovery for assigning disk names to ensure that the long CTD names default to EBN naming while retaining the mixed naming schemes.</p>
1113006	<p>On a system connected to an XP128 or XP1024 disk array managed by the bundled ASL (Array Support Library) for XP, <code>vxconfigd</code> may core dump at <code>set_asso()</code> when a command like <code>vxctrl enable,vxdisk scandisks</code> or <code>vxconfigd -k</code> is issued to have it re-scan the disks.</p> <p>The solution is to remove the extra free from the ASL code.</p>
1138767	<p><code>vxsnap(1M)</code> command core dumps during <code>vxsnap print</code> operation with a segmentation violation</p> <p>A check is introduced to verify that the snapshot volume record pointer is not NULL before using it further.</p>

Incident	Description
1158584	<p>Mirroring for rootdisk system volumes fail even with a Base-VxVM license.</p> <p>The solution is to rectify the check that identifies system volumes on the root disk and permit mirroring of rootdisk volumes with a Base-VxVM license.</p>
1158588	<p>Breaking off a root disk mirror with vxbrk_rootmir(1M) command fails with the following error messages:</p> <pre data-bbox="717 520 1342 684">vxbrk_rootmir fails with an error vxbrk_rootmir: 16:04: Checking specified disk(s) for presence and type vxbrk_rootmir: ERROR: rootdg disk group contains more than 1 disk (2)</pre> <p>The solution is to rectify the check that verifies whether volumes on the root disk mirror are the same as those present on the root disk and permit breaking off the mirror. Also, the code is modified to confirm that there is no invalid mirror disk in the boot disk group.</p>
1003391	<p>The system panics at volpage_pageread_done() when recovering a large volume.</p> <p>The panic is resolved by freeing the SIO after the done function is called.</p>
1274174	<p>A system panics in function vxtask_created() with a data page fault when mirroring a VxVM volume.</p> <p>The variable is initialized properly.</p>
1274257	<p>The system panics during a failover when the vxmapadm disable path/ ctrlr command is executed on a LUN, which has non-A/A arrays connected. The following is the panic string "Panic string: fdmp_select_path:2a".</p> <p>If the vxmapadm disable path/ctrlr command is executed while failover is in progress, the failover operation is restarted to make sure that the I/O is not routed to the disabled path.</p>
1299393	<p>During a disk rescan operation, the panic occurs in voldios_start() function.</p> <p>During an I/O operation, the reference counts on the disk structure are incremented. The disk structure is not deleted because of a non-zero reference count and the panic is prevented.</p>

Incident	Description
1101741	<p>All vx commands are hung when an A/A-A array is undergoing a path failover while vxconfigd is restarting. A vxconfigd thread is seen sleeping indefinitely in function gendmpopen() with a call to sleep_spinunlock().</p> <p>During path failover, instead of the DMP device, the underlying paths are opened to read PGR keys so the second call to gendmpopen() is avoided.</p>
1274260	<p>While setting the path attribute nonpreferred, the vxcdmpadm(1M) command stops responding.</p> <p>When the nonpreferred attribute change is requested, while I/O operations are executing, the corresponding paths are marked correctly. As a result, the I/O operations complete within the specified time.</p>
1314312	<p>In a CVM cluster environment, the abort processing stops responding when the master node is aborted. As a result, all the VxVM commands also stop responding.</p> <p>The "CVM_MSG_DCO_INST_DISTRIBUTE" message purge performs the corresponding cleanup operation.</p>
1171737	<p>System panics in vol_kmsg_get_packed_sz() routine when host bus adapter (HBA) cables from the array are removed.</p> <p>The issue has been resolved to ensure that the proper message allocations are done.</p>
1197127	<p>In a CVM cluster, a node encounters data TLB fault panic while a plex is dissociated from a shared mirror volume. The panic occurs in function volmv_msg_detach()</p> <p>The code has been changed to make sure that the disk group pointer gets reinitialized.</p>
1171796	<p>The CVM master node crashes in function gen_shared_update_cur_pri() when HBA cables are removed from the disk array with I/Os continuously going on the shared mirrored volume.</p> <p>The issue has been fixed so that the NODE QUIESCED flag is reset with dmp exclusive spin lock (vxdmp_spin) held.</p>

Incident	Description
1154217	<p>In CVM environment, a shared volume is mirrored across two A/A-A arrays or enclosures. The dd(1) command issues I/Os on the shared volumes. While I/Os are in progress, if one of the arrays is turned off or the cable is pulled out, the master node panics in vxdmp_hp_kmem_zalloc () function.</p> <p>The solution in case of a non-A/A array is to not allow I/Os when the <code>cur_pri_path</code> of dmp node is set to NULL.</p>
1130126	<p>Under very heavy I/O load on mirrored volumes, a read request may return success to user space even though the I/O has not been performed. This can result in data corruption.</p> <p>A timestamp is properly initialized on the DMP buffer so that the I/O request is queued for retrial if it times out.</p>
1274148	<p>Following the first reboot after an ignite installation with a VxVM root disk, MSA LUNs are discovered to have duplicate disk IDs.</p> <p>The solution is to create the <code>/etc/vx/.aasci3</code> file from the post install script, which is run following the swinstall of the VRTSvxvm fileset.</p>
1274185	<p>Initializing a VxVM boot disk for IA platforms using <code>vxdisksetup(1M)</code> may fail with the following errors:</p> <pre data-bbox="721 964 1296 1090">vxfs mkfs: ERROR: V-3-25495: read failure at devid/blknum 0/2400191 vxfs mkfs: ERROR: V-3-21080: vxfs: No such device or address</pre> <p>The solution is to add the code to append the s2 partition to the <code>vxdisk list cXtYdZs2</code> output.</p>
1274272	<p>In a CVM environment, if <code>vxdmpadm</code> is used to disable and then re-enable a DMP path to a non-A/A LUN which belongs to a shared disk group, subsequent I/O sent to the DMP path will never return. When this problem scenario happens to all DMP paths of the LUN, I/O to the dmpnode will hang altogether.</p> <p>DMP resets the flag while enabling the paths for LUNs belonging to non-A/A category.</p>

Incident	Description
1274276	<p>In a CVM environment, running <code>vxdmpadm [disable enable] <path/ctlr></code> when a path failover happens in parallel may result in a system panic in <code>dmp_dev_strategy()</code>. Disabling/enabling a DMP path multiple times is known to be a typical scenario for this panic as the disable/enable operation triggers the path failover.</p> <p>A reference count has been added to keep track of outstanding I/Os to a disabled path during path quiesce. The disabled path is closed/unopened only when there is no more outstanding I/Os, thus avoiding the race.</p>
1299403	<p>In a CVM environment, after a node fails and the CVM master completes node recovery, layered volumes are left in a SYNC state permanently. The plex level recovery does not take place to move the volumes to ACTIVE state. The layered volumes remaining in SYNC state can cause huge performance overhead. This recovery issue is specific to layered volumes in a shared disk group.</p> <p>The <code>vxrecover</code> has been changed to specify the <code>-f</code> option for the resync operation if the DRL has sub-volumes. This enables the recovery of the top-level volume in a layered volume hierarchy.</p>
1362201	<p>In a Cluster Volume Manager (CVM) environment, master hangs in 'master selection' state during random shutdown of nodes.</p> <p>CVM reconfiguration code is modified to allow the transfer of the master selection message even when the kernel messaging is disabled.</p>
1362204	<p>In a Cluster Volume Manager (CVM) environment, a Data Page fault panic was observed during CVM reconfiguration.</p> <ul style="list-style-type: none"> ■ The code in the sender thread, which accesses the message after sending it, is changed to access the message before sending it. ■ Code that handles errors while sending a message still needs to access the message after it is sent. This code verifies the existence of the message using the message ID before accessing the entire message.

Incident	Description
1415504	<p>In a campus cluster environment, a full resynchronization operation occurs on a detached plex when it reattaches to the cluster. This leads to a performance degradation.</p> <p>The code has been modified such that a fast resynchronization occurs instead of a full resynchronization.</p>
1416347	<p>Executing the vxctl(1M) enable or vxdisk(1M) scandisks command results in a system panic in function vol_change_disk() during a new storage allocation in a cluster environment.</p> <p>A check has been introduced to detect the NULL pointer before dereferencing it. When the NULL point is detected, check has been made not to access disk pointer values and return VKE_ENOENT error code.</p>
1416349	<p>In a Cluster Volume Manager (CVM) environment, Data Page fault panic was observed in function vol_putdisk() during LUN provisioning.</p> <p>Voldisk deletion code path is modified to acquire appropriate locks before removing the voldisk structure from the list.</p>
1427506	<p>In the CVM environment, with the heavy load of I/Os on the nodes, I/Os on all the nodes hang followed by uncorrectable write errors:</p> <p>The following message is seen in syslog:(/var/adm/syslog/syslog:</p> <pre>Nov 10 18:48:53 eig3 vmunix: WARNING: VxVM vxio V-5-3-0 vol_klog_flush_childdone : I/O Timedout on disk <diskname>, 1 I/Os hung on the disk!</pre> <pre>Nov 10 18:48:53 eig3 vmunix: WARNING: VxVM vxio V-5-3-0 vol_klog_flush_childdone: I/O hung; disallowing all I/Os to the disk. : I/O hung; disallowing all I/Os to the disk...</pre> <p>The default I/O timeout value is increased to accommodate scenarios when an I/O takes more time than the I/O time out value.</p>
1450048	<p>DMP stats daemon wakes up every second for statistics monitoring and leads to a system panic in function dmp_get_iocount()</p> <p>Dynamic Multipathing (DMP) code is fixed by adding locks to synchronize threads for read and write access to 'dmp_stats_table', thus avoiding the race condition.</p>

Incident	Description
1453894	<p>Data TLB fault panic occurs in function scsi_strategy_real() after extra paths are added to an existing LUN on the system.</p> <p>The code has been fixed to perform the second open (in read/write mode) on the new paths even if the existing dmpnode has been opened in read/write mode.</p>
1459367	<p>Enhance the vxdisk(1m) offline/online operation to offline/online multiple disks through a single command.</p> <p>vxdisk(1m) is modified with the following changes:</p> <ul style="list-style-type: none"> ■ The new filename option “-l filename” is introduced in vxdisk. ■ The vxdisk command sends a list of disk access names to vxconfigd in a single request for online/offline operations. <p>vxconfigd(1m) is modified with the following changes:</p> <ul style="list-style-type: none"> ■ vxconfigd can receive a list of disk access names for online/offline requests. ■ vxconfigd uses multiple threads to process the online/offline requests in parallel.
1158584	<p>Mirroring for rootdisk system volumes fail even with a Base-VxVM license.</p> <p>The solution is to rectify the check that identifies system volumes on the root disk and permit mirroring of rootdisk volumes with a Base-VxVM license.</p>
901010	<p>After adding a new path to a CLARiiON array online, vxdisk scandisks command may take more than an hour to complete if the CLARiiON array is in the Active/Passive Explicit Failover (A/PF) mode.</p> <p>DDL is modified to do a SCSI pass-through open instead of a regular open for discovery before issuing SCSI inquiry.</p>
1019507	<p>The vxdisk list command shows the capacity of the disk incorrectly.</p> <p>The correct disk geometry is calculated after comparing both the mode sense data and data obtained from the explicit DIOC_IOCTL ioctl(2).</p>

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 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 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 software 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
1425413, 1425405	The <code>qiomkfile</code> file had incorrect permissions set. Permissions on the <code>qiomkfile</code> are corrected.
1510791	After an <code>open(..)</code> failed with ENOENT because the parent directory was removed, subsequent <code>open(..)</code> still failed with ENOENT even if the parent directory was re-created. A negative DNLC entry is not added if the directory inode is marked for removal.
1465699	A Data page fault panic occurred in the <code>vx_statvfs()</code> function. The <code>vx_statvfs</code> code is modified to lock and validate the VFS structures.

Incident	Description
1507758	An oversized transaction processing goes into an infinite loop and leads to a system hang with a vx_trancommit() function failure. Reservation sizes are increased only gradually after vx_trancommit() failures.
1274317	The intent log replay during a full fsck operation took approximately 2 to 3 hours to execute. Tail pointers are added to the linked lists. As a result, the fsck utility takes less time to execute.
1417404	The vxupgrade command allowed / and /stand filesystems to be upgraded to disk layout Version greater than 5. A check is added to prevent the vxupgrade command from upgrading / and /stand filesystems to disk layout Version greater than 5.

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
1093618	When running DBED commands with Oracle Release 11gR1, they exit with errors. The following error message is displayed: ERROR V-81-3221 Unsupported Oracle Release. The DBED scripts have been enhanced to support Oracle 11gR1.
1488232	The dbed_ckptcreate (1M) script displays the following error: SFORA vxreptadm ERROR V-81-6509 Invalid input format. Expect 4 columns, got 0 columns. SFORA vxreptadm ERROR V-81-6510 Expected input format: <ThreadNumber> <ArchiveSequence> <ArchiveFormat> <ArchiveDest> The script has been fixed to retrieve the correct redolog sequence number when some Oracle instances are offline.
1488217	The DBDST (Database Dynamic Storage Tiering) commands such as dbdst_file_move and dbdst_convert cannot be run from a CVM (Veritas Cluster Volume Manager) slave node. The DBDST commands have been modified to ship the VxVM commands to the CVM master node through the vxdbd daemon. The command can then be run on the master node and the results returned to the slave node.

Incident	Description
1488208	<p>The dbed_vmclonedb (1M) and dbed_clonedb (1M) scripts fail with error messages when run with Oracle Release 10gR1.</p> <p>SFORA dbed_vmclonedb ERROR V-81-4882 An error occurred while reconfiguring Oracle instance 'ORAINST'.</p> <p>SFORA dbed_vmclonedb ERROR V-81-4881 Log file is at /tmp/dbed_vmclonedb.29130/startup.log.</p> <p>The code is modified to correctly handle the parameters by checking the Oracle Release version number.</p>
1488212	<p>The DBED scripts, dbed_clonedb (1M) and dbed_vmclonedb (1M) fail for Oracle Flashback databases.</p> <p>The DBED commands fail for Oracle Flashback databases.</p> <p>The following error messages are displayed:</p> <pre>ORA-38774: cannot disable media recovery - flashback database is enabled ALTER DATABASE OPEN * ERROR at line 1: ORA-38760: This database instance failed to turn on flashback database SFORA dbed_clonedb WARNING V-81-4921 Database clone4 created output warnings during restart.</pre> <p>This issue was fixed by modifying the code to turn off the flashback database mode before switching to the NOARCHIVELOG mode.</p>
1488204	<p>The DBED commands such as dbed_clonedb(1M), dbed_vmclonedb(1M) fail for databases that have more than 1000 datafiles or tablespaces.</p> <p>Error messages similar to the following are displayed:</p> <pre>SFORA dbed_clonedb ERROR V-81-4882 An error occurred while reconfiguring Oracle instance 'nashoff'.</pre> <p>SFORA dbed_clonedb ERROR V-81-4881 Log file is at /tmp/oralog.out.21766.</p> <p>This DBED code has been modified to remove checks for any upper limits for datafiles and tablespaces. The limits are retrieved from the MAXDATAFILES and MAXTABLESPACES attributes from the Oracle database.</p>

Incident	Description
1488223	<p>On HP-UX IPF systems, the <code>odm.check</code> script reports that the Oracle ODM library is not linked to the Veritas ODM library even if it is correctly linked.</p> <p>The <code>odm.check</code> script has been modified to correctly check for the link to the <code>libodm.so</code> library.</p>
1488234	<p>After rebooting all the nodes in a VCS cluster, the DBED repository service group state is not automatically onlined. The VCS command <code>hastatus -sum</code> lists the repository service group state as “PARTIAL.”</p> <p>The DBED code has been changed to not run the database repository <code>init</code> scripts in a VCS environment to avoid the conflict.</p>

Known issues

Veritas Storage Foundation known issues in the SF 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*. You can view them at the following URL:

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

Veritas Storage Foundation 5.0 MP2 known issues

Veritas Storage Foundation known issues

During the install of SF from DVD media, the installer may display an install failed error messages for VRTSvxvm package [1527735]

When you run the install of SFCFS and SF from DVD media, the installer may display an install failed error messages for VRTSvxvm package. The installer may be intervened by a timeout (600 seconds) during the VRTSvxvm package installation from a DVD media depending on the speed of DVD driver.

Workaround: Use installer -timeout 1800. The timeout default value (600 seconds) will be overridden to 1800 seconds.

Task Assistant tab is not available in the VEA GUI [1528943]

After upgrading from SF 4.1 to SF 5.0 the Task Assistant tab is not available in VEA GUI. Use the System tab to perform operations using VEA GUI.

vxinstalltemplate binary of VRTSalloc core dumps [1526254]

The vxinstalltemplate binary of VRTSalloc core dumps during 5.0 MP2 SF stack installation. A core file with name core.vxinstalltemplate is generated.

Workaround: To resolve this issue, run the following command:

```
# /usr/sbin/vxtemplate -r -C -d \
/optVRTSalloc/config/alloc_capabilities.txt install
```

Information about a managed host is not updated in Central Server GUI [1519528]

After you upgrade from SF or SFORA 4.1MP2 to 5.0 MP2, information about a managed host is not updated in the Central Server GUI.

Workaround: You can use the following steps to push the latest patch on the MH after which the host information is discovered correctly.

- 1 Remove the UNOF patch of VRTSmh on the managed host. For example:

```
# swremove UNOF_36026
```
- 2 From the central server, push the latest available patch to the managed host using vxdompackage as given below. For example:

```
# /opt/VRTSscs/adm/vxdompackage add --hostfile  
/hostfile --imagefile \  
/opt/VRTSscs/adm/vxdom_images/hpux/VRTSMH_1.1.tar  
where hostfile is the file containing the managed host name.
```

Warning messages displayed during verification of VRTS packages

Warning messages are displayed for certain VRTS packages when you run the swverify command to verify the installation. These warnings may be ignored.

If you are upgrading from SF 5.0 MP1, the following warning message is displayed for the following VRTS package:

VRTSrvw [1454379]

The following warning message is displayed:

```
WARNING: Fileset "VRTSrvw.VRW-FILESET, l=/,r=5.0" had file warnings.
```

If you are upgrading from SF 5.0, the following warning message is displayed for the VRTS package:

VRTSacclib [1454382]

The following warning message is displayed:

```
WARNING: Fileset "VRTSacclib.library,l=/,r=5.0.00.0" had file  
warnings.
```

Warning messages may appear in the swinstall.log and swremove.log files (1529476)

Warning messages may appear in the swinstall.log files after you have installed the product using the common product installer script. For example:

```
WARNING: The dependencies for fileset  
PHCO_38836.VXFS-ENG-A-MAN,r=1.0"cannot be resolved (see  
previous lines).
```

The operation on this fileset will still be attempted even though there are unresolved dependencies because the "enforce_dependencies" option is set to "false".

```
WARNING: New Install PHCO_38836.VXFS-ENG-A-MAN,r=1.0
```

```
WARNING: 1 of 1 filesets had Warnings.
```

```
WARNING: The Analysis Phase had warnings. See the above  
output for details.
```

```
WARNING: The dependencies for fileset  
"PHCO_38850.VXFS-RUN,r=1.0"cannot be resolved (see  
previous lines). The operation on this fileset will still  
be attempted even though there are unresolved
```

dependencies because the "enforce_dependencies" option is set to "false".

* Summary of Analysis Phase:

WARNING: New Install PHCO_38850.VXFS-RUN,r=1.0

WARNING: 1 of 3 filesets had Warnings.

* 2 of 3 filesets had no Errors or Warnings.

WARNING: The Analysis Phase had warnings. See the above output for details.

WARNING: The dependencies for fileset

"PHKL_38829.VXVM-KRN,r=1.0" cannot be resolved (see previous lines). The operation on this fileset will still be attempted even though there are unresolved dependencies because the "enforce_dependencies" option is set to "false".

* Summary of Analysis Phase:

WARNING: New Install PHKL_38829.VXVM-KRN,r=1.0

WARNING: 1 of 1 filesets had Warnings.

WARNING: The Analysis Phase had warnings. See the above output for details.

WARNING: The dependencies for fileset

"PHCO_38830.VXVM-RUN,r=1.0" cannot be resolved (see previous lines). The operation on this fileset will still be attempted even though there are unresolved dependencies because the "enforce_dependencies" option is set to "false".

* Summary of Analysis Phase:

WARNING: New Install PHCO_38830.VXVM-RUN,r=1.0

WARNING: 1 of 4 filesets had Warnings.

* 3 of 4 filesets had no Errors or Warnings.

WARNING: The Analysis Phase had warnings. See the above output for details.

Warning messages may appear in the swremove.log file while removing the patches

```
* The fileset
"PHCO_38831.VMPRO-PRG,l=/opt/VRTSvmp, r=1.0" requires
the selected fileset "PHKL_38829.VXVM-KRN,l=/,r=1.0"
as a corequisite.

WARNING: The dependencies for fileset
"PHKL_38829.VXVM-KRN,l=/,r=1.0" cannot be resolved (see
previous lines). The operation on this fileset will still
be attempted even though there are unresolved
dependencies because the "enforce_dependencies" option is
set to "false".
* The fileset
"PHCO_38831.VMPRO-PRG,l=/opt/VRTSvmp, r=1.0" Storage
Foundation Cluster File System Release Notes 71
Known issues
requires the selected fileset
"PHCO_38830.VXVM-RUN,l=/,r=1.0"
as a corequisite.

WARNING: The dependencies for fileset
"PHCO_38830.VXVM-RUN,l=/,r=1.0"
cannot be resolved (see previous lines).
The operation on this fileset will still be attempted
even
though there are unresolved dependencies because the
"enforce_dependencies" option is set to "false".
* Summary of Analysis Phase:

WARNING: Remove PHKL_38829.VXVM-KRN,l=/,r=1.0
WARNING: Remove PHCO_38830.VXVM-RUN,l=/,r=1.0
WARNING: 2 of 46 filesets had Warnings.
* 44 of 46 filesets had no Errors or Warnings.
WARNING: The Analysis Phase had warnings. See the above
output for details.
These warning messages are harmless and may be ignored. The required
patches are correctly installed despite these warning messages.
```

Veritas Storage Foundation for Oracle known issues

Warning messages may appear in the swinstall.log [1540372]

Warning messages may appear in the swinstall.log file after you have installed the product using the common product installer script. For example:

```
WARNING: The dependencies for fileset "PHKL_38795.ODM-KRN,r=1.0"
cannot be resolved (see previous lines).
The operation on this fileset will still be attempted even though
there are unresolved dependencies because the "enforce_dependencies"
option is set to "false". * The software
"PHKL_38794.VXFS-KRN,l=/,r=1.0" is not in the correct state
(AVAILABLE, INSTALLED, or CONFIGURED). You will need to recopy,
reinstall or configure this software before operating on software
that depends on it. The corequisite
"PHKL_38794.VXFS-KRN,fa=HP-UX_B.11.23_PA" for fileset
"PHKL_38795.ODM-RUN,r=1.0" cannot be successfully resolved.
```

```
WARNING: The dependencies for fileset "PHKL_38795.ODM-RUN,r=1.0"
cannot be resolved (see previous lines). The operation on this
fileset will still be attempted even though there are unresolved
dependencies because the "enforce_dependencies" option is set to
"false".
```

Summary of Analysis Phase:

```
WARNING: New Install PHKL_38795.ODM-KRN,r=1.0
WARNING: New Install PHKL_38795.ODM-RUN,r=1.0
WARNING: 2 of 3 filesets had Warnings.
* 1 of 3 filesets had no Errors or Warnings.
WARNING: The Analysis Phase had warnings. See the above output for
details.
```

These warning messages are harmless and may be ignored. The required patches are correctly installed despite these warning messages.

Offline checkpoint clonedb operation for 10gr1 fails for 1000 or more tablespaces [1526006]

The checkpoint dbd_clonedb fails if there are 1000 or more tablespaces in the primary Oracle Database. This issue is specific to Oracle 10gr1 only. The following error message is displayed:

```
SFORA dbd_clonedb ERROR V-81-4919 Database ncln7 has reported an
error.
SFORA dbd_clonedb ERROR V-81-4881 Log file is at
/tmp/oralog.out.17868.
```

Aries error message is displayed due to the HP-UX Aries Emulator [1484070]

Aries error are generated on HP-UX IA architecture by vxstorage_stats and dbd_analyzer. The following error message is displayed when vxstorage_stats and dbd_analyzer is executed:

```
[HP ARIES32]: Internal Error (code : 021) [PID : <Process PID number]
[HP ARIES32]: Got masked synchronous signal 11.
[HP ARIES32]: Aborting process!!
```

Warning messages displayed during verification of VRTS packages

Warning messages are displayed for certain VRTS packages when you run the swverify command to verify the installation. These warnings may be ignored.

If you are upgrading from SF Oracle 5.0 MP1, the following warning messages are displayed for the VRTS package.

VRTSvrw [1454379]

The following warning message is displayed:

```
WARNING: Fileset "VRTSvrw.VRW-FILESET, l=/,r=5.0" had file warnings.
```

If you are upgrading from SF Oracle 5.0, the warning messages are displayed for the following VRTS packages.

VRTSdbcom [1469234]

The following warning messages are displayed:

```
WARNING: Directory "/opt/VRTS/man/man1m" should have mode "755" but the actual mode is "555".
WARNING: Directory "/opt/VRTS/man/man1m" should have group,gid "sys,3" but the actual group,gid is "other,1".
WARNING: Fileset "VRTSdbcom.DBCOM,l=/,r=5.0.00.00-GA" had file warnings.
```

VRTSdbed [1469234]

The following warning messages are displayed:

```
WARNING: Directory "/etc/vx" should have mode "755" but the actual mode is "555".
WARNING: Directory "/etc/vx" should have owner,uid "root,0" but the actual owner,uid is "bin,2".
WARNING: Directory "/etc/vx" should have group,gid "sys,3" but the actual group,gid is "bin,2".
WARNING: Directory "/opt/VRTS/man/man4" should have mode "755" but the actual mode is "555".
WARNING: Directory "/opt/VRTS/man/man4" should have group,gid "sys,3" but the actual group,gid is "other,1".
WARNING: Directory "/opt/VRTSdbed" should have mode "700" but the actual mode is "750".
```

```
WARNING: Directory "/opt/VRTSdbed" should have owner,uid "root,0" but the actual owner,uid is "oracle,107".
WARNING: Directory "/opt/VRTSdbed" should have group,gid "sys,3" but the actual group,gid is "oinstall,1001".
WARNING: Fileset "VRTSdbed.DBED,l=/,r=5.0.00.00-GA" had file warnings.
```

VRTSacclib [1454382]

The following warning message is displayed:

```
WARNING: Fileset "VRTSacclib.library,l=/,r=5.0.00.0" had file warnings.
```

VRTSodm [1501804]

The following warning message is displayed:

```
WARNING: Fileset "VRTSodm.ODM-KRN,l=/,r=5.0.01.00" had file warnings.
```

VRTSdbdoc [1513034]

The following warning message is displayed:

```
WARNING: Directory "/opt/VRTS" should have mode "755" but the actual mode is "555".
WARNING: Directory "/opt/VRTS" should have owner,uid "root,0" but the actual owner,uid is "bin,2".
WARNING: Directory "/opt/VRTS" should have group,gid "sys,3" but the actual group,gid is "bin,2".
WARNING: Directory "/opt/VRTS/docs" should have mode "755" but the actual mode is "555".
WARNING: Directory "/opt/VRTS/docs" should have owner,uid "root,0" but the actual owner,uid is "bin,2".
WARNING: Directory "/opt/VRTS/docs" should have group,gid "sys,3" but the actual group,gid is "bin,2".
WARNING: Fileset "VRTSdbdoc.DBDOC,l=/,r=5.0.0.0" had file warnings.
```

Cannot unmount single-host clone in HA environment after failover [818522]

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

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

Veritas Volume Manager known issues

Harmless Warnings

Normal boot time warnings for Veritas Volume Manager [1514814]

At normal boot time, warning messages are displayed. The following warning messages are harmless and do not indicate any problem.

```
WARNING: "NOTICE: VxVM vxio V-5-0-1046 changing UDID for disk
<diskname>"  
WARNING: "VxVM vxconfigd ERROR V-5-1-12826 /etc/vx/uuid/bin/osuuid
list: invalid guid:"
```

DMP Issues

vxesd Compatibility Issues [1499955]

In certain configurations, vxesd dumps core and does not work properly. This happens due to the compatibility issues between the vxesd component of DMP and the HP-UX libraries. vxesd is an optional component of DMP and does not affect any DMP functionality.

Veritas Storage Foundation 5.0 MP1 known issues

Veritas Storage Foundation known issues

Software disc cannot be ejected during installation [622442]

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, vxreloacd, 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/vxreloacd
```

Uninstalling the VRTSmapro package [617740]

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:

```
# /optVRTSmapro/bin/vxmapping_prov.config -r
```

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 must be 022. Symantec recommends setting it 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, you are prompted to enter the fully-qualified host name manually to continue the installation.

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

Unconfigure VEA Action Agent after a MANAGED host installation [616057]

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.

After upgrading Volume Replicator, replications may stop responding [308183]

Replication may stop responding after upgrading to Volume Replicator 5.0, 5.0 MP1, or 5.0 MP2, if the disk group version is not upgraded at the same time.

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 known issues in this MP2 release of Veritas Storage Foundation for Oracle.

File fragmentation check in the qio_convertdbfiles command may report errors [819430]

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

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

Problems uninstalling or upgrading Veritas Storage Foundation for Oracle when Veritas Storage Foundation Cluster File System is installed on the same system [840486]

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

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

To uninstall the Veritas Storage Foundation for Oracle packages

- 1 Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
- 2 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:

```
# /sbin/init.d/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.

dbed_vmclosedb -p failed to create clonedb with modified pfile [853792]

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

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

```
SFORA dbed_vmclosedb 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 processes a pfile correctly at startup even if a parameter value contains unquoted special characters. However, the pfile parser uses by Veritas Storage Foundation for Oracle strictly enforces the pfile specification contained in the Oracle Reference Manual.

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

Cannot use Web GUI to view snapplan log [861696]

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.

Database FlashSnap archive log destinations [862092, 862687]

With Oracle Release 10g and above, Database FlashSnap clones do not support DB_RECOVERY_FILE_DESTINATION as the sole mandatory archive log destination. This issue will not be detected by FlashSnap validation with dbed_vmchecksnap, or by the snapshot command dbed_vmsnap. However, recovery fails when attempting to clone a database using the snapshot, and the message ORA-01195: online backup of file 1 needs more recovery to be consistent may appear in the log file.

Workaround: Define a mandatory log archive destination that is not DB_RECOVERY_FILE_DESTINATION and set the ARCHIVELOG_DEST parameter of the snapplan to this value.

Veritas 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

- Verify 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.

- 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.

- Update the information that VxVM holds about the device.

```
# vxdisk scandisks
```

- Check that the cloned disk is 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
```

- Import the cloned disk into the new disk group newdg, and update the disk's UDID.

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

- Check that the state of the cloned disk is 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 [784343]

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.

Handling intermittently failing paths

The `dmp_health_time` and `dmp_path_age` tunables control how DMP handles intermittently failing paths. The default values in VxVM 5.0 MP2 of `dmp_health_time` and `dmp_path_age` are 60 and 300 seconds respectively. The value of `dmp_health_time` represents the minimum time in seconds for which a path must stay healthy. If a path changes state between enabled and disabled on a shorter time scale than this, DMP marks the path as intermittently failing and disables I/O on the path. I/O is not re-enabled on an intermittently failing path until `dmp_path_age` seconds have elapsed without further outage. The minimum configurable value of `dmp_path_age` is 0, which prevents DMP from detecting intermittently failing paths.

Cluster issues

Handling intermittently failing paths in a Campus Cluster

In remote mirror configurations, a site is reattached when its disks come back online. Recovery is then initiated for the plexes of a volume that are configured at that site. Depending on the configuration, recovery of the plexes can take a considerable time and consume considerable resources. To minimize the frequency of having to perform a site reattachment operation, you should use the `vxsdmpadm 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 automatically reattach 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, `vxreloacd`, is running, `vxsited` tries to reattach the site, and lets `vxreloacd` try to use the available disks in the disk group to relocate the failed subdisks. If `vxreloacd` relocates the failed subdisks, it starts to recover the plexes at the site. When it recovers all the plexes, they are put into the ACTIVE state, and the state of the site is set to ACTIVE.

If `vxrelocl` is not running, `vxsited` reattaches a site only when all the disks at that site become accessible. After the site is reattached, `vxsited` sets the site state to ACTIVE, and starts recovering the plexes. When all the plexes have been recovered, they are put into the ACTIVE state.

Note: `vxsited` does not try to reattach a site that you have explicitly detached with the `vxmg 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 start recovery of plexes at those sites. To send mail to other users, add the user name to the line that starts `vxsited` 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 `vxsited` daemon, and prevent it from restarting. To kill the daemon, run the following command from the command line:

```
# ps -afe
```

Find 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 `/sbin/init.d/vxvm-recover` startup script.

Replacing a disk in a site-consistent disk group [536853]

If you use the `vxdiskadm` command to replace a disk in a 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.

Domain controller mode in CVM clusters [603213]

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

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

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

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

Veritas Enterprise Administrator issues

Volume tags not displayed [602953]

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

Search does not return any objects for non-Administrator users [840452]

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.

Workaround: Add the user to the Administrator group.

Veritas Web GUI issues

Incorrect error message when importing a disk group [607096]

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

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

Workaround: Refresh the page.

Error when creating a volume set [615960]

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

Workaround: Refresh the page.

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 MP2 release of the 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 MP2 release of the Veritas Volume Replicator.

Synchronizing volumes and RVG with large volumes [840217]

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 synchronized and it actually synchronizes only a portion of the volume having size larger than 1TB.

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

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

Issue with VVR VEA in the Japanese locale [616709]

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

Bunker DG: If protocol is Storage the Bunker DG is expected to have been imported on the Primary host.

Protocol: Protocol should be set to Storage when Bunker storage is directly accessible from the Primary host.

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)`, `vxdmpadm(1M)`, `vxdmpinq(1M)`, `vxpool(1M)`, `vxresize(1M)`, `vxtemplate(1M)`, and `vxvoladm(1M)` manual pages are updated in this Maintenance Pack to correct 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 on 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 as follows:

```
# 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, enter the following:

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

- 2 List the details of the SRVM configuration file.

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

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

```
-rw-r--r-- 1 oracle dba 10569216 Jan 20 14:29 /db/srvvm.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 svrctl configuration:

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

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

Then add the second instance:

```
$ svrctl add instance -d KPRDADV1 -i KPADV1R2 -n node2
$ svrctl 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.

```
$ svrctl 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 you configure 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 the HA environment by running the following commands in the order presented:
 - Unconfigure the SFUA repository.
`# /opt/VRTSdbcom/bin/sfua_db_config -o unconfig_cluster`
 - Import the repository disk group.
 - Start the repository disk volume.
 - Mount the repository file system.
 - Run the following 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_vmcchecksnap -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 GUI 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
```