Veritas Storage Foundation™ and High Availability Read This First

Linux for IBM™ Power and x86_64

5.0 Release Update 4

symantec™

Veritas Storage Foundation™ Read This First

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 RU4

Document version: 5.0 RU4.2

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://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:

www.symantec.com/techsupp/

Contacting Technical Support

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

www.symantec.com/techsupp/

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:

https://licensing.symantec.com

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/assistance_care.jsp

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	customercare_apac@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:

www.symantec.com

Select your country or language from the site index.

Contents

Chapter 1

Storage Foundation - Overview

This chapter includes the following topics:

Overview of this release

The Veritas Storage Foundation 5.0 RU4 release is based on a mix of previous 5.0 MP3 rolling patches and 5.0 release updates. This 5.0 RU4 release includes support for x64-bit Linux platforms and Linux for IBM Power™ platforms running the following operating systems:

Table 1-1 Operating system support in Veritas Storage Foundation 5.0 RU4

Operating system	Architecture
SUSE Linux Enterprise Server 10 Service Pack 3 (SLES10 SP3)	Linux on IBM™ Power platforms x86_64-bit platforms
SUSE Linux Enterprise Server 11 (SLES11)	Linux on IBM™ Power platforms

This document provides release information about the products in the Veritas Storage Foundation 5.0 Release Update 4 product line:

- Veritas Storage Foundation™ (Basic, Standard, Standard HA, Enterprise, and Enterprise HA)
- Veritas Storage Foundation™ for Oracle (Standard and Enterprise Editions)

- Veritas Storage Foundation™ for Sybase (Standard, Enterprise, and HA Editions)
- Veritas Storage Foundation™ for DB2 (Standard, Enterprise, and HA Editions)
- Veritas™ Volume Replicator (VVR)
- Veritas™ Volume Manager (VxVM)
- Veritas™ File System (VxFS)
- Veritas Storage Foundation™ Cluster File System (SFCFS) (Standard and HA Editions)
- Veritas™ Cluster Server (x32-bit and x64-bit)
- Veritas Storage Foundation™ Cluster File System for Oracle RAC (SFCFS for Oracle RAC)

There is no Veritas Storage Foundation support for Xen or VMWare in the base Linux operating systems for this release. There is support for Xen in this 5.0 RU4 release for SLES 10 SP3 only.

Although VMWare is not supported on Linux on IBM Power™, Veritas Storage Foundation is supported in the VMWare guest for SLES 10 SP3 (with certain restrictions such as DMP), provided that the VMWare guest supports the Linux SUSE operating systems available for this 5.0 RU4 release.

About Release 5.0 RU4 Installation options

The Veritas Storage Foundation 5.0 RU4 release includes requirements for both full, fresh installations and in-place upgrade installations:

- For platforms running SLES 10 (SP3) on an x86_64-bit architecture, full installations using the product installer are supported as well as upgrades using the 5.0 RU4 `installmp` script
- For platforms running SLES 10 (SP3) or SLES 11 GA on Linux for IBM Power™, only full installations using the product installer are supported in the 5.0 RU4 release

Each of the products in the Veritas Storage Foundation 5.0 RU4 release is activated by a single license key. You must obtain a license key before installing any products.

Refer to the installation guide for your product:

- *Veritas Storage Foundation Installation Guide*
- *Veritas Cluster Server Installation Guide*
- *Veritas Storage Foundation Cluster File System Installation Guide*

- *Veritas Storage Foundation Cluster File System for Oracle RAC Installation Guide*

See the *Veritas Storage Foundation Installation Guide* for license key information.

About the Simple Admin utility

Veritas Storage Foundation has an optional utility, called Simple Admin, that you can use with Veritas File System and Veritas Volume Manager. The Simple Admin utility simplifies storage management by providing a single interface to the administrator and by abstracting the administrator from many of the commands needed to create and manage volumes, disks groups, and file systems.

You can download the Simple Admin utility for Veritas Storage Foundation from the following URL:

http://www.symantec.com/business/products/agents_options.jsp?pcid=2245&pvid=203_1

About Storage Foundation Manager

Storage Foundation Manager (SFM) is a free license add-on to Veritas Storage Foundation that provides centralized application, server and storage management capabilities across a heterogeneous infrastructure.

You can download SFM packages and SFM Add-ons from the following URL:

http://go.symantec.com/vom

About Release 5.0 RU4 Documentation

Release Update 4 (5.0 RU4) is cumulative and based on Storage Foundation 5.0 MP3 and several 5.0 MP3 release patches (RP1, RP2, RP3) and release updates 5.0 RU1 and RU3 . The documents updated for this 5.0 RU4 release are:

- Veritas Storage Foundation Read This First
- Veritas Storage Foundation Getting Started Guide

The full documentation set for all previous releases is also available online from the Symantec Veritas Storage Foundation website:

http://www.symantec.com/business/support/overview.jsp?pid=15107

For the latest information on updates, patches, and known issues regarding this release, see the Late-Breaking News TechNote on the Symantec Technical Support website:

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

See the Hardware Compatibility List (HCL) for information about hardware support for this 5.0 RU4 release. The hardware TechNote is available at:

http://entsupport.symantec.com/docs/283282

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

This document does not contain release information for Veritas Cluster Server or Veritas Storage Foundation for Oracle RAC.

See the *Veritas Cluster Server Release Notes*.

System requirements

This section describes the system requirements for the 5.0 RU4 release.

Storage Foundation supported Linux operating systems

This section describes the supported operating systems for Storage Foundation 5.0 RU4.

Storage Foundation operates on the following operating systems distributed by SUSE:

- SUSE Linux Enterprise Server 10 Service Pack 3 (SLES10 SP3) for Linux on IBM™ Power platforms and x86_64-bit platforms
- SUSE Linux Enterprise Server 11 (SLES11) GA

The following table lists the supported kernel versions and architecture for each Linux on IBM™ Power and x64-bit operating system version in this 5.0 RU4 release.

Table 1-2 Supported Linux operating system and kernel versions

Operating system	Kernel	Architecture
SUSE Linux Enterprise Server 10 Service Pack 3 (SLES10 SP3)	2.6.16.60-0.54.5	ppc64 x86_64
SUSE Linux Enterprise Server 11 (SLES11) GA	2.6.27.19-5	ppc64

Symantec supports only SUSE distributed kernel binaries. Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported SUSE service packs is available in the following Late-Breaking News TechNote. The TechNote also includes any updates

to the supported operating systems and software. Read this Late-Breaking News TechNote before you install Symantec products.

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

Operating system fresh installs and upgrades for 5.0 RU4

If your system is running an older version of SUSE Linux Enterprise Server, you must upgrade it before attempting to install the Veritas Storage Foundation 5.0 RU4 software.

The 5.0 RU4 release supports only fresh installs of Storage Foundation using the `installer` script for SLES 10 SP3 Linux on IBM™ Power or SLES 11 for Linux on IBM™ Power.

Both fresh installs as well as upgrades are supported for SLES 10 SP3 for x86_64-bit platforms in this release.

Consult the SUSE documentation as well as the Installation section of this guide for more information on upgrading your system.

Storage Foundation Cluster File System supported Linux operating systems

This section describes the Storage Foundation Cluster File System supported Linux operating systems.

Storage Foundation Cluster File System operates on the following Linux on IBM™ Power operating systems and kernel binaries distributed by SUSE:

- SUSE Linux Enterprise Server 10 (SLES 10) with SP3 (2.6.16.60-0.54.5 kernel) on ppc64 and x86_64-bit platforms
- SUSE Linux Enterprise Server 11 (SLES 11) GA (2.6.27.19-5 kernel) on ppc64

Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Storage Foundation supported DB2 versions

DB2 versions 9.5 and 9.7 are supported on the Linux operating systems listed above for this 5.0 RU4 release.

Storage Foundation supported Oracle versions

Oracle versions 10gR2 (10.2) and 11g R1 are supported on SLES 10 (SP3) on x64 and Oracle version 10gR2 (10.2) for Linux on IBM™ Power platforms in this 5.0 RU4 release.

Storage Foundation supported Sybase versions

Sybase ASE version 15 is supported only on the SLES 10 SP3 operating system for PPC in this 5.0 RU4 release.

Storage Foundation Cluster File System Oracle versions

Oracle versions 10g Release 2 and 11g Release 1 are supported for use with Storage Foundation Cluster File System for Oracle RAC.

Storage Foundation supported Veritas Cluster Server versions

Veritas Cluster Server is supported on the Linux operating systems listed above for this 5.0 RU4 release.

Software and hardware requirements

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

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

The hardware compatibility list (HCL) is available at:

http://entsupport.symantec.com/docs/332581

The hardware TechNote is available at:

http://entsupport.symantec.com/docs/283282

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

VxVM licenses

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

Table 1-3 describes the levels of licensing in Veritas Volume Manager and supported features.

Table 1-3 Levels of licensing in Veritas Volume Manager and supported features

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

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

To see the license features that are enabled in VxVM

- Enter the following command:

```
# vxdctl license
```

Cross-Platform Data Sharing licensing

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

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

Component product release notes

In addition to reading these Release Notes, review all component product release notes before installing the product.

Chapter 2

Storage Foundation - Installation

This chapter includes the following topics:

Installation

Storage Foundation 5.0 RU4 Installation

Storage Foundation 5.0 Release Update 4 provides support for:

- SuSE Linux Enterprise Server 10 Service Pack 3 (SLES10 SP3) for x64-bit platforms and Linux on IBM™ Power platforms
- SuSE Linux Enterprise Server 11 (SLES11) GA for Linux on IBM™ Power platforms

Full, fresh installations as well as upgrade installations are supported in this release as follows:

- For platforms running SLES 10 (SP3) on an x86_64-bit architecture, full CPI installations are supported as well as in-place upgrades using the 5.0 RU4 `installmp` script
- For platforms running SLES 10 (SP3) or SLES 11 GA on a ppc architecture, only full installations are supported in the 5.0 RU4 release

The following sections summarize the installation options for this release.

Veritas Installation Assessment Service

The Veritas Installation Assessment Service (VIAS) utility assists you in getting ready for a Veritas Storage Foundation and High Availability Solutions installation. The VIAS utility allows the preinstallation evaluation of a configuration, to validate it prior to starting an installation or upgrade.

https://vias.symantec.com/

Simplified installation and configuration

Installation and configuration procedures have been simplified, based on usability testing.

Support for SuSE Linux Enterprise Server 10

The Storage Foundation software stack has been enhanced in 5.0 RU4 to support SuSE Linux Enterprise Server 10 Service Pack 3 (SLES10 SP3) for x64-bit platforms and Linux on IBM™ Power platforms.

Support for SuSE Linux Enterprise Server 11

The Storage Foundation software stack has been enhanced in 5.0 Release Update 4 to support SuSE Linux Enterprise Server 11 (SLES11) GA for Linux on IBM™ Power platforms.

Installing or upgrading Storage Foundation and High Availability products to release 5.0 RU4

The only supported upgrade paths to Storage Foundation (SF) 5.0 Release Update 4, Storage Foundation Cluster File System (SFCFS) 5.0 Release Update 4, or Veritas Cluster Server 5.0 Release Update 4 (VCS) are from Storage Foundation 5.0MP3, including any release packs that are supported on SLES10, SLES 10 SP1, or SLES 10 SP2. For any other configuration, upgrade is not supported. You must perform a fresh installation.

Table 2-1 shows links to the installation sections.

Table 2-1 Installation scenarios

Product	Installation scenario	Instructions
Storage Foundation Cluster File System	Installing SFCFS for the first time	See "Installing and configuring Storage Foundation Cluster File System" on page 29.

Table 2-1 Installation scenarios *(continued)*

Product	Installation scenario	Instructions
Storage Foundation Storage Foundation for Oracle Storage Foundation for Sybase Storage Foundation for DB2	Installing SF or SF for databasese for the first time.	See "Installing and configuring Storage Foundation or Storage Foundation for databases for the first time using the common product installer" on page 24.
Storage Foundation Cluster File System for Oracle	Installing SFCFS RAC for the first time.	See "Installing Storage Foundation Cluster File System for Oracle RAC" on page 49.
Veritas Cluster Server	Installing VCS for the first time.	See "Installing Veritas Cluster Server" on page 77.

Upgrades are supported for SLES 10 x64-bit platforms in the 5.0 RU4 release only when upgrading Storage Foundation from the same underlying operating system. For example, a customer running Storage Foundation 5.0 MP3 on SLES 10 SP2 can perform an in-place upgrade to 5.0 RU4 using the `installmp` installer script.

Note: Upgrading Storage Foundation on a previously-upgraded operating system is not supported as an in-place upgrade in this release. Customers running Storage Foundation on SLES 10 SP1, for example, cannot upgrade to the 5.0 RU4 release on SLES 11. In this case, customers must uninstall their existing Storage Foundation release, upgrade their operating system, and then perform a fresh installation 5.0 RU4.

Table 2-2 shows the supported upgrade paths for SF, SF for databases, SFCFS, and VCS. Refer to the appropriate section for your product for details about using the `installmp` script. Upgrade the OS to SLES 10 SP3 before using the `installmp` script to install 5.0 RU4.

Table 2-2 Supported upgrade paths to version 5.0RU4

Operating system	Architecture	Product version	Upgrade path
SLES10 SP1 or SLES10 SP2	x64	SF or SFHA 5.0MP3, 5.0MP3RP2, 5.0MP3RP3, or 5.0RU3 SF for Oracle, SF for DB2 or SF for Sybase 5.0MP3, 5.0MP3RP2, 5.0MP3RP3, or 5.0RU3	Upgrade OS to SLES10 SP3. Upgrade Veritas product to 5.0RU4 using the installmp script. See "Upgrading Storage Foundation software from 5.0 MP3 to 5.0 RU4 using the product installer" on page 28.
SLES10 SP1 or SLES10 SP2	x64	SFCFS or SFCFS HA 5.0MP3, 5.0MP3RP2, 5.0MP3RP3, or 5.0RU3	Upgrade OS to SLES10 SP3. Upgrade Veritas product to 5.0RU4 using the installmp script. See "Upgrading Storage Foundation Cluster File System from release 5.0MP3 or later" on page 37.
SLES10 SP1 or SLES10 SP2	x64	VCS 5.0MP3, 5.0MP3RP2, 5.0MP3RP3, or 5.0RU3	Upgrade OS to SLES10 SP3. Upgrade Veritas product to 5.0RU4 using the installmp script. See "Upgrading Veritas Cluster Server" on page 77.

For upgrade information about SFCFS for Oracle RAC, including upgrade paths, refer to the appropriate section.

See "Upgrading Storage Foundation Cluster File System for Oracle RAC" on page 71.

About the common product installer

The product installer is the recommended method to license and install the Veritas products. The installer also enables you to 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.

See "About installation scripts" on page 23.

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.

About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.0 Release Update 4 provides several installation scripts.

To install the Veritas Storage Foundation products 5.0RU4 on a system that already has Veritas Storage Foundation 5.0MP3, including maintenance packs and rolling patches, use the `installmp` script.

To install a fresh installation on a system, or to upgrade from Veritas Storage Foundation and High Availability Solutions version prior to 5.0 Release Update 4, the recommended installation method is to use the common product installer. To use the common product installer, run the `installer` command.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from an electronic download site, which does not include the common product installer, use the appropriate product installation script.

The following product installation scripts are available:

Veritas Cluster Server (VCS)	`installvcs`
Veritas Volume Replicator (VVR)	`installvvr`
Veritas Storage Foundation (SF)	`installsf`
Veritas Storage Foundation for Oracle (SFORA)	`installsfora`
Veritas Storage Foundation for DB2 (SFDB2)	`installsfdb2`

Veritas Storage Foundation for Sybase (SFSYB) **Note:** In this release, Veritas Storage Foundation for Sybase is supported only for SLES 10 SP3 operating system for PPC.	`installsfsyb`
Veritas Storage Foundation Cluster File System (SFCFS)	`installsfcfs`
Symantec Product Authentication Service (AT)	`installat`
Veritas Volume Manager	`installvm`

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

Installing and configuring Storage Foundation or Storage Foundation for databases for the first time 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 the installation of Storage Foundation on a single system.

Veritas Storage Foundation for Oracle, Veritas Storage Foundation for Sybase, and Veritas Storage Foundation for DB2 can also be installed using this procedure. These products are not available on all platforms.

Note: The common product installer does not automatically install VCS Sybase Agent in this RU4 release. To use the VCS Sybase Agent, install the agent RPM manually.

The Veritas 5.0 RU4 release operates on the following operating system and hardware:

- SUSE Linux Enterprise Server 10 (SLES 10) SP3 (2.6.16.60-0.54.5) on Linux for IBM™ Power and x86_64-bit platforms
- SUSE Linux Enterprise Server 11 (SLES 11) GA (2.6.27.19-5 or later) on Linux for IBM™ Power.

For information about performing in-place upgrades of Veritas Storage Foundation,

To install Storage Foundation

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.

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

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.

   Veritas Storage Foundation for Oracle (not available on SLES 11 ppc64) and Veritas Storage Foundation for DB2 can also be installed using this procedure. Select the number corresponding to one of those products, if desired.

   Note: Do not select the "Storage Foundation Cluster File System for Oracle RAC" option unless you have the correct license and setup.

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 then press Return.

   ```
   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 should 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-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 RPMs or all RPMs. Optional RPMs 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 should see output similar to the following:

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

1) Install required Veritas Storage Foundation rpms -
   491 MB required
2) Install all Veritas Storage Foundation rpms -
   625 MB required

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

11 Configure Storage Foundation when prompted.

```
Are you ready to configure SF? [y,n,q] (y) y
```

Note: Symantec recommends that you do not configure the software during installation. Configuration of software should be done after the installation using the Common Product Installer script `-configure` option available from the `/opt` directory. For example: `opt/VRTS/install/installsf -configure`

---

12 You have the option of specifying the default name of a disk group. If you specify a name, it is used for Veritas Volume Manager commands when a disk group is not 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
```

13 If you responded `y`, then enter the information for the default disk group name.

```
Specify a default disk group name for system host1. [?] dg001
```

14 You are prompted to confirm the default disk group.

Note: If `nodg` is displayed, then the host will be configured to have no default disk group.

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

15 Verify the fully qualified hostname of the systems.

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

16 Enabling Veritas Storage Foundation Management Server management simplifies and improves management of complex data center resources, reducing planned and unplanned down time.

To enable centralized management using Storage Foundation Manager, download Veritas Storage Foundation Management Server from:

http://go.symantec.com/vom

Press **Enter** to continue.

See the *Veritas Storage Foundation Manager Installation Guide* for more information.

17 The installation and configuration complete automatically, and the processes are started.

Check the log file, if needed, to confirm the installation and configuration.

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

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

Note: If you choose not to use the Veritas product installer, you will need to edit `allow_unsupported_modules` in `/etc/modprobe.d/unsupported-modules`. Please refer to Novell support document 7002793.

Upgrading Storage Foundation software from 5.0 MP3 to 5.0 RU4 using the product installer

Use the following procedure to upgrade to Storage Foundation 5.0 Release Update 4. This procedure can be used to upgrade on a standalone system, or on the nodes of a cluster.

The 5.0 RU4 release supports only fresh installs of Storage Foundation using the `installer` script for SLES 10 SP3 and SLES 11 for Linux on IBM™ Power.

Both fresh installs as well as in-place upgrades are supported for SLES 10 SP3 for x86_64-bit platforms in this release.

For an upgrade of Storage Foundation Cluster File System, See "Upgrading Storage Foundation Cluster File System from release 5.0MP3 or later" on page 37. for the steps for a full or phased upgrade:

To upgrade Storage Foundation 5.0 Release Update 4

1 Load and mount the disc.

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

3 To upgrade the Storage Foundation software, invoke the `installmp` command using the option that corresponds to your configuration.

To upgrade the local system, enter the following command:

```
# ./installmp
```

Note: If you are upgrading multiple systems in a cluster, select to upgrade the systems simultaneously.

To upgrade more than one system using secure shell (SSH) utilities, enter the following command from one node in the cluster:

```
# ./installmp system_name1 system_name2 ...
```

To upgrade more than one system using remote shell (RSH) utilities, enter the following command from one node in the cluster:

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

4 After the initial system checks have completed successfully, press Enter to start the requirement checks for the upgrade.

5 After the requirement checks have completed successfully, press Enter to begin upgrading Storage Foundation.

6 Reboot each of the nodes on which you upgraded Storage Foundation.

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

8 If you set the value of the vol_vvr_use_host_byte_order tunable to 1, reboot the system.

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

10 If you are upgrading a cluster, restore any VCS configuration files as described in the *Veritas Cluster Server 5.0 Installation Guide* and *Veritas Cluster Server 5.0 Release Notes*.

Installing and configuring Storage Foundation Cluster File System

Use the procedures in this section to install and configure Storage Foundation Cluster File System or Storage Foundation Cluster File System High Availability.

Installing Storage Foundation Cluster File System using the common product installer

The product installer is the recommended method to license and install Storage Foundation Cluster File System.

The following sample procedure is based on the installation of a Veritas Storage Foundation Cluster File System HA cluster with two nodes: "system01" and "system02." If you are installing on standalone systems only, some steps are unnecessary, and these are indicated.

Default responses are enclosed by parentheses. Press Return to accept defaults.

To install the Storage Foundation Cluster File System

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.

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

4. From this directory, type the following command to install if you are 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. From the Installation menu, choose the `I` option for Install and enter the number for Veritas Storage Foundation Cluster File System. Press **Return**.
6. You are prompted to enter one or more system names to install SFCFS.

   ```
   Enter the system names separted by spaces on which to install
   SFCFS: system01 system02
   ```

7. 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.

8 Enter the product license information.

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

```
Enter a SFCFS license key for system01?
```

9 Enter `y` to accept another license key or enter `n` to proceed.

```
Do you want to enter another license key for system01?
[y,n,q] (n) n
```

10 You can choose to install required RPMs or all RPMs.

11 A list includes the items in the selected option. Press **Return** to continue.

12 If you are installing SFCFS for the first time proceed to configuring the product.

See "Configuring Storage Foundation Cluster File System" on page 31.

Configuring Storage Foundation Cluster File System

This section describes configuring Storage Foundation Cluster File System using the Veritas product installer. If you configured Storage Foundation Cluster File System during the installation process, you do not need to perform the procedure in this section.

To configure the product, run the Veritas product installer or the appropriate installation script using the `-configure` option.

To configure Storage Foundation Cluster File System

1 To invoke the common installer, run the `installer` command with the `configure` option, as shown in this example:

```
# ./installer -configure
```

2 You are prompted to enter the system names (in the following example, "system01" and "system02") on which the software is to be installed. Enter the system name or names and then press Return.

```
Enter the system names separated by spaces on which to
configure SFCFS: system01 system02
```

3 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 the ssh configured for password free logins, or configure rsh and use the -rsh option.

4 The procedure checks system licensing, and you can enter additional licenses, if needed.

```
Checking system licensing

SFCFS license registered on system01

Do you want to enter another license key for system01?
[y,n,q] (n) n
```

5 Any running SFCFS processes are stopped. Press **Return** to continue.

6 Starting I/O Fencing in enabled mode requires manual intervention after SFCFS Configuration. I/O Fencing can be configured in disabled mode now and it does not require any manual intervention after SFCFS Configuration.

Determine at this time if you plan to configure I/O Fencing in enabled mode or disabled mode, as well as the number of network interconnects (NICS) required on your systems. If you configure I/O Fencing in enabled mode only a single NIC is required, though at least two is recommended.

```
Will you be configuring I/O Fencing in enabled mode?
[y,n,q,?] (y) n
```

7 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.

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 NICs available on the first system and reports them.

9 Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat link
on host1: [b,?] eth1
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,?] eth2

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`, be sure that the same NICs are available on each system; the installer may not verify this.

Notice that in this example, `eth0` is not selected for use as a private heartbeat NIC because it already in use 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`, or enter `y` to configure.

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. Refer to the *Veritas Cluster Server Installation Guide* for more information on configuring a VxSS Root Broker.

```
Would you like to configure SFCFS to use
Veritas Security Services? [y,n,q] (n) n
```

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

Enter `y` or `n` to configure SMTP notification.

```
Active NIC devices discovered on host1: eth0 Enter the NIC for the
SF Notifier to use on host1: [b,?] (eth0) eth0 Is eth0 to be the
public NIC used by all systems [y,n,q,b,?] (y) y
```

Enter the domain-based hostname of the SMTP server (for 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**

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

Enter `y` or `n` to configure SNMP notification.

```
Active NIC devices discovered on host1: eth0 Enter the NIC
for the SF Notifier to use on host1: [b,?] (eth0) eth0
Is bge0 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
```

14 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
```

15 You are now 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 by running the `vxdctl defaultdg` *`diskgroup`* command on a system.

See the `vxdctl (1M)` manual page and the *Veritas Volume Manager Administrator's Guide* for more information.

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
```

16 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. [?] diskgroup001
```

17 Validate the default disk group information, and press Return.

18 You may be prompted to verify the fully qualified hostname of the systems.

```
Is the fully qualified hostname of system
 "system01" = system01.veritas.com"? [y,n,q] (y)
```

19 Enter `y` to accept the fully qualified domain name.

```
Is the fully qualified hostname of system
 "system02" = system02.veritas.com"? [y,n,q]
```

20 View the information about enabling Veritas Storage Foundation Manager, then press Return to continue.

21 The Veritas Storage Foundation Cluster File System software is verified and configured.

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

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

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

22 After the configuration completes, restart the Storage Agent.

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

23 The installation script prompts for a reboot if there are one or more errors. Reboot the system (or systems) if the install script prompts you to do so.

24 Before installing Oracle binaries (ORACLE_HOME), consider these points:

- Local installations provide a comfort level using traditional installation methods and the possibility of improved protection against a single point of failure.
- CFS installations provide a single Oracle installation to manage, regardless of number of nodes. This scenario offers a necessary reduction in storage requirements and easy addition of nodes.

Select the location based on your high availability requirements. Symantec generally recommends using local installations.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide* for more information on Oracle Disk Manager.

25 Configure the Storage Foundation Cluster File System and Cluster Volume Manager agents as required.

For more information about configuring agents, see the *Storage Foundation Cluster File System Administrator's Guide*.

To use volumes as part of an Replicated Volume Group (RVG), configure the required RVG agents. The CVMVolDg resource does not support managing or monitoring volumes that are part of RVG.

For more information about RVG agents, see the *Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide*.

Upgrading Storage Foundation Cluster File System from release 5.0MP3 or later

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation Cluster File System: phased and full.

See "Phased upgrade for Release Update 4" on page 37.

See "Full upgrade for Release Update 4" on page 45.

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

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

Preparing to upgrade to the Release Upgrade 4

If you are upgrading an installed Veritas Storage Foundation Cluster File System 5.0MP3 or 5.0MP3RPx release , preserve the existing configuration information.

To preserve the existing configuration information, perform the following actions:

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

Phased upgrade for Release Update 4

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time.

Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

Each phase of the phased upgrade should be performed on more than one node of the cluster.

The stages of the phased upgrade procedure include the following steps:

- Freeze service group operations and stop cluster failover operations.
- Select a two or more nodes to upgrade, and leave a group of one or more nodes running.
- Take the selected group of nodes offline and prepare them for the upgrade.
- Upgrade the Veritas Storage Foundation Cluster File System software on the selected group of nodes.
- Take the second group of nodes offline.
- Bring the first group of nodes online.
- Upgrade the second group of nodes.
- Bring the second group of nodes online and restart cluster failover services. The cluster is fully restored.

Performing the phased upgrade for Release Update 4

This section describes how to perform a phased upgrade for Release Update 4

To select the nodes for the upgrade

1. Select one or more nodes to upgrade first.
2. Leave a group of one or more nodes running.

Upgrading the Veritas Storage Foundation Cluster File System software to Release Update 4

To take the selected group of nodes offline and prepare them for the upgrade

1. Log in as superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so that you can execute all product commands.
3. Switch the failover service group from the first sub-cluster (nodes that are being upgraded) to the second sub-cluster using the following command:

   ```
   # hagrp -switch failover_group_name -to second_sub_cluster_node_name
   ```

4. From any node in the cluster, make the cluster configuration writable.

   ```
   # haconf -makerw
   ```

5. Enter the following command to freeze high availability service group operations on each node:

   ```
   # hasys -freeze -persistent node_name
   ```

6 Make the configuration read-only

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

7 Stop cluster operations on each node in the group being upgraded, by entering the following command:

```
# hastop -local
```

8 Stop all VCS components, ODM, GMS, GLM using the following commands:

```
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxglm stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

9 Check if each node's root disk is under VxVM control by running this command:

```
# df -v /
```

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

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

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

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

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

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

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

11 On each node, use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

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

12 On each node in the cluster, unmount all Storage Checkpoints and file systems:

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

13 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

  ```
  # vxrvg -g diskgroup stop rvg_name
  ```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

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

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

14 Check if the VEA service is running:

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

15 If the VEA service is running, stop it:

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

16 If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.

17 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.

18 On each node, stop all VxVM volumes by entering the following command for each disk group:

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

To upgrade the Veritas Storage Foundation Cluster File System and High Availability software

1 If your operating system is not at a supported level, you must upgrade the operating system.

Upgrade the operating system as follows:

- Disable the startup scripts before upgrading the operating system.

```
# insserv -r vcs
# insserv -r vxfen
# insserv -r vxgms
# insserv -r gab
# insserv -r llt
```

- Upgrade the operating system to the supported OS and service pack level for the first sub-cluster nodes. For instructions, see the operating system documentation.

2 Insert the appropriate media disc into your system's DVD-ROM drive.

3 If volume management software is running on your system, the software disc automatically mounts as `/mnt/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

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

4 Change to the top-level directory on the disc:

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

5 To upgrade the Storage Foundation Cluster File System, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:

- To install on the local system only, enter the following command:

```
# ./installmp
```

- To install on more than one system using secure shell (ssh) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ...
```

- To install on more than one system using remote shell (rsh) utilities, enter the following command:

  ```
  # ./installmp node_name1 node_name2 ... -rsh
  ```

6 After the initial system checks are complete, press **Return** to start the requirement checks.

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

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

  The second group of nodes to be upgraded is stopped. The nodes are stopped before the first group of nodes is rebooted.

9 Reboot the system or systems in the first group of nodes.

Upgrading the remaining nodes

This section describes how to upgrade the remaining nodes.

Take the second group of nodes offline.

Bring the first group (with the newly installed patches) online.

Upgrade the second group of nodes.

To take the second group of nodes offline

1 Unfreeze all the VCS service groups using the command:

```
# hagrp -unfreeze <group_name> -persistent
```

2 Offline all SFCFS resources on nodes to be upgrade by running the following commands on one of the cluster nodes.

```
# hagrp -offline service_group -sys system01
# hagrp -offline service_group -sys system02
```

3 Stop cluster operations on each node in the second group being upgraded, by entering the following command:

```
# hastop -local
```

4 Stop all VCS components, ODM, GMS, GLM using the following commands:

```
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxglm stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

To bring the first group of nodes online

- Bring the first group of nodes online.

  See "Bringing the upgraded group of nodes online" on page 43.

To upgrade the second group of nodes

1 To upgrade the second group of nodes, perform the upgrade of the Veritas Storage Foundation Cluster File System and High Availability software on the second group of nodes. Perform the procedure starting from step 9 of the section "To take the selected group of nodes offline and prepare them for the upgrade"

  See "Upgrading the Veritas Storage Foundation Cluster File System software to Release Update 4" on page 38.

2 Then bring the second group of nodes online.

  See "Bringing the upgraded group of nodes online" on page 43.

Bringing the upgraded group of nodes online

Use the following procedure to bring the upgraded group of nodes online.

To bring the upgraded group of nodes online

1 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.

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

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

4 When the first group of cluster nodes come up, no GAB ports are OPEN. That is, the `gabconfig -a` command shows no GAB ports. For example:

```
# gabconfig -a
 GAB Port Memberships
 ===============================================================
```

Perform the following command to form a cluster after the upgraded nodes are rebooted in the first group of the cluster.

```
# gabconfig -xc
```

GAB ports a, b, d and h now appear in the `gabconfig -a` command output.

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

```
# haconf -makerw
```

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

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

7 Make the configuration read-only:

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

8 On the first group of the cluster, bring the VCS service groups online:

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

After you bring the cvm service group ONLINE, then all of the GAB ports v, w and f come ONLINE. Also, all of the CFS mounts service groups come ONLINE automaticaly. Only failover service groups need to be brought ONLINE manually.

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

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

10 If you have stopped any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, restart each RVG:

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

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

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

12 Check if the VEA service was restarted:

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

13 If the VEA service is not running, restart it:

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

Full upgrade for Release Update 4

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

The stages of the full upgrade procedure are:

- Take all nodes in the cluster offline and install the software patches.
- Bring all the nodes (with the newly installed patches) online to restart cluster failover services. The cluster is fully restored.

Performing the full upgrade to Release Update 4

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

To prepare for a full upgrade to Release Update 4.

1 Log in as superuser.

2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.

3 Stop high-availability cluster operations. This command can be executed from any node in the cluster, and stops cluster operations on all the nodes.

```
# hastop -all
```

4 Check if each node's root disk is under VxVM control by running this command:

```
# df -v /
```

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

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

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

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

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

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

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

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

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

7 Unmount all Storage Checkpoints and file systems:

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

8 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

  ```
  # vxrvg -g diskgroup stop rvg_name
  ```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

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

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

9 Check if the VEA service is running:

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

10 If the VEA service is running, stop it:

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

11 If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.

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 VxVM volumes by entering the following command for each disk group:

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

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

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

15 Continue to the procedure to upgrade the Veritas Storage Foundation Cluster File System software.

To upgrade the Veritas Storage Foundation Cluster File System software

1 Insert the appropriate media disc into your system's DVD-ROM drive.

2 If volume management software is running on your system, the software disc automatically mounts as `/mnt/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

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

3 Upgrade your operating system, and patch it to a supported kernel version.

4 Change to the top-level directory on the disc:

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

5 To upgrade the Storage Foundation Cluster File System, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:

- To install on the local system only, enter the following command:

```
# ./installmp
```

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

```
# ./installmp node_name1 node_name2 ...
```

- To install on more than one system using remote shell (RSH) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ... -rsh
```

6 After the initial system checks are complete, press **Return** to start the requirement checks.

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

8 Shut down and reboot the systems.

To bring the upgraded cluster online and restore components

1 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.

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

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

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

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

5 If you have stopped any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, restart each RVG:

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

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

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

7 Check if the VEA service was restarted:

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

8 If the VEA service is not running, restart it:

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

Installing and configuring Storage Foundation Cluster File System for Oracle RAC and Oracle RAC

You can install Storage Foundation Cluster File System for Oracle RAC on clusters of up to 16 nodes.

The following packages are installed on each cluster node:

- Veritas Cluster Server (VCS)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)
- Oracle Disk Manager (ODM)

You can configure the following components for Storage Foundation Cluster File System for Oracle RAC:

- Veritas Cluster Server (VCS)

  Note: You can not configure VCS to manage Oracle Clusterware.

- CVM (Veritas Volume Manager enabled for clusters)
- CFS (Veritas File System enabled for clusters)

Installing Storage Foundation Cluster File System for Oracle RAC

The following procedure describes the installation of an Storage Foundation Cluster File System for Oracle RAC cluster with two nodes: 'galaxy' and 'nebula'.

To install on multiple systems, set up the systems such that commands between systems execute without prompting for password or confirmation.

The product installer 'installsfcfsrac' is the recommended program to license and install Storage Foundation Cluster File System for Oracle RAC.

Note: Default responses are enclosed in parentheses. Press 'Return' to accept default values.

To install Storage Foundation Cluster File System for Oracle RAC

1. Insert the product disc with the Storage Foundation Cluster File System for Oracle RAC software into a drive connected to the system.
2. Navigate to the directory containing the installation program.

   ```
   # cd /dvd_mnt/distribution_arch/\
   storage_foundation_cluster_file_system_for_oracle_rac/
   ```

3. Depending on the installation program you use, type the appropriate command:

   Using `installer` program:

   ```
   # ./installer galaxy nebula
   ```

   Choose "I" for "Install/Upgrade a Product" and enter the number displayed against the product name. Press Return.

   Using `installsfcfsrac` program:

   ```
   # cd storage_foundation_cluster_file_system_for_oracle_rac
   # ./installsfcfsrac galaxy nebula
   ```

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

   If the installer hangs or asks for a password, stop the installer and set up SSH or RSH communications appropriately. Then, run the installer again.

5. Enter the product license information.

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

   ```
   Enter a SFCFSRAC license key for galaxy?
   ```

6 Enter y to accept another license key, otherwise enter n to proceed.

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

The installer displays the option to install RPMs.

7 Enter the appropriate option to install the RPMs. Based on your requirement, you may install all the RPMs or select the RPMs you want to install.

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

The installer displays the list of RPMs that will be installed. Review the list of RPMs.

8 Enter n when prompted to configure Storage Foundation Cluster File System for Oracle RAC.

Note: Symantec recommends that you do not configure the software during installation. Configure the software after the installation using the `-configure` option available in the `/opt` directory. For example:
`/opt/VRTS/install/installsfcfsrac -configure`

```
Are you ready to configure SFCFSRAC? [y,n,q] (y) n
```

For instructions:

See "Configuring Storage Foundation Cluster File System for Oracle RAC" on page 51.

On completion of installation of the selected packages, the installation logs are created. The installation logs can be referred in the event of any issues encountered during the installation.

Configuring Storage Foundation Cluster File System for Oracle RAC

After installation, configure the product by running the common product installation program `installer` or the product-specific installation program `installsfcfsrac` with the `-configure` option.

No configuration changes are made to the systems until all configuration questions are completed and confirmed.

To configure Storage Foundation Cluster File System for Oracle RAC

1. Log into the system as the root user and change to the directory containing the installation program 'installsfcfsrac'.

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

2. Run the following command to configure Storage Foundation Cluster File System for Oracle RAC:

   - If you are using SSH:

     ```
     # ./installsfcfsrac -configure
     ```

   - If you are using RSH:

     ```
     # ./installsfcfsrac -rsh -configure
     ```

3. Enter the names of the systems on which you want to configure Storage Foundation Cluster File System for Oracle RAC. Press Return.

   ```
   Enter the system names separated by spaces on which to
   configure SFCFSRAC: galaxy nebula
   ```

   During the initial system check, the installer checks that communication between systems is set up appropriately.

   The installer requires that SSH commands used between systems execute without prompting for passwords or confirmations. If the installer hangs or asks for a password, stop the installer and set up SSH or RSH as required. Then, run the installer again.

4. Enter additional licenses, if required.

   ```
   Checking system licensing
   SFCFSRAC license registered on galaxy
   Do you want to enter another license key for galaxy? [y,n,q] (n) n
   ```

   If there are any Storage Foundation Cluster File System for Oracle RAC processes running, these processes are stopped. Enter Return to continue.

5 Press Return to continue.

All systems are configured to create one cluster.

Enter the unique cluster name and Cluster ID number.

```
Enter the unique cluster name: [?] raccluster_101
Enter the unique Cluster ID number between 0-65535: [b,?] 76
```

The installer discovers the NICs available on the first system and reports them.

```
Discovering NICs on galaxy ... discovered eth0 eth1 eth2 eth3
```

6 Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat link
on galaxy: [b,?] eth1
Are you sure you want to use eth1 for the first private
heartbeat link? [y,n,q,b,?] (y) y
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
galaxy: [b,?] eth2
Are you sure you want to use eth2 for the second private
heartbeat link? [y,n,q,b,?] (y) y
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 you answer `y`, be sure that the same NICs are available on each system; the installer may not verify this.

Notice that in this example, eth0 is not selected for use as a private heartbeat NIC because it is already in use as the public network interface.

7 Review the information and enter `y` to confirm.

```
Is this information correct? [y,n,q] y
```

If the information is not correct, enter `n`. The installer prompts you to enter the information again.

8 If you want to configure the product to use Veritas Security Services, enter `y`, otherwise enter `n`.

Warning: Before configuring a cluster to operate using Veritas Security Services, another system must already have Veritas Security Services installed and must 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 SFCFSRAC to use
Veritas Security Services? [y,n,q] (n) n
```

9 Enter `y` to set the username and password, otherwise enter `n`.

```
Do you want to set the username and/or password for the Admin user
(default username = admin, password= password? [y,n,q] (n)
```

Note: To add users (Administrator, Operator, or Guest), you need the user name, password, and user privileges.

For more information, see the *Veritas Cluster Server Installation guide*.

---

10 Enter `y` if you want to add another user, otherwise enter `n`.

```
Do you want to add another user to the cluster? [y,n,q] (y)
```

11 Enter `y` if the information is correct, otherwise enter `n`.

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

12 Enter y to configure SMTP notification. If you do not want to configure SMTP notification, enter n.

```
Do you want to configure SMTP notification? [y,n,q] (y) y
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the SF Notifier to use on galaxy: [b,?] (eth0) eth0
Is eth0 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 SNMP recipient
(example: user@yourcompany.com): [b,?] user@mycompany.com
Enter the minimum severity of events for which mail should be sent
to user@mycompany.com [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
```

Add other SMTP recipients, or respond **n** to continue.

Verify and confirm that the information is correct, by entering **y**, or enter it again.

13 Enter y to configure SNMP notification. If you do not want to configure SNMP notification, enter n.

```
Do you want to configure SNMP notification? [y,n,q] (y) y
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the SF Notifier to use on galaxy: [b,?] (eth0) eth0
Is eth0 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,?] galaxy
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
```

14 Enter n if you want to use the operating system device naming scheme.

Enter y if you want to use enclosure-based naming scheme. The enclosure-based naming scheme is a feature of Veritas Volume Manager. You can reference disks using a symbolic name that is more meaningful than the operating system device name.

For more information, see the *Veritas Volume Manager Administrator's Guide*.

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

15 Enter `y` if you want to specify the name of the default disk group at this time.

This step does not create the disk group. The default name specified is assumed by Veritas Volume Manager if a disk group is not specified while running commands. After installation, 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
```

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.

For more information, see the `vxdctl (1M)` manual page.

16 If you responded `y` in the previous step, enter the name of the default disk group.

```
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. [?] diskgroup001
```

17 Validate the default disk group information, and press Return.

18 Verify the fully qualified hostname of the systems.

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

19 Enabling Veritas Storage Foundation Management Server management simplifies and improves management of complex data center resources, reducing planned and unplanned down time.

To enable centralized management using Storage Foundation Manager, download Veritas Storage Foundation Management Server from:

http://go.symantec.com/vom

Press **Enter** to continue.

See the *Veritas Storage Foundation Manager Installation Guide* for more information.

20 The installation and configuration process is completed and the processes are started.

You can check the log file to confirm the installation and configuration.

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

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

Preparing to install Oracle RAC

Before you install Oracle RAC, make sure that you review the recommendations and perform the following tasks:

- Recommendations before installing Oracle RAC software
- Creating operating system groups and users
- Creating the Oracle user and groups
- Creating CRS_HOME
- Creating ORACLE_HOME
- Verifying the OCR and Vote-disk shared volumes

Recommendations before installing Oracle RAC software

Review the following recommendations before installing Oracle RAC software. Symantec and Oracle recommend local installations.

About the location of ORACLE_HOME

`ORACLE_HOME` is the location where the Oracle database binaries are installed. Select the location based on your high availability (HA) requirements.

Before installing Oracle binaries on ORACLE_HOME locally on each node or on a cluster file system on shared storage, review the following information:

- Local installations provide improved protection against a single point of failure.
- Storage Foundation Cluster File System for Oracle RAC installations provide a single Oracle installation to manage, regardless of the number of nodes. This results in reduced storage requirements and easy addition of nodes.

About the location of CRS_HOME

CRS_HOME is the location where Oracle Clusterware binaries are installed.

Creating operating system groups and users

Before starting Oracle installation, you need to create the following operating system groups and users:

- OSDBA group (dba) is necessary for Oracle database software installation.
- OSOPER group (oper) is an optional group. It must be created for users having a limited set of database administrative privileges.
- Oracle Inventory group (oinstall) is necessary for all the installations.
- Oracle user (Oracle software owner user/oracle) is necessary for performing Oracle Clusterware and database software installation. This user must have the Oracle Inventory group as the primary group and the OSDBA and OSOPER groups as the secondary groups.

For more information on creating Oracle user and groups, see the Oracle installation guide.

Creating the Oracle user and groups

On each system, create a local group and local user for Oracle. Be sure to assign the same group ID, user ID, and home directory for the user on each system.

The following procedure creates the group 'oinstall' (Oracle Inventory group) and the user 'oracle' (Oracle software owner user).

Note: When you create the user and group, make sure that you specify a user and group ID that is not in use.

To create the operating system Oracle user and group on each system

1 Create the 'oinstall' group on each system.

```
# groupadd -g 1000 oinstall

# groupadd -g 1001 dba
```

2 Create the Oracle user and the user home directory on each system:

```
# useradd -g oinstall -u 1000 \
-G dba -md /home/oracle oracle
```

Creating CRS_HOME

The following procedure provides instructions on creating CRS_HOME.

To create CRS_HOME on each system

1 Log in as root user on a system.

```
# su - root
```

2 On one of the nodes, create a disk group:

```
# vxdg init crsbindg sdc
```

3 Create the volume in the group for CRS_HOME:

```
# vxassist -g crsbindg make crsvol 1024M
```

where 1024M is the size in MB for CRS_HOME in the sample command.

For more information on exact size requirements for CRS_HOME, see the Oracle product documentation.

4 Create a VxFS file system on which to install Oracle Clusterware:

```
# mkfs -t vxfs /dev/vx/rdsk/crsbindg/crsbinv
```

5 Create the mount point for CRS_HOME:

```
# mkdir /oracle/crsbin
```

6 Mount the file system using the device file for the block device:

```
# mount -t vxfs /dev/vx/rdsk/crsbindg/crsbinvol /oracle/crsbin
```

7 Export the CRS_HOME directory as /oracle/crsbin for the oracle user.

8 Assign ownership of the directory to the user 'oracle' and the group 'oinstall':

```
# chown -R oracle:oinstall /oracle/crsbin
```

```
# chmod -R 755 /oracle/crsbin
```

9 On each cluster node, repeat step 1 through step 8.

10 After creating the volume and file system for CRS_HOME, add them to the VCS configuration. This automatically starts the volume and file system when the system starts.

For more information on configuring a volume and file system under VCS:

Creating ORACLE_HOME

The following procedure provides instructions on creating ORACLE_HOME:

To create ORACLE_HOME on each system

1 Log in as root user on a system.

```
# su - root
```

2 On one of the nodes, create a disk group:

```
# vxdg init orabindg sde
```

For shared ORACLE_HOME, run the following command on the CVM master:

```
# vxdg -s init orabindg sde
```

3 Create the volume in the group for ORACLE_HOME:

```
# vxassist -g orabindg make orabinvol 500M
```

For more information on exact size requirements for ORACLE_HOME, see the Oracle product documentation.

4 Create a VxFS file system on which to install database:

```
# mkfs -t vxfs /dev/vx/rdsk/orabindg/orabinvol
```

5 Create the mount point for ORACLE_HOME:

```
# mkdir /oracle/orabin
```

6 Mount the file system using the device file for the block device:

```
# mount -t vxfs /dev/vx/dsk/orabindg/orabinvol /oracle/orabin
```

For each shared ORACLE_HOME, run the following command:

```
# mount -t vxfs -o cluster /dev/vx/dsk/orabindg/orabinvol \
/oracle/orabin
```

7 Export the ORACLE_HOME directory as /oracle/orabin for the oracle user.

8 Assign ownership of the directory to the user 'oracle' and the group 'oinstall':

```
# chown -R oracle:oinstall /oracle/orabin
```

```
# chmod -R 755 /oracle/orabin
```

9 For each local ORACLE_HOME, repeat step 1 through step 8 on each cluster node.

For each shared ORACLE_HOME, repeat step 5 through step 6 on each cluster node.

10 After creating the volume and file system for ORACLE_HOME, add them to the VCS configuration. This automatically configures the volume and file system when the system starts.

For more information on configuring a volume and file system under VCS:

Configuring private IP addresses

The CRS daemon requires a private IP address on each system to enable communications and heartbeating.

For information on configuring private IP addresses on all of the cluster nodes, see the Oracle product documentation.

Obtaining public virtual IP addresses for use by Oracle

Before starting the Oracle installation, configure a virtual IP address for each node. Register an IP address and an associated host name in DNS for each public network interface.

For information on configuring virtual IP addresses on all of the cluster nodes, see the Oracle product documentation.

Creating OCR and Vote-disk volumes

Create the OCR and Vote-disk shared volumes. The CFS directory is not supported by Oracle for OCR and Vote-Disk.

The installation of Oracle Clusterware requires a predefined location for the OCR and Vote-disk components.

To create the OCR and Vote-disk shared volumes

1. Log in as root user.
2. On the master node, create a shared disk group. For example:

   ```
   # vxdg -s init ocrvotedg sdd
   ```

   where `ocrvotedg` is the OCR Vote-disk group.
3. Create volumes in the shared group for OCR and Vote-disk. For example:

   ```
   # vxassist -g ocrvotedg make ocrvol 500M

   # vxassist -g ocrvotedg make vdvol 500M
   ```

   where `ocrvotedg` is the OCR Vote-disk group.

   where `ocrvol` is the OCR volume.

   where `vdvol` is the Vote-disk volume.

   where `500M` is the size of the volumes in MB.

   See the documentation for minimum size requirements for OCR and Vote-disk.

   The OCR volume can be mirrored to provide high availability. During the Oracle Clusterware operation, if one of the mirrors fail, then the other mirror can be used immediately without interrupting Oracle Clusterware functionality.

   For more information on creating mirrored volumes, see the *Veritas Volume Manager Administrator's Guide*.

   Mirroring of the OCR volume is not a requirement from Oracle. More than one volume can be specified for Vote-disk, if Normal Redundancy option is selected in the Specify Voting Disk Location screen of the Oracle Clusterware installer.
4. Assign ownership of the volumes using the `vxedit` command. For example:

   ```
   # vxedit -g ocrvotedg set user=oracle group=oinstall mode=660 ocrvol

   # vxedit -g ocrvotedg set user=oracle group=oinstall mode=660 vdvol
   ```

   where `ocrvotedg` is the OCR Vote-disk group.

   where `root` and `oracle` are the user names.

   where `oinstall` is the group name.

where `640` and `644` are the mode values.

where `ocrvol` is the OCR volume.

where `vdvol` is the Vote-disk volume.

After creating shared volumes for OCR and Vote-disk, you may proceed with the VCS configuration. This automatically starts the OCR and Vote-disk when the system starts.

For more information on configuring a volume and file system under VCS:

Verifying the OCR and Vote-disk shared volumes

Verify the OCR and Vote-disk shared volumes that have been created and configured.

To verify the OCR and Vote-disk shared volumes

- Run the following command to verify that the OCR and Vote-disk shared volumes are present:

```
# ls -l /dev/vx/rdsk/ocrvotedg/*

crw-r----- 1 root oinstall 199, 3 Jun 26 15:58
/dev/vx/rdsk/ocrvotedg/ocrvol
crw-r--r-- 1 oracle oinstall 199, 4 Jun 26 15:58
/dev/vx/rdsk/ocrvotedg/vdvol
```

These shared volumes are used during the Oracle Clusterware installation. During the Oracle Clusterware installation, use the following locations whenever requested by the Oracle Clusterware installer.

For OCR, use: /dev/vx/rdsk/ocrvotedg/ocrvol

For Vote-disk, use: /dev/vx/rdsk/ocrvotedg/vdvol

Installing Oracle Clusterware and database software

For information on installing Oracle Clusterware and database software, see the Oracle product documentation.

Note: Storage Foundation Cluster File System for Oracle RAC does not support Symantec's implementation of SCSI-3 PGR-based I/O fencing. Oracle Clusterware is expected to handle any split-brain situations. See the following TechNote for more information: http://entsupport.symantec.com/docs/306411

Completing the post-installation tasks

Perform the following tasks after installing Oracle RAC:

- Relinking with ODM
- Creating Oracle databases
- Increasing the peer inactivity timeout of LLT
- Setting the start and stop init sequence for VCS and Oracle Clusterware
- Configuring LLT to use bonded network interfaces (optional)

Relinking with ODM

After installing Oracle database, you must relink Oracle database with Veritas Extension for Oracle Disk Manager (ODM).

If ORACLE_HOME is on a shared file system, run the following commands from any node, otherwise run them on each node.

ORACLE_HOME is the location where Oracle database binaries have been installed.

To configure Veritas Extension for Oracle Disk Manager

1. Log in as `oracle` user.
2. If the Oracle database is running, then shut down the Oracle database.
3. Verify that the file `/opt/VRTSodm/lib64/libodm.so` exists.
4. Link Oracle's ODM library present in `ORACLE_HOME` with the Veritas Extension for Oracle Disk Manager library:

   For Oracle RAC 10g:

   - Change to the `$ORACLE_HOME/lib` directory:

     ```
     # cd $ORACLE_HOME/lib
     ```

   - Back up `libodm10.so` file.

     ```
     # mv libodm10.so libodm10.so.oracle-`date '+%m_%d_%y-%H_%M_%S'`
     ```

   - Link `libodm10.so` file with the Veritas ODM library:

     ```
     # ln -s /opt/VRTSodm/lib64/libodm.so libodm10.so
     ```

5 Start the Oracle database.

6 To confirm that the Oracle database starts with Veritas Extension for ODM, check the alert log for the following text:

```
Veritas <version> ODM Library
```

where *<version>* is the ODM library version shipped with the product.

The alert log location depends on the Oracle version used.

For more information on the exact location of the alert log, see the Oracle documentation.

Creating Oracle databases

This section provides instructions for creating Oracle RAC 10g and Oracle RAC 11g database tablespaces. You can create database tablespaces on shared raw VxVM volumes or on CFS.

Before you create database tablespaces:

- Make sure that CRS daemons are running.
  To verify the status of Oracle Clusterware, enter:

  ```
  # $CRS_HOME/bin/crsctl check crs
  ```

  The following text displays a sample output that verifies the status of CRS daemons:

  ```
  Cluster Synchronization Services appears healthy
  Cluster Ready Services appears healthy
  Event Manager appears healthy
  ```

- Verify that all private IP addresses required by Oracle Clusterware on each node are up.

Creating database tablespaces on shared raw VxVM volumes

This section provides instructions for creating database tablespaces on shared raw VxVM volumes.

To create database tablespace on shared raw VxVM volumes

1. Log in as the root user.
2. On any node in the cluster, enter the following command to locate the CVM master:

```
# vxdctl -c mode

mode: enabled: cluster active - MASTER
master: galaxy
```

   The above sample output indicates that `galaxy` is the CVM master.

3. On the CVM master, identify the spare disks that can be used for creating shared disk group for Oracle database tablespaces:

```
# vxdisk -o alldgs list

DEVICE TYPE DISK GROUP STATUS
sda auto:none - - online invalid
sdb auto:none - - online invalid
sdc auto:cdsdisk - tempdg online shared
sdd auto:none - ocrvotedg online shared
sde auto:cdsdisk - - online shared
sdf auto:cdsdisk - - online shared
```

   The above sample output indicates that the shared disks `sde` and `sdf` are free and may be used for Oracle database tablespaces.

   Check if the disks are of sufficient size. If the size is not sufficient for the available disks, then you may need to add additional disks to the system.

   For more information on size requirements, see the Oracle documentation.

4. On the CVM master node, create a shared disk group:

```
# vxdg -s init oradatadg sde sdf
```

5 Create a volume in the shared disk group for each of the required tablespaces.

For example, enter:

```
# vxassist -g oradatadg make VRT_volume01 1000M

# vxassist -g oradatadg make VRT_volume02 10M
.
.
.
```

For more information, see the Oracle documentation specific to the Oracle database release to determine the tablespace requirements.

6 Define the access mode and permissions for the volumes that store Oracle data. For each volume required for Oracle database tablespaces, run the `vxedit` command as follows:

```
# vxedit -g disk_group set group=group user=user mode=660 \
<volume_name>
```

For example:

```
# vxedit -g oradatadg set group=oinstall user=oracle mode=660 \
VRT_volume01
```

In this example, `VRT_volume01` is the name of one of the volumes.

Repeat the command to define access mode and permissions for each volume in the `oradatadg`.

For more information about the command, see the `vxedit` (1M) manual page.

To automatically start the shared disk group by VCS, you need to configure the shared disk group under VCS.

For more information on configuring a volume and file system under VCS:

7 Create the database using the Oracle documentation.

Creating database tablespaces on CFS

If you plan to use CFS to store the Oracle database, use the following procedure to create the file system.

To create database tablespaces on CFS

1. Log in as the root user.
2. On any node in the cluster, enter the following command to locate the CVM master:

```
# vxdctl -c mode

mode: enabled: cluster active - MASTER
master: galaxy
```

   The above sample output indicates that `galaxy` is the CVM master.

3. On the CVM master, identify the spare disks that can be used for creating shared disk group for Oracle database tablespaces:

```
# vxdisk -o alldgs list

DEVICE TYPE DISK GROUP STATUS
sda auto:none - - online invalid
sdb auto:none - - online invalid
sdc auto:cdsdisk - tempdg online shared
sdd auto:none - ocrvotedg online shared
sde auto:cdsdisk - - online shared
sdf auto:cdsdisk - - online shared
```

   The above sample output indicates that shared disks `sde` and `sdf` are free and can be used for Oracle database tablespaces.

   Check if the disks are of sufficient size. If the size is not sufficient for the available disks, then you may need to add additional disks to the system.

   For more information on size requirements, see the Oracle documentation.

4. Create a shared disk group. For example, enter:

```
# vxdg -s init oradatadg sde sdf
```

5. Create a single shared volume that is large enough to contain a file system for all tablespaces.

   The following command assumes 6.8 GB of space for the tablespaces:

```
# vxassist -g oradatadg make oradatavol 6800M
```

   For more information about tablespace sizes, see the Oracle documentation specific to the Oracle database release.

6 Create a VxFS file system in this volume:

```
# mkfs -t vxfs /dev/vx/rdsk/oradatadg/oradatavol
```

7 Create a mount point for the shared file system:

```
# mkdir /oradata
```

8 From the same node, mount the file system:

```
# mount -t vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol \
/oradata
```

To automatically start the file system by VCS, you need to configure the file system under VCS.

For more information on configuring a volume and file system under VCS:

9 Set `oracle` as the owner of the file system, and set `775` as the permission:

```
# chown oracle:oinstall /oradata
```

```
# chmod 775 /oradata
```

10 On the other nodes, complete steps 7 through 8.

11 Create the Oracle database using the Oracle documentation.

Increasing the peer inactivity timeout of LLT

Storage Foundation Cluster File System for Oracle RAC does not support Symantec's implementation of I/O fencing. Oracle Clusterware must handle any split-brain situations. In the presence of two clusterwares (VCS and Oracle Clusterware), there is a high possibility of data corruption due to the lack of co-ordination between the clusterwares.

Note: To prevent data corruption, you must modify the LLT peer inactivity timeout settings.

For instructions, refer to the following technote:

http://entsupport.symantec.com/docs/306411

Setting the start and stop init sequence for VCS and Oracle Clusterware

VCS and Oracle Clusterware are interdependent services in Storage Foundation Cluster File System for Oracle RAC. The mounts and volumes on which Oracle

Clusterware resides may be controlled by VCS. Moreover, the OCR and Vote disks used by Oracle Clusterware are configured under VCS. This implies that VCS must start before Oracle Clusterware. Likewise, the volumes and mount points must not be in use by Oracle Clusterware when VCS attempts to take them offline. Therefore, Oracle Clusterware must stop before VCS. Since there is no inherent coordination between VCS and Oracle Clusterware in Storage Foundation Cluster File System for Oracle RAC, you need to ensure that VCS and Oracle Clusterware are started and stopped in the correct order by modifying the numbering of the start and stop scripts for these services in the appropriate run levels.

To start VCS before Oracle Clusterware, modify the numbering such that the VCS start script (S*vcs) ranks lower in number to the Oracle Clusterware start script (S*init.crs) in the appropriate run levels.

Note: If the sequence for the start script is not set correctly, Oracle Clusterware fails to start as the binaries may not be available when the Oracle Clusterware start script is invoked.

To stop Oracle Clusterware before VCS, modify the numbering such that the Oracle Clusterware stop script (K*init.crs) ranks lower in number to the VCS stop script (K*vcs) in the appropriate run levels.

Note: If the sequence for the stop script is not set correctly, the 'shutdown -r now' command hangs indefinitely in VCS as a result of shared volumes and mount points in use by Oracle Clusterware.

Configuring LLT to use bonded network interfaces (optional)

This is an optional task that may be performed after installation.

If you have configured LLT to use a single bonded network interface, GAB reports jeopardy membership even if there is more than one interface beneath the bonded interface.

To prevent GAB from reporting jeopardy membership, it is recommended that you add the following line in the /etc/llttab file:

```
set-dbg-minlinks 2
```

After you update the /etc/llttab file, when LLT is restarted, GAB does not report jeopardy membership even if only one bonded interface is specified in the /etc/llttab file.

For more information, see the following technote:

http://entsupport.symantec.com/docs/308107

For more information, see the following documents:

Veritas Volume Manager Administrator's Guide

Veritas Storage Foundation for Cluster File System Administrator's Guide

Veritas Storage Foundation Installation Guide

Upgrading Storage Foundation Cluster File System for Oracle RAC

You can upgrade to 5.0 RU4 from the following versions:

- Storage Foundation Cluster File System for Oracle RAC 5.0 MP2
- Storage Foundation Cluster File System for Oracle RAC 5.0 MP3
- Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 RP2
- Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 RP3
- Storage Foundation Cluster File System for Oracle RAC 5.0 RU3

Storage Foundation Cluster File System for Oracle RAC software must be at the same version across all nodes in an Storage Foundation Cluster File System for Oracle RAC cluster, in this case 5.0 Release Update 4.

Note: Symantec strongly recommends upgrading all Storage Foundation Cluster File System for Oracle RAC component products to the same version at the same time. In a CVM/CFS environment, dependencies between Storage Foundation Cluster File System for Oracle RAC component products will not be met if you do not upgrade all components to the same version.

Supported upgrade paths for SLES distributions

Table 2-3 lists the supported upgrade paths for SLES distributions.

Table 2-3 Supported upgrade paths for SLES distributions

Architecture	From Storage Foundation Cluster File System for Oracle RAC version	On SLES version	To SLES and Storage Foundation Cluster File System for Oracle RAC version
x86_64	Storage Foundation Cluster File System for Oracle RAC 5.0 MP2 Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 RP2 Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 RP3	SLES 9 SP3 SLES 9 SP4	SLES 10 SP3 Storage Foundation Cluster File System for Oracle RAC 5.0 RU4
x86_64	Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 RP2 Storage Foundation Cluster File System for Oracle RAC 5.0 MP3 RP3	SLES 10 SP1 SLES 10 SP2	SLES 10 SP3 Storage Foundation Cluster File System for Oracle RAC 5.0 RU4
ppc64	Storage Foundation Cluster File System for Oracle RAC 5.0 RU3	SLES 10 SP2	SLES 10 SP3 Storage Foundation Cluster File System for Oracle RAC 5.0 RU4

Upgrading Storage Foundation Cluster File System for Oracle RAC and operating system

This section describes how to perform a full upgrade to version 5.0 RU4.

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

To prepare for a full upgrade

1. Log in as superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
3. Stop the applications that are not managed by VCS. Use native application commands to stop the application.
4. Stop Oracle Clusterware.

   ```
   # /etc/init.d/init.crs stop
   ```

5. Stop high-availability cluster operations. This command can be executed from any node in the cluster, and stops cluster operations on all the nodes.

   ```
   # hastop -all
   ```

6. Check if each node's root disk is under VxVM control by running this command:

   ```
   # df -v /
   ```

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

   - Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

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

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

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

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

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

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

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

9 Unmount all Storage Checkpoints and file systems:

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

10 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

  ```
  # vxrvg -g diskgroup stop rvg_name
  ```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

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

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

11 Check if the VEA service is running:

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

12 If the VEA service is running, stop it:

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

13 If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.

14 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.

15 Stop all VxVM volumes by entering the following command for each disk group:

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

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

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

17 Disable the startup scripts before upgrading the operating system.

```
# insserv -r vcs
# insserv -r vxfen
# insserv -r vxgms
# insserv -r gab
# insserv -r llt
```

18 Upgrade the operating system. For instructions, see the operating system documentation.

To upgrade Storage Foundation Cluster File System for Oracle RAC

1 Insert the appropriate media disc into your system's DVD-ROM drive.

2 If volume management software is running on your system, the software disc automatically mounts as `/mnt/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

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

3 Change to the top-level directory on the disc:

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

4 Invoke the `installmp` command from one of your cluster nodes:

```
# ./installmp node_name1 node_name2 [-rsh]
```

5 After the initial system checks are complete, press **Return** to start the requirement checks.

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

7 Shut down and reboot the systems.

8 Upgrade Oracle RAC, if required.

See "Upgrading the Oracle database" on page 76.

9 Relink the ODM library.

See "Relinking with ODM" on page 64.

To bring the upgraded cluster online and restore components

1. If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.
2. If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.
3. If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.
4. Restart all the volumes by entering the following command for each disk group:

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

5. If you have stopped any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, restart each RVG:

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

6. Remount all VxFS file systems and Storage Checkpoints on all nodes:

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

7. Check if the VEA service was restarted:

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

8. If the VEA service is not running, restart it:

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

9. Start the applications that are not managed by VCS. Use native application commands to start the applications.

Upgrading the Oracle database

For instructions on upgrading the Oracle database, see the appropriate Oracle Metalink documents on the Oracle Web site.

To upgrade Oracle database to 10.2.0.4	See Oracle Metalink Doc ID 316889.1

To upgrade Oracle database to 11.1.0.6	See Oracle Metalink Doc ID 429825.1

Note: After upgrading the database, relink the Oracle library with the Veritas ODM library.

See "Relinking with ODM" on page 64.

Installing Veritas Cluster Server

There are no changes for fresh installs of Veritas Cluster Server in this 5.0 RU4 release.

Refer to the chapter Installing and configuring VCS in the *Veritas Cluster Server 5.0 MP3 Installation Guide* for more information.

Upgrading Veritas Cluster Server

If you are currently running a cluster with VCS 5.0 MP3 on SLES10SP2, you can use the installer to upgrade to VCS 5.0 RU4 on SLES 10 SP3. The procedure to upgrade a cluster to VCS 5.0 RU4 is described as follows.

Performing the pre-upgrade tasks

Perform the following procedures before you upgrade VCS.

To perform pre-upgrade tasks

1. Log in as superuser on one of the nodes where you want to upgrade VCS.
2. Mount the software disc
3. Verify that /opt/VRTS/bin is set in your PATH environment variable to execute all product commands.
4. Make the VCS configuration writable.

   On a node that you want to upgrade, type the following command:

   ```
   # haconf -makerw
   ```

5. Freeze all the service groups in the configuration.

   To freeze the service groups, on each node that you selected to upgrade, type:

   ```
   # hagrp -freeze groupname -persistent
   ```

6. Save the VCS configuration.

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

Stopping VCS and its components

Perform the following steps to stop VCS and its components:

To stop VCS and its components

1. Stop VCS if it is already running.
2. On each node, run the following command:

   ```
   # hastop -local -force
   ```

3. Stop I/O fencing on each node:

   ```
   # /sbin/vxfenconfig -U
   ```

4. Stop GAB on each node:

   ```
   # /sbin/gabconfig -U
   ```

5. Stop LLT on each node:

   ```
   # /sbin/lltconfig -Uo
   ```

Upgrading the Linux operating system

Upgrade the Linux operating system, if necessary, on each node from SLES10 SP2 to SLES10 SP3.

Perform the following procedure to upgrade VCS to version 5.0RU4.

Refer to Linux documentation for upgrade details.

To upgrade VCS with installmp

1. Log in as superuser.
2. Insert the 5.0 RU4 software disc into the disc drive of one of the nodes.
3. Navigate to the folder that contains the installmp program.
4. Upgrade to VCS 5.0RU4.

   Enter the following command:

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

   The installer begins with a copyright message and specifies the directory where the logs are created.
5. Review the output as the program verifies that the upgrade can proceed on the systems. The installvcs program checks the operating system level and system-to-system communication. It also creates logs for the VCS update.

6 Review the output as the installer checks for existing rpms. The installer lists the rpms that will be installed or upgraded.

Press **Enter** at the prompt

The installer is now ready to upgrade VCS.

7 To upgrade to VCS 5.0RU4, enter **y**at the following prompt

Are you sure you want to install RU4? [y,n,q] **(y)**

8 View the output as the program attempts to stop VCS and uninstall the rpms.

The installer also uninstalls the deprecated rpms.

Progress indicators show the status of these tasks.

The program now attempts to install VCS 5.0 RU4 rpms. A progress indicator shows the status of the task.

Starting VCS

After the `installmp` program upgrades the product, the installer prompts to reboot or manually start up the product processes.

- If you have not already configured the product, use the `-configure` option with the appropriate product installation script in the `/opt/VRTS/install/` directory.
- If you want to manually start VCS and its components, enter the following commands:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /sbin/gabconfig -cx
# /etc/init.d/vxfen start
```

- If you choose to reboot the upgraded cluster nodes, run the following command to properly restart your systems:

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

VCS should restart on all the nodes.

Completing the upgrade

The upgrade is complete after starting the cluster server processes.

To complete the upgrade

1. Note the locations of the summary and log files that the program creates. Upgrade log files and summary file are saved at:

   ```
   /opt/VRTS/install/logs/installvcs-unique string/
   ```

2. Other files specific to the installer program are also created in the `/opt/VRTS/install/logs/installvcs-uniquestring/` directory.

3. After VCS is up and running, make the VCS configuration writable. Type the following command:

   ```
   # haconf -makerw
   ```

4. Unfreeze all the service groups in the configuration. On a node, for each group, type the following command:

   ```
   # hagrp -unfreeze groupname -persistent
   ```

5. Save the VCS configuration:

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

Creating new VCS accounts if you used native OS accounts

VCS has deprecated the `AllowNativeCliUsers` attribute. To use native OS accounts with VCS, use the `halogin` command. After you run the `halogin` command, VCS encrypts and stores your VCS credentials in your home directory for a specific time period. After you run the `halogin` command, you need not authenticate yourself every time you run a VCS command.

In secure clusters, the command also sets up a trust relationship and retrieves a certificate from an authentication broker. See the *Veritas Cluster Server User's Guide* for information on assigning user privileges to OS user groups for clusters running in secure mode and clusters not running in secure mode.

Perform the following procedure if you used the `AllowNativeCliUsers` attribute. Ensure that each native user running VCS commands has a home directory on the system from which the user runs VCS commands.

To set up VCS authentication for clusters running in secure mode

1 Set the configuration (main.cf) mode to read/write.

```
# haconf -makerw
```

2 Assign proper privileges to the OS users or user groups.

Each OS user must perform steps 3 and 4.

3 If the user executes VCS commands from a remote host, set the following environment variables:

VCS_HOST—Name of the VCS node on which you run commands. You may specify the virtual IP address associated with the cluster.

VCS_DOMAIN—Name of the VxSS domain to which the user belongs.

VCS_DOMAINTYPE—Type of VxSS domain: unixpwd, nt, nis, nisplus, or vx.

4 Run the `halogin` command:

```
$ halogin vcsusername password
```

To set up VCS authentication for clusters not running in secure mode

1 Set the configuration (main.cf) mode to read/write.

```
# haconf -makerw
```

2 Create VCS user accounts for all users and assign privileges to these users.

3 Each VCS user must run the halogin command:

```
$ halogin vcsusername password
```

Removing 5.0RU4

This section provides instructions for removing 5.0RU4. Follow the procedure for the Veritas product that you have installed.

The uninstalling steps are the same as for 5.0MP3. For additional details, see the 5.0MP3 installation guide for your product.

To uninstall Veritas Volume Replicator, see the 5.0MP3 *Storage Foundation Installation Guide*.

Storage Foundation or Storage Foundation for databases	See "Uninstalling Storage Foundation or Storage Foundation for databases " on page 82.
Storage Foundation Cluster File System	See "Uninstalling Storage Foundation Cluster File System" on page 84.

Veritas Cluster Server	See "Uninstalling Veritas Cluster Server" on page 88.
Storage Foundation Cluster File System for RAC	See "Uninstalling Storage Foundation Cluster File System for Oracle RAC" on page 86.

Uninstalling Storage Foundation or Storage Foundation for databases

This section provides instructions for uninstalling Storage Foundation or Storage Foundation for databases. You need to complete the preparatory tasks before you uninstall Storage Foundation or Storage Foundation for databases.

Preparing to uninstall Storage Foundation or Storage Foundation for databases

Perform the steps in the following procedure before you uninstall Storage Foundation or Storage Foundation for databases. Some steps do not apply if you are uninstalling in a stand-alone configuration. These steps are noted as applying only to an HA configuration.

To prepare to uninstall Storage Foundation or Storage Foundation for databases from a cluster

1. Log in as the root user. In an HA configuration, log in on any node in the cluster.
2. Verify that the following directories are set in your PATH environment variable in order to execute the necessary commands:

   ```
   /opt/VRTS/bin

   /opt/VRTSvcs/bin
   ```

3. In an HA configuration, back up the following configuration files:

   ```
   # mv /etc/lltab /etc/lltab.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
   ```

4. In a stand-alone configuration for Storage Foundation for Oracle or Storage Foundation for DB2, use the `sfua_db_config` command with the `-o dropdb` option to remove the database and unmount the database repository volume.

   ```
   # /opt/VRTSdbcom/bin/sfua_db_config -o dropdb
   # /opt/VRTSdbcom/config/sfua_rep_mount stop
   ```

5 In an HA configuration, stop VCS:

```
# hastop -all
```

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

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

7 Unmount all Storage Checkpoints and file systems:

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

8 If the root disk is encapsulated, remove rootability.

For details, see the 5.0MP3 *Storage Foundation Installation Guide*.

9 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g disk_group stopall
```

To verify that no volumes are open:

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

Uninstalling Storage Foundation or Storage Foundation for databases

Perform the steps in the following procedure to uninstall Storage Foundation or Storage Foundation for databases from a cluster.

To uninstall Storage Foundation or Storage Foundation for databases

1 Log in as the root user on any node in the cluster.

2 Navigate to the directory that contains the uninstallation program:

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

3 Start the uninstallation program.

For Storage Foundation:

```
# ./uninstallsf
```

For Storage Foundation for DB2:

```
# ./uninstallsfdb2
```

For Storage Foundation for Oracle

```
# ./uninstallsfora
```

For Storage Foundation for Sybase

```
# ./uninstallsfsyb
```

4 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall Storage Foundation, for example, `host1`:

```
Enter the system names separated by spaces from which to
uninstall Storage Foundation: host1
```

5 Confirm the uninstallation at the following prompt:

```
Are you sure you want to uninstall SF [y,n,q] (y)
```

The installer stops the Storage Foundation or Storage Foundation for databases processes and uninstalls the packages.

6 Reboot the nodes:

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

Uninstalling Storage Foundation Cluster File System

This section provides instructions for uninstalling Storage Foundation or Storage Foundation for databases. You need to complete the preparatory tasks before you uninstall Storage Foundation or Storage Foundation for databases.

Preparing to uninstall Storage Foundation Cluster File System

Perform the steps in the following procedure before you uninstall Storage Foundation or Storage Foundation for databases.

To prepare to uninstall Storage Foundation Cluster File System

1. Log in as the root user on any node in the cluster.
2. Verify that the following directories are set in your PATH environment variable in order to execute the necessary commands:

   ```
   /opt/VRTS/bin
   ```

   ```
   /opt/VRTSvcs/bin
   ```

3. Back up the following configuration files:

   ```
   # mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
   ```

4. Stop VCS:

   ```
   # hastop -all
   ```

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

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

6. Unmount all Storage Checkpoints and file systems:

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

7. Stop all VxVM volumes by entering the following command for each disk group:

   ```
   # vxvol -g disk_group stopall
   ```

   To verify that no volumes are open:

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

Uninstalling Storage Foundation Cluster File System

Perform the steps in the following procedure to uninstall Storage Foundation Cluster File System.

To uninstall Storage Foundation Cluster File System

1. Log in as the root user on any node in the cluster.
2. Navigate to the directory that contains the uninstallation program:

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

3. Start the uninstallation program:

   ```
   # ./uninstallsfcfs galaxy nebula
   ```

4. Confirm the uninstallation at the following prompt:

   ```
   Are you sure you want to uninstall SFCFS [y,n,q] (y)
   ```

   The installer stops the Storage Foundation Cluster File System processes and uninstalls the packages.
5. Reboot the nodes:

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

Uninstalling Storage Foundation Cluster File System for Oracle RAC

This section provides instructions for uninstalling Storage Foundation Cluster File System for Oracle RAC. You need to complete the preparatory tasks before you uninstall Storage Foundation Cluster File System for Oracle RAC.

Preparing to uninstall Storage Foundation Cluster File System for Oracle RAC from a cluster

Perform the steps in the following procedure before you uninstall Storage Foundation Cluster File System for Oracle RAC from a cluster.

To prepare to uninstall Storage Foundation Cluster File System for Oracle RAC from a cluster

1. Log in as the root user on any node in the cluster.
2. Verify that the following directories are set in your PATH environment variable in order to execute the necessary commands:

   ```
   /opt/VRTS/bin

   /opt/VRTSvcs/bin
   ```

3 Back up the following configuration files:

```
# mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`
# mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
```

4 On all the nodes, stop the CFS-dependant applications that are not under VCS control using application specific commands.

For example, to stop Oracle Clusterware:

```
# /etc/init.d/init.crs stop
```

5 Stop VCS:

```
# hastop -all
```

6 Verify that port h is not open:

```
# gabconfig -a
```

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

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

8 Unmount all Storage Checkpoints and file systems:

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

9 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g disk_group stopall
```

To verify that no volumes are open:

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

Uninstalling Storage Foundation Cluster File System for Oracle RAC from a cluster

Perform the steps in the following procedure to uninstall Storage Foundation Cluster File System for Oracle RAC from a cluster.

To uninstall Storage Foundation Cluster File System for Oracle RAC from a cluster manually

1. Log in as the root user on any node in the cluster.
2. Navigate to the directory that contains the uninstallation program:

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

3. Start the uninstallation program:

   ```
   # ./uninstallsfcfsrac galaxy nebula
   ```

4. Press Enter to uninstall Storage Foundation Cluster File System for Oracle RAC.

   ```
   Do you want to uninstall SFCFSRAC from these systems [y,n,q] (y)
   ```

   The installer checks the RPMs installed on the system.
5. Confirm the uninstallation at the following prompt:

   ```
   Are you sure you want to uninstall SFCFSRAC [y,n,q] (y)
   ```

   The installer stops the Storage Foundation Cluster File System for Oracle RAC processes and uninstalls the packages.
6. Reboot the nodes:

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

Uninstalling Veritas Cluster Server

This section provides instructions for uninstalling Veritas Cluster Server. You need to complete the preparatory tasks before you uninstall Veritas Cluster Server.

Prerequisites before you uninstall VCS

Review the following prerequisites before you uninstall VCS:

- Before you remove VCS from any node in the cluster, shut down the applications that depend on VCS. For example, applications such as Java Console or any high availability agents for VCS
- Before you remove VCS from fewer than all nodes in a cluster, stop the service groups on the nodes from which you uninstall VCS. You must also reconfigure VCS on the remaining nodes.

- If you have manually edited any of the VCS configuration files, you need to reformat them.

Refer to the 5.0MP3 VCS Installation Guide for details.

Uninstalling Veritas Cluster Server

Perform the steps in the following procedure to uninstall Veritas Cluster Server

To uninstall Veritas Cluster Server

1. Log in as the root user on any node in the cluster.
2. Navigate to the directory that contains the uninstallation program:

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

3. Start the uninstallation program:

   ```
   # ./uninstallvcs galaxy nebula
   ```

4. Confirm the uninstallation at the following prompt:

   ```
   Are you sure you want to uninstall VCS [y,n,q] (y)
   ```

   The installer stops the Veritas Cluster Server processes and uninstalls the packages.

5. Reboot the nodes:

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

Chapter 3

Changes in Storage Foundation

This chapter includes the following topics:

Changes in Storage Foundation

Storage Foundation 5.0 Release Update 4 (RU4) is based on Veritas Storage Foundation 5.0 MP3 RP2. It includes support for SuSE Linux Enterprise Server 10 Service Pack 3 (SLES10 SP3) and SuSE Linux Enterprise Server 11 (SLES11) with DBED. Review previous versions of the Storage Foundation 5.0 (Linux) Release Notes at the following URL:

http://entsupport.symantec.com/docs/306947

Storage Foundation 5.0 RU4 package changes for fresh installs

For fresh installs, the following packages are replaced by VRTSsfmh:

- VRTSaa

- VRTSccg
- VRTSdcli
- VRTSmh

The following packages are obsolete in this 5.0 RU4 release and will not be installed during fresh installs. These packages will be uninstalled during an in-place upgrade:

- VRTScmccc
- VRTScmcs
- VRTSjre
- VRTSjre15
- VRTSweb
- VRTSvrw

Storage Foundation changes to support compatible disk layout (CDL) disks

Although currently undergoing development discussions, there is no partition table library available from IBM or Linux to manipulate CDL-formatted disks. Due to the absence of a library for reading, writing, modifying, tagging, deleting, or creating CDL partitions, VxVM can only support CDL-formatted DASD disks in a limited way:

- VxVM can format CDL-formatted disks as simple disks only.
  For example, there is no support for `cdsdisk` or sliced disk format.

The `cdsdisk` format is not supported for DASD disks because these disks are not SCSI compatible.

The Sliced disk format is not supported for DASD disks because there is no library available to create separate partitions for the private and public region from within VxVM.

- The `vxdisk` list displays all 4K formatted disks, including those that are not under VxVM control, as `TYPE:auto:simple`.

However, the disks not under VxVM control are displayed as `STATUS:online invalid`. Due to the non-availability of a library to manipulate CDL partitions, VxVM is not able to put a partition tag on the disks under its control. Therefore, it assumes that all disks could potentially be VxVM disks. As a result, disks not under VxVM control do not have a valid private region and are identified as `STATUS:online invalid`.

Thin Storage Reclamation support

Thin Storage is an array vendor solution for allocating storage to applications only when the storage is truly needed, from a pool of free storage. Thin Storage attempts to solve the problem of under utilization of available array capacity.

Thin Storage Reclamation-capable arrays and LUNs allow the administrators to release once-used storage to the pool of free storage. Storage is allocated from the free pool when files are created and written to in the file system. However, this storage is not released to the free pool when files get deleted; the administrator must perform the operation of reclaiming this storage for the free pool.

Veritas File System supports reclamation of the free blocks in the file system on Veritas Volume Manager-backed file systems. The operation of reclamation can be done on a disk, LUN, full file system, or part of a file system using the `vxdisk` and `fsadm` commands, and the `vxfs_ts_reclaim` API.

Veritas Volume Manager

Veritas Volume Manager includes the following changes in 5.0 Release Update 4:

Support for Xen

The VxVM 5.0RU4 release supports Xen for SLES 10 SP3 on x86_64.

See "Supported features" on page 151.

Boot disk recovery

The process to create a failback disk has been simplified. You can now use the VxVM commands to break off the root mirrors and make the mirror disk bootable. The resulting mirror of the boot disk can be used as a failback disk in case of upgrade failure.

See the `vxrootadm`(1M) manual page for more information

Enhancements to the Dynamic Multipathing feature

This release provides a number of enhancements to the Dynamic Multipathing (DMP) features of VxVM. These enhancements simplify administration, and improve display of detailed information about the connected storage.

Improved Dynamic Multipathing device naming

The DMP device naming feature has been enhanced to provide a more consistent and user friendly approach for naming the DMP devices.

The following enhancements apply regardless of the specified naming scheme:

- DMP now enables you to assign customized names for DMP devices. You can specify customized names for individual devices, or you can use a file containing user-defined names to assign multiple names.
- You can specify a DMP device name to commands using the name of any of its subpaths. The output displays the DMP device name assigned.
- In a symmetric cluster, the DDL-generated enclosure-based names for DMP devices are now consistent across all the nodes in the cluster.
- Device names can be made persistent. This is the default for the enclosure-based naming (EBN) scheme.

The following enhancements apply to the EBN naming scheme:

- DDL generates the device name in the format *enclosure_index*. If you specify the `use_avid` argument, the name is generated with the Array Volume ID for the index number to provide a more meaningful name.

Default behavior for I/O throttling

By default, DMP is now configured with no I/O throttling. In previous releases, I/O throttling was set to on. Use the `vxdmpadm setattr` command with the `recoveryoption` keyword to configure I/O throttling for DMP.

Specifying a minimum number of active paths

You can now configure a minimum redundancy level, which is the minimum number of paths for the devices under an enclosure. Use the `redundancy` option of the `vxdmpadm getdmpnode` command to display any devices that have fewer than the minimum number of paths. You can also configure DMP to notify you when the number of active paths falls below the configured minimum.

Enhanced subpaths listing

The `vxdmpadm getsubpaths` command now provides the ability to list all subpaths known to DMP, subpaths of an enclosure, or subpaths through an array port or pwwn. To list the paths through an array port, specify either a combination of enclosure name and array port id, or the array port WWN.

The default listing of the `vxdmpadm getsubpaths` command is sorted by enclosure name, then by DMP node and within that by pathname. The new option `-s` enables

you to sort the output based on path name, DMP node name, enclosure name, or host controller name.

Enhanced I/O statistics

The following enhancements have been made to I/O statistics:

Queued and Erroneous I/O counts

The `vxdmpadm iostat show` command now provides options to display queued I/O counts (`-q` option) and erroneous I/O counts (`-e` option). These options are applicable for DMP node, path and controller.

Filter zero entries

The `vxdmpadm iostat show` command now provides the `-z` option to filter out entities for which all data entries are zero. This option is especially useful in a cluster environment, when many paths are required for failover capabilities, but the paths are not being used for I/O.

Specifying units for statistics data

You can now specify the units in which the statistics data is displayed. The `-u` option accepts `k`, `m` and `g` arguments to display throughput in kilo-, mega-, and giga- system blocks. The `us` argument displays average read/write time in microseconds. By default, the read/write times are displayed in milliseconds up to 2 decimal places. The throughput data is displayed in terms of 'BLOCKS' and the output is scaled, meaning that the small values are displayed in small units and the larger values are displayed in bigger units, keeping significant digits constant. The `bytes` argument to `-u` option can be used to display throughput in exact number of bytes.

Cumulative I/O statistics

The `vxdmpadm iostat` command now has a `groupby` clause to provide cumulative I/O statistics listing per dmpnode, controller, array port id, host-array controller pair and enclosure. If the `groupby` clause is not specified then the statistics are displayed per path.

Miscellaneous improvements to DMP I/O statistics

The following improvements have been made to the DMP I/O statistics output:

The way in which average read/write time is calculated has been corrected.

By default, the average read/write time is displayed in milliseconds up to two places after the decimal point. Use the new option `-u us` to display the average read/write time in microseconds.

The average I/O size is set to 512 bytes.

If the output exceeds 80 columns, then the output is wrapped.

Making DMP restore options persistent

The restore policy, restore interval, and restore period are now persistent across reboot. In addition to being set as options to the `vxdmpadm start restore` command, these attributes can also be set using the `vxdmpadm settune` command. The new tunables are: `dmp_restore_policy`, `dmp_restore_interval`, and `dmp_restore_cycles`.

In addition, there is a new tunable, `dmp_enable_restore`, which enables the path restoration thread to be started.

New log file location for DMP events

There is a new log file location available for tracking DMP events.

Extended device attributes displayed in vxdisk list

The `vxdisk list` command now displays extended device attributes like hardware mirrors for certain arrays.

Display use_all_paths attribute for an enclosure

Display value of `use_all_paths` attribute for an enclosure.

Viewing information about the ASLs installed on the system

The `/usr/lib/vxvm/diag.d/vxcheckasl` command has been enhanced to provide all the information regarding the ASLs (all those installed in the system), the devices (all seen by OS) and all the possible ways in which these ASLs can interact with these devices.

Displaying the count of LUNs in an enclosure

The `vxdmpadm listenclosure` command now displays the count of LUNs in its default output.

Displaying LUN serial number

The `vxdmpadm getdmpnode` command now includes the option `-v` to display the LUN serial number along with other information.

Displaying HBA details

The `vxdmpadm getctlr` output has been enhanced to display HBA vendor details and the Controller ID. For iSCSI devices, the Controller ID is the IQN or IEEE-format based name. For FC devices, the Controller ID is the WWN. Because the WWN is obtained from Event Source Daemon, this field is blank if the Event Source Daemon is not running.

New exclude and include options for the vxdmpadm command

The `vxdmpadm` command now includes `exclude` and `include` commands to suppress or unsuppress devices from VxVM, respectively.

This provides a command line interface for these operations, which previously required user interaction.

New command for reporting DMP node information

The `vxdmpadm` command now includes the `list` command to display information about a DMP node, including all of the attributes that are set for that DMP node. The `vxdmpadm list` command can be used for a specified dmpnode, all dmpnodes, all dmpnodes on the path name or dmpnodename, or all dmpnodes in an enclosure.

Setting attributes for all enclosures

The `vxdmpadm setattr` command now has the `all` option for enclosure, array type and arrayname. The `all` option allows you to set the attributes (`iopolicy`, `failover_policy`, `recoveryoption`) on all the enclosures specified. Also, `vxdmpadm setattr arraytype` *`array_type`* sets the attribute for all array types derived from the given *array_type*.

Support for ALUA JBOD devices

DDL has now improved the support for JBOD devices to include ALUA JBOD devices. DMP now provides immediate basic support for any ALUA compliant array.

Full support still requires an array support library (ASL) for that array. See the Hardware Compatibility List (HCL) for details about supported arrays.

Enhancements to the Cluster Volume Manager

Cluster Volume Manager has been enhanced.

Distributed Volume Recovery

In Cluster Volume Manager (CVM), upon a node crash, the mirror recovery is initiated by the CVM master. Prior to this release, the CVM master also performed all of the recovery I/O. In this release, the CVM master can distribute recovery tasks to other nodes in the cluster. Distributing the recovery tasks is desireable in some situations so that the CVM master can avoid an I/O or CPU bottleneck.

When distribution of volume recovery is turned on, the master distributes recovery tasks in a round-robin fashion to other nodes in the cluster. By default, all of the nodes in the cluster can participate in volume recovery. You can also exclude particular nodes from the volume recovery.

To turn on the distribution of volume recovery, add the keyword `distribute` in the file `/etc/default/vxrecover`. You can also specify the `distribute` keyword when running `vxrecover`.

To disable this feature, remove the `distribute` keyword from the file `/etc/default/vxrecover`. If the feature is not enabled, the master performs the resynchronization tasks locally, as in the previous release.

The distributed volume resynchronization functionality does not depend on the disk group version, and hence works with older version disk groups.

Campus Cluster enhancements

The Campus Cluster feature provides the capability of mirroring volumes across sites, with hosts connected to storage at all sites through a Fibre Channel network.

In this release, the following enhancements have been made to the Campus Cluster feature:

Site tagging of disks or enclosures

The following enhancements to `vxdisk` are related to site tagging:

- Site tagging operations on multiple disks or enclosures are now supported.
- New option to rename a site tag on a disk or enclosure.

Automatic site tagging

The `vxdg settag` command now provides an option for automatic tagging of a site. You can specify that an enclosure is automatically tagged with a particular site name. When you add a LUN belonging to that enclosure to a disk group, the LUN is tagged with the site name specified.

Site renaming

The `vxdg` command has a new `renamesite` option. The `renamesite` option renames the existing site record that is configured on the disk group. The `renamesite` option also associates all of the volume objects to the new site.

Veritas Volume Replicator

Veritas Volume Replicator includes the follwing change to the minimum size for the Storage Replication Log. The minimum size for the Storage Replication Log (SRL) is increased to 130MB in this release.

Veritas File System

Veritas File System includes the following changes in 5.0 Release Update 4:

Unsupported features in this release:

- Data Management API (DMAPI)
- Veritas File System as the root file system

Loopback device support

This release introduces loopback device support. However, loopback device support has the following limitations and restrictions:

- VxFS does not support running a loopback device backed by a file.
- VxFS does not support running Storage Foundation Cluster File System on top of loopback devices.

Mount mntlock and mntunlock options

You can specify the `mntlock` option with the `mount` command, which prevents a file system from being unmounted by an application. This option is useful for an application that does not want the file systems that the application is monitoring to be improperly unmounted by other applications or administrators. Clustering applications, such as Veritas Cluster Server (VCS), are particularly expected to benefit.

The `mntunlock` option of the `vxumount` command reverses the `mntlock` option if you previously locked the file system.

Sendfile support

This release supports the use of the `sendfile`() system call with VxFS. You no longer need to set the Apache configuration option `EnableSendfile Off` when using Apache with VxFS.

Veritas Storage Foundation for Oracle

Veritas Storage Foundation for Oracle includes the following change in 5.0 Release Update 4.

Support for new database

Veritas Storage Foundation for Oracle provides support for the Oracle™ 11g database in the 5.0 Release Update 4 release.

The following features of Veritas Storage Foundation for Databases are supported in the Oracle 11g environment:

- Storage Checkpoints
- Storage mapping
- Database cloning (`clonedb`)
- Database Flashsnap
- Tiered storage for databases (DBDST)
- Quick I/O
- Veritas Extension for Oracle Disk Manager (ODM)

Veritas Storage Foundation for DB2

Veritas Storage Foundation for DB2 includes the following change in 5.0 Release Update 4.

Support for new database

Veritas Storage Foundation for DB2 provides support for the DB2 9.7 database in the 5.0 Release Update 4 release.

The following features of Veritas Storage Foundation for Databases are supported in the DB2 9.7 environment:

- Storage checkpoints

- Storage mapping
- Database cloning (`clonedb`)
- Database Flashsnap
- Tiered storage for databases (DBDST)
- Concurrent I/O

Veritas Storage Foundation Cluster File System

Veritas Storage Foundation Cluster File System includes the following changes in 5.0 Release Update 4:

Number of parallel fsck threads to run during recovery is tunable

In prior releases the number of parallel fsck threads that could be active during recovery was set to 4. In this release the default depends on the number of CPUs in the system, but is tunable within given limits.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide*.

Quick I/O and Veritas ODM

Quick I/O and Veritas Extension for Oracle Disk Manager (ODM) are now enabled by default for Storage Foundation and Storage Foundation Cluster File System. Existing licenses will enable these features after 5.0 Release Update 4 is installed or upgraded on the system. This will not require the addition of a new license.

Chapter 4

Storage Foundation - Software Limitations

This chapter includes the following topics:

Software limitations

Veritas Storage Foundation software limitations

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

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/283859

5.0 RU4 Veritas Storage Foundation software limitations

The following software limitation exists in the 5.0 Release Update 4 release of Veritas Storage Foundation.

unmount can hang when inotify watches are used (1590324)

If `inotify` watches are used, then an unmount can hang in the `vx_softcnt_flush()` call. The hang occurs because inotify watches increment the I-count variable and cause the `v_os_hold` value to remain elevated until the inotify watcher releases the hold.

5.0 RU3 Veritas Storage Foundation software limitations

There are no new additional Veritas Storage Foundation software limitations in the 5.0 RU3 release.

5.0 RU1 Veritas Storage Foundation software limitations

The following are additional software limitations in the 5.0 Release Update 4 release of Veritas Storage Foundation.

Xen platform, Security-enhanced Linux, and database support limitations

The Veritas Storage Foundation 5.0 RU1 release on SuSE Linux Enterprise Server (SLES) 11 is not supported on the Xen platform for Linux.

There is no support for the Xen kernel on SLES11 in this release.

Security-enhanced Linux (SE Linux) is not supported in the Veritas Storage Foundation 5.0 RU1 release.

Because of the database and Xen support limitations on SuSE Linux Enterprise Server (SLES) 11, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation Cluster File System for Oracle RAC is not available on SuSE Linux Enterprise Server (SLES) 11 at the time of this release.

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

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

odmmkfile command must be run as a privileged user

The `odmmkfile` command must be run as a privileged user. This behavior has changed from the previous releases where this command is run by any user.

5.0 MP2 Veritas Storage Foundation software limitations

There are no additional Veritas Storage Foundation software limitations in the 5.0 MP2 release.

Veritas Volume Manager 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.

5.0 RU4 Veritas Volume Manager software limitations

There are no new additional software limitations in this 5.0 RU4 release of Veritas Volume Manager.

5.0 RU3 Veritas Volume Manager software limitations

The following software limitation exists in this 5.0 RU3 release of Veritas Volume Manager.

Limited root disk encapsulation support (1845502)

Root disk encapsulation is not supported in this release.

5.0 MP3 Veritas Volume Manager software limitations

The following incident is the new software limitation for the 5.0 MP3 release of Veritas Volume Manager.

Devices and some paths are not discovered properly with IBM's DS4700 disk array after a reboot (1205369)

On a system with IBM's DS4700 disk array, use fewer than 30 LUNs to ensure that the disk array discovers all of the devices and paths after a reboot.

5.0 MP2 Veritas Volume Manager software limitations

There are no additional Veritas Volume Manager software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas Volume Manager software limitations

There are no additional Veritas Volume Manager software limitations in the 5.0 MP1 release.

Veritas File System software limitations

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

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

5.0 Release Update 4 Veritas File system software limitations

The following are new additional software limitations in this 5.0 Release Update 4 release of Veritas File System.

Support of 32 terabyte file systems

Only Veritas Storage Foundation Enterprise and Veritas Storage Foundation Enterprise HA support file systems that are greater than 32 TB.

5.0 MP3 Veritas File system software limitations

There are no additional Veritas File System software limitations in the 5.0 MP3 release.

5.0 MP2 Veritas File system software limitations

There are no additional Veritas File System software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas File system software limitations

There are no additional Veritas File System software limitations in the 5.0 MP1 release.

Veritas Storage Foundation Cluster File System software limitations

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

http://entsupport.symantec.com/docs/283857

Veritas Storage Foundation Cluster File System 5.0 RU4 software limitations

There are no new additional software limitations in the 5.0 RU4 release of Veritas Storage Foundation Cluster File System.

Veritas Storage Foundation Cluster File System 5.0 RU3 software limitations

There are no new additional software limitations in the 5.0 RU3 release of Veritas Storage Foundation Cluster File System.

Veritas Storage Foundation Cluster File System 5.0 MP3 software limitations

This section lists the software limitations in this release of Veritas Storage Foundation Cluster File System.

Compatibility with previous versions of Veritas File System

A disk layout Version 7 file system created with VxFS 5.0 software will not be accessible if the VxFS 5.0 file system software is removed and the system is reverted to VxFS 4.1.

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

The `VX_CONCURRENT` advisory cannot be set on a file that is actively open by Quick I/O or ODM. A file that has the `VX_CONCURRENT` advisory set may not be concurrently opened by Quick I/O or ODM. Quick I/O and ODM access are not allowed for any files on a file system that is mounted with the `-o cio` mount option.

Consistent distribution and kernel version for Storage Foundation Cluster File System

All the nodes in a SFCFS cluster must be at the same OS version and patch level. In addition, mixing nodes running 32-bit kernel with nodes running 64-bit kernel is not supported with SFCFS.

5.0 MP2 Veritas Storage Foundation Cluster File System software limitations

There are no additional Veritas Storage Foundation Cluster File System software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas Storage Foundation Cluster File System software limitations

There are no additional Veritas Storage Foundation Cluster File System software limitations in the 5.0 MP1 release.

Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

The following sections list software limitations for Veritas Storage Foundation Cluster File System for Oracle RAC.

5.0 RU4 Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

There are no additional software limitations in the 5.0 RU4 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

5.0 RU3 Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

There are no additional software limitations in the 5.0 RU3 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

5.0 MP3 Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

The following are software limitations in the 5.0 MP3 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

I/O Fencing

Veritas Storage Foundation Cluster File System (SFCFS) for Oracle RAC is not licensed to support the I/O Fencing feature. If you use the product installer or installation scripts, you are not given the option during installation to enable I/O Fencing, and the I/O Fencing driver will be disabled. If you install the product manually, use the `/etc/vxfen.d/vxfenmode_disabled` file as a template for the `/etc/vxfenmode` configuration file.

If you subsequently enable I/O Fencing, importing a shared disk group will fail with the following error message:

```
VxVM vxdg ERROR V-5-1-10978 Disk group shared_dg: import failed:
License has expired or is not available for operation
```

To verify the state of the I/O Fencing driver, use the `/sbin/vxfenadm -d` command.

Storage Foundation Cluster File System for Oracle RAC does not support Symantec's implementation of SCSI-3 PGR based I/O fencing and Oracle Clusterware (CRS) is expected to handle any split-brain situations. More information is available at the following URL:

```
http://entsupport.symantec.com/docs/306411
```

Avoiding the serial split brain condition

Veritas Storage Foundation Cluster File System (SFCFS) for Oracle RAC does not include support for the I/O Fencing feature. As a result, SFCFS requires that two heartbeat links be configured during installation. When a node is reduced to a single heartbeat connection, SFCFS can no longer discriminate between the loss of a system and the loss of the final network connection. This "jeopardy" state affects all applications that use the cluster file system mount points. (Jeopardy cannot be avoided as the I/O Fencing feature is not available.)

If a node fails after the jeopardy state has been notified, all the cluster nodes cease to be members of the shared disk group, but the mounted cluster file systems are

not disabled on all the nodes in the cluster. This is fine because Oracle's native fencing will evict the problem node from the cluster and prevent potential data corruption. The surviving cluster nodes will be unaffected and continue to operate.

Oracle's native fencing will evict the problem node from the cluster if either of the following events occur:

- Simultaneous failure of both heartbeat links to a node.
- A node hangs and is unable to respond to heartbeat messages.

5.0 MP2 Veritas Storage Foundation Cluster File System for Oracle RAC software limitations

The following are software limitations in the 5.0 MP2 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

Storage Foundation Management Server is not supported for centralized servers

Storage Foundation Management Server is only supported for managed hosts; it is not supported for centralized servers.

Veritas Storage Foundation for Oracle software limitations

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

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

5.0 RU4 Veritas Storage Foundation for Oracle software limitations

There are no additional Veritas Storage Foundation for Oracle software limitations in the 5.0 RU4 release.

5.0 RU3 Veritas Storage Foundation for Oracle software limitations

There are no additional Veritas Storage Foundation for Oracle software limitations in the 5.0 RU3 release.

Veritas Storage Foundation for DB2 software limitations

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

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

5.0 RU4 Veritas Storage Foundation for DB2 software limitations

There are no additional Veritas Storage Foundation for DB2 software limitations in this 5.0 RU4 release.

5.0 RU3 Veritas Storage Foundation for DB2 software limitations

There are no additional Veritas Storage Foundation for DB2 software limitations in this 5.0 RU3 release.

5.0 MP3 Veritas Storage Foundation for DB2 software limitations

There are no additional Veritas Storage Foundation for DB2 software limitations in the 5.0 MP3 release.

5.0 MP1 Veritas Storage Foundation for DB2 software limitations

The following are additional software limitations in the 5.0 MP1 release of Veritas Storage Foundation for DB2.

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.

Cannot restore if tablespace is converted from Quick I/O to regular file after backup (25272)

If you convert a tablespace from a Quick I/O file to a regular file after backing up the database, you will not be able to restore the tablespace from that backup. For example, if you take a backup of a database that has a DMS tablespace with Quick I/O files as containers, and later convert the Quick I/O files to regular files, restoring the database from that backup will fail.

Workaround

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

Selected utilities require setuid (643964)

Some Veritas Storage Foundation for Databases programs are setuid binaries because they are meant to be run as a database administrator and the APIs used are root access-only Symantec internal APIs. The affected binaries are used mainly for information query purposes.

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

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

Repository hostnames are case-insensitive (851129)

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

Veritas Storage Foundation for Databases software limitations

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

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

5.0 RU3 Veritas Storage Foundation for Databases software limitations

There are no additional Veritas Storage Foundation for Databases software limitations in this 5.0 RU3 release.

5.0 MP3 Veritas Storage Foundation for Databases software limitations

There are no additional Veritas Storage Foundation for Databases software limitations in the 5.0 MP3 release.

5.0 MP2 Veritas Storage Foundation for Databases software limitations

The following are additional software limitations in the 5.0 MP2 release of Veritas Storage Foundation for Databases.

Deep mapping unsupported

Deep mapping on EMC SYMCLI is currently not supported.

Use of buffered I/O

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

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

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

5.0 MP1 Veritas Storage Foundation for Databases software limitations

The following are additional software limitations in the 5.0 MP1 release of Veritas Storage Foundation for Databases.

Some features stop working after a Global Cluster failover (563603)

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

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

Differing locales produces unintelligible characters in GUI (605487)

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

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.

Avoid using UNIX VEA via PC-based UNIX emulators

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

CLI database state changes are delayed in GUI (604685)

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

Workaround

Start or stop the database from the GUI, or do a manual rescan from the GUI after starting or stopping with CLI.

Deep mapping unsupported

Deep mapping on EMC SYMCLI is currently not supported.

Command Line Interface limitations

This section describes command line interface limitations.

The following commands do not support multipartition databases on SMP architectures:

- `db2ed_clonedb`
- `db2ed_vmchecksnap`
- `db2ed_vmclonedb`

Database cloning limitation (285139)

If you clone a database and mount it, ensure that the directory where the mount point resides is owned by the instance owner of the cloned database.

If the directory where the mount point resides is not owned by the instance owner, an error message is displayed when you attempt to remove and unmount the cloned database.

For example:

```
$ db2ed_clonedb -I inst01 -S prod -T clone -c \
    Checkpoint_1105997700  -m /mnt
```

where /mnt is created by root and the owner has been changed to `inst01`.

When you attempt to remove and unmount the clone database, you will get the following error message:

```
$ db2ed_clonedb -T clone -o umount -d
    rm: Unable to remove directory /mnt Permission denied
```

This error message does not affect the functionality of `db2ed_clonedb`. The clone database has been removed and unmounted even when you receive this error message.

To avoid this error, create a directory under `/` as root and change the owner of the directory to the instance owner. Then, specify a mount point under the newly created directory. For example, instead of using the mount point /mnt as in the above example, specify a mount point under `/mnt`, such as `/mnt/clone`:

```
$ db2ed_clonedb -I inst01 -S prod -T clone -c \
    Checkpoint_1105997700  -m /mnt/clone
```

Note: If your mount point is under the directory `/tmp`, you will not encounter this problem.

No support for Intelligent Storage Provisioning

The Standard, Standard HA, Enterprise, and Enterprise HA versions of Veritas Storage Foundation for DB2 do not support Intelligent Storage Provisioning (ISP).

Using VxVM volumes as DB2 DEVICE containers

For this release, Symantec recommends that the VxVM raw device interface be used when configuring a raw device as a DB2 container. Using VxVM raw interface ensures fast and reliable I/O support. Also, VxVM's 5.0 raw device interface supports full Linux native asynchronous I/O.

Make sure the DB2 instance owner has read/write access to the VxVM volume. Also, the DB2 instance owner should have permissions to change the directory to `/dev/vx`. By default, only the root user can access the `/dev/vx` directory.

The following example shows how a DMS tablespace is created on a VxVM volume.

```
CREATE  REGULAR  TABLESPACE USERRAW \
    PAGESIZE 4K  \
    MANAGED BY DATABASE  \
    USING (DEVICE '/dev/vx/dsk/db2dg/db2vol1' 128000)  \
    EXTENTSIZE 32 PREFETCHSIZE 32 \
    BUFFERPOOL IBMDEFAULTBP
```

An alternative method to use a VxVM volume is to go through the Linux raw binding interface. For further information, refer to the man page for raw commands on Linux. Using Linux raw binding is not a recommended method for VxVM.

Deep mapping on hard disk arrays

This release of Veritas Storage Foundation 5.0 for DB2 will not offer the deep mapping feature because partner disk array software on Linux platforms have been unavailable. The storage mapping display will show disk arrays only to the LUN level.

During installation, the `VRTSvail` package is not installed. However, this package is included in this release, so that this issue can be corrected in a future maintenance pack release.

Use of buffered I/O

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

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

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

Veritas recommends the use of Concurrent I/O with DB2. In DB2 UDB version 8.2.2 or above, Concurrent I/O is used with VxFS when DBAs choose "Direct I/O" for a given tablespace. When Direct I/O is specified this way, the DB2 server activates Concurrent I/O semantics in a way that is transparent to users.

See the *Veritas Storage Foundation for DB2 Administrator's Guide* to learn how to use Concurrent I/O with DB2 UDB on Linux, UNIX, and Windows.

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, 643964)

Some Veritas Storage Foundation for Databases programs are setuid binaries because they are meant to be run as a database administrator and the APIs used are root access-only Symantec internal APIs. The affected binaries are used mainly for information query purposes.

For these reasons, the following programs are setuid-enabled in Veritas Storage Foundation for Oracle:

- `/opt/VRTSdbed/.dba/dbed_analyzer`
- `/opt/VRTSdbed/.dba/vxckptplan`
- `/opt/VRTSdbcom/bin/vxstorage_stats`
- `/opt/VRTSdbcom/.dba/vxdbd_start`
- `/opt/VRTSdbcom/.dba/vxckpt_ismounted`

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

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

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

Veritas Storage Foundation for Sybase software limitations

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

http://entsupport.symantec.com/docs/290244

5.0 RU4 Veritas Storage Foundation for Sybase software limitations

The are no additional software limitations in this 5.0 RU4 release of Veritas Storage Foundation for Sybase. However, *Veritas Storage Foundation for Sybase* is supported only for SLES 10 SP3 on ppc.

5.0 RU3 Veritas Storage Foundation for Sybase software limitations

The are no additional software limitations in this 5.0 RU3 release of Veritas Storage Foundation for Sybase.

Veritas Volume Replicator software limitations

Software limitations in the 5.0 release are listed in the *Veritas Volume Replicator 5.0 Release Notes*, which is available at the following URL.

http://entsupport.symantec.com/docs/290244

5.0 RU4 Veritas Volume Replicator software limitations

There are no additional Veritas Volume Replicator software limitations in this 5.0 RU4 release.

5.0 RU3 Veritas Volume Replicator software limitations

There are no additional Veritas Volume Replicator software limitations in this 5.0 RU3 release.

5.0 MP3 Veritas Volume Replicator software limitations

There are no additional Veritas Volume Replicator software limitations in this release.

5.0 MP2 Veritas Volume Replicator software limitations

There are no additional Veritas Volume Replicator software limitations in the 5.0 MP2 release.

5.0 MP1 Veritas Volume Replicator software limitations

There are no Veritas Volume Replicator software limitations in the 5.0 MP1 release.

Chapter 5

Storage Foundation - Known Issues

This chapter includes the following topics:

Known issues

The following are additional known issues for Veritas Storage Foundation.

Veritas Storage Foundation known issues

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

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

Veritas Storage Foundation 5.0 RU4 known issues

The following known issue exists in this 5.0 RU4 release.

32-bit JRE requirement (1870929)

Installation of the 32-bit JRE ibm-java-ppc-jre-6.0-6.0.ppc is required in this 5.0 RU4 release.

DMP device status detection issue (1870300)

Dynamic multipathing cannot detect a re-enabled SLES11 device status; the OS removes stale device(s) after a timeout. When restored, the devices are recreated, but the device name may be different. Without rescandisks, DMP fails to recognize the disk(s) because the device name was changed.

Veritas Storage Foundation 5.0 RU3 known issues

The following known issue exists in this 5.0 RU3 release.

VRTSvcssy package and Sybase Agent installation (1877566)

CPI does not automatically install the VCS Sybase Agent in this RU3 release.

Workaround

Install the VCS Sybase Agent RPM manually.

Note: This issue is fixed in the 5.0 RU4 release.

Veritas Storage Foundation 5.0 MP3 known issues

The following are new additional known issues in this 5.0 MP3 release of Veritas Storage Foundation.

While configuring Storage Foundation Management Server and the Cluster Management Console through the CPI for authentication passwords, some special characters are not passed correctly through the CPI (1245237)

While configuring Storage Foundation Management Server and the Cluster Management Console through the Common Package Interface (CPI) for authentication passwords, some special characters are not correctly passed through the CPI to the nodes, even though these special characters are accepted by authentication.

The following special characters are not correctly passed through the CPI to the nodes:

- \' (single quote)
- \" (double quote)
- \@ (at sign)
- \$ (dollar)
- \\ (slash)
- * (star)

Workaround

There is no workaround for this issue. When entering authentication passwords, do not use any of the special characters listed above.

PATH setting when using product installer to install on remote nodes

There is a problem with some versions of `rsh` and `ssh` that causes the PATH to not get set correctly on remote shell invocations. As a result some of the commands fails to run. Before running the product installer, ensure that the `PATH` is set correctly for the `root` user. For bash shells, it can be set in the `~/.bashrc` or the `~/.bash_profile` file.

To ensure that the PATH has been set properly

- Run the following command for each remote node you want to install and check that the command succeeds.

  ```
  # rsh remote_node lsmod
  ```

Using ODM with Storage Foundation or Storage Foundation Cluster File System

Starting with 5.0MP3, Quick I/O and Veritas Extension for Oracle Disk Manager (ODM) are now enabled by default for Storage Foundation and Storage Foundation Cluster File System.

To use this functionality, you may need to manually install the required packages and patches. Refer to the chart below for your product and install scenario.

Installing ODM

To install ODM to use with Storage Foundation or Storage Foundation Cluster File System

1. Install the ODM rpms
2. If you are installing SFCFS, install the GMS rpms.
3. If you reboot the machine during or after the installation, GMS is automatically loaded. If you do not reboot, you must manually load GMS.

Veritas Storage Foundation 5.0 MP1 known issues

The following known issues exist in the Veritas Storage Foundation 5.0 MP1 release.

Veritas Storage Foundation Basic known issues

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

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

Veritas Storage Foundation Basic 5.0 RU4 known issues

There are no additional known issues in this 5.0 RU4 release of Veritas Storage Foundation.

Veritas Storage Foundation Basic 5.0 RU3 known issues

There are no additional known issues in this 5.0 RU3 release of Veritas Storage Foundation.

Veritas Storage Foundation Basic 5.0 MP3 known issues

There are no new additional known issues in the 5.0 MP3 release of Veritas Storage Foundation.

Veritas Volume Manager known issues

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

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

Veritas Volume Manager 5.0 RU4 known issues

The following additional known issues are described for this 5.0 RU4 release.

vxconfigd daemon issue (1893139)

When using an EMC CX4-240 ALUA iSCSI target on a SLES10 SP3 x86_64 platform in a clustered environment, after connecting the iscsi array through the software initiator, the `vxconfigd` daemon does not start due to a Novell issue. When a second iscsi session is opened and a scandisks operation is performed, it takes approximately 30 minutes to stabilize the connection and many connection errors may take place:

```
ISCSI_ERR_CONN_FAILED (1011).
   connection1:0: detected conn error (1011)
```

Workaround

There is currently no workaround available.

Bugzilla opened with Novell (bugzilla #564524)

DMP detection issue (1870300, 1594277, 1753740)

DMP cannot detect re-enabled OS devices if the device names are changed on a SLES 11 operating system. The parameter `remove_on_dev_loss` of

the `scsi_transport_fc` module is removed in SLES 11, and the OS device files are removed after a device loss with `dev_loss_tmo`. When the device comes back online, the device names may have changed and DMP will then not recognize its status with the restored daemon.

Workaround

For a workaround and/or patch from Novell regarding a fix, refer to the LBN at:

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

Veritas Volume Manager 5.0 RU3 known issues

The following additional known issues are described for this 5.0 RU3 release.

Root disk encapsulation failure (1845502)

Root disk encapsulation is not supported in this release.

LVM volume not converted by vxvmconvert utility (1809789)

Because of changes to the `pvmove` command in the LVM package, `vxvmconvert` cannot convert LVM diskgroups to VxVM diskgroups after LVM version 2.02.32. LVM version 2.02.32 is the last known working version.

There are no known workarounds for this issue."

Restrictions on acceptable root disk layouts for root disk encapsulation (1837396)

There are known restrictions for acceptable root disk layouts for root disk encapsulation. To root encapsulate a system where the swap partition is not the last partition on the rootdisk requires that there is space for a private partition at the end of the disk (typically 32MB). Disks with > 32MB free at the end of the root disk can be encapsulated successfully.

Veritas Volume Manager 5.0 MP3 known issues

The following are new additional known issues in this 5.0 MP3 release of Veritas Storage Foundation.

Limitation of automatic site reattachment feature (1256764)

The site does not reattach automatically in the case where the site storage has disconnected and reconnected to a CVM slave node, but the master node never lost connection to the site storage.

Deport operation on a shared disk group fails (1368377)

With all primary paths inaccessible, the deport operation on a shared disk group fails to clear the PGR keys as the DMP database is not up-to-date. The deport operation succeeds but the PGR keys are not cleared as the DMP database is not updated to reflect the inaccessibility of failed primary paths.

Workaround

Running `vxdisk scandisks` before the DG deport operation triggers DMP reconfiguration which updates the DMP database such that a disk is accessible through active paths.

vxvmconvert fails for LVM volumes created on whole CCISS disks (1253830)

`vxvmconvert` on 5.0 MP3 only works for the LVM volumes created on CCISS device disk partitions, not whole CCISS disks. If a Sistina Physical Volume was created on a whole CCISS disk, LVM and LVM2 volumes created on that physical CCISS disk cannot be converted to VxVM volumes by `vxvmconvert`. The operation will fail and can result in the destruction of LVM partitions. For the Physical Volumes created on CCISS device disk partitions, `vxvmconvert` will succeed.

I/O failures result in the disk failing flag (1205380)

In some DMP failover scenarios, I/O retry causes the disk failing flag to be set, although there is nothing wrong with the disks except for the failing flag.

Workaround

Clear the failing flag using the `vxedit` command.

Manually installing the VRTSvxvm patch requires a reboot (1180992)

After you manually install the `VRTSvxvm` patch, you must reboot the system.

Plex disabled in DS4000 disk array (924680)

This issue has been identified in the 5.0 MP1 related to the DS4000 disk array. If there is a high I/O load to the array, a device inquiry may fail. This failure causes the DMP node to be disabled. When the DMP node is disabled, all I/O to the DMP node fails.

Autotagging can cause the reattach of a site to fail in a Campus Cluster (1470548)

Using the autotagging feature for a Campus Cluster can cause the site attach to fail. VxVM displays an error message such as the following:

```
VxVM vxdg ERROR V-5-1-10128 tagid already assigned to disk
```

Workaround:

Tag disks manually; do not use autotagging.

DMP issues of Veritas Volume Manager

The following are DMP issues in the 5.0 MP1 release of Veritas Volume Manager.

Identification of ATA and SATA disks (862137)

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

Workaround

Disable DMP for ATA and SATA disks.

To disable multipathing for ATA and SATA disks

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

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

2 Run the `vxdiskadm` command and select option 17 (`Prevent multipathing/Suppress devices from VxVM's view`).

3 Select option 7 (`Prevent multipathing of disks by specifying a VID:PID combination`).

4 Enter `ATA:*` as the VID:PID combination.

5 Exit from `vxdiskadm`, and reboot the system.

Handling intermittently failing paths

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

Veritas Volume Manager cluster issues

The following cluster issue exists in this release of Veritas Volume Manager.

Handling intermittently failing paths in a Campus Cluster

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

Veritas File System known issues

The Veritas File System 5.0 Release Update 4 known issues are listed by release in this section. Known issues in the Veritas File System 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes* which is available at the following URL:

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

Veritas File System 5.0 RU4 known issues

There are no new additional known issues for this 5.0 RU4 release.

Veritas File System 5.0 RU3 known issues

The following new additional known issues exist in the 5.0 RU3 release of Veritas File System.

The disk space used by sparse files is higher on SLES11 compared to SLES10 (1833123)

The current mechanism of space allocation for sparse files is based on the memory page size of the underlying operating system. The page size on SLES11 is 64Kb, which makes the disk space used by sparse files on SLES11 higher compared to SLES10 which has a 4Kb page size.

Veritas File System 5.0 MP3 known issues

The following are new additional Veritas File System known issues in this 5.0 MP3 release.

Possible error messages using 32-bit commands on raw devices on 64-bit system (1079725)

On 64-bit systems, using 32-bit commands on raw devices sometimes produces error messages in the system log similar to the following example:

```
ioctl32(mkfs.vxfs:7471): Unknown cmd fd(3) cmd(564f4c02){16}
arg(fff51138) on /dev/sdc2
```

You can safely ignore these messages.

Unmounting a VxFS file system takes long time after an iSCSI dynamic multipathing failover

Unmounting a VxFS file system can take several minutes after an iSCSI dynamic multipathing failover.

Workaround

The following procedure eliminates the long unmount time.

To eliminate the long unmount time

1 After a failover, run the `iostat` command to verify that all I/Os are drained:

```
# iostat
```

2 Unmount the file system or stop VCS:

```
# umount /MyFileSystem
```

Or:

```
# hastop -local
```

fsapadm enforceckpt core dumps if you do not specify a Storage Checkpoint

The `fsapadm` command core dumps if you do not specify a Storage Checkpoint when specifying the `enforceckpt` keyword.

Workaround

The proper usage is as follows:

```
fsapadm enforceckpt [-f strict] mountpoint storage_checkpoint
```

Veritas Storage Foundation Cluster File System known issues

Veritas Storage Foundation Cluster File System known issues in the 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*.

The *Veritas Storage Foundation Cluster File System 5.0 Release Notes* can be viewed at the following URL:

http://entsupport.symantec.com/docs/283857

Veritas Storage Foundation Cluster File System 5.0 RU4 known issues

The following known issues exist in the 5.0 RU4 release of Veritas Storage Foundation Cluster File System.

During a link failure, the MultiNICA agent does not failover the virtual IP to other active devices (1923602)

On SLES10 (SP3) PPC only, the virtual IP configured for a MultiNICA resource type does not fail over to other active devices. This is due to miiagent binary being unable to detect the device state correctly.

Workaround

Rename the existing `miiagent` in `/opt/VRTSvcs/bin/MultiNICA` to **miiagent1**. Create a new shell script by the name miiagent with execute permissions as a replacement. The contents should be as follows:

```
#!/bin/sh echo 99
```

With this change, the modified `miiagent` indicates that it cannot determine the state of the NIC card.

timezone reporting (1885725)

If the timezone of a system is changed on the fly when CVMVolDg resources are already in the ONLINE state (while the cluster is running), the CVMVolDg resources are reported as OFFLINE in the next monitor cycle. In turn, CVMVolDg and CFSMount resources enter a faulted state.

Workaround

Use the `hastop -all` to bring down the cluster before changing the timezone and then restart the cluster with the `hastart` command on all the nodes.

Veritas Storage Foundation Cluster File System 5.0 RU3 known issues

There are no additional software issues in 5.0 RU3 release of Veritas Storage Foundation Cluster File System.

Veritas Storage Foundation Cluster File System 5.0 MP3 known issues

The following are new additional Veritas Storage Foundation Cluster File System known issues in this 5.0 MP3 release.

CFSMount resource may fault

During cluster startup in a SFCFS for Oracle RAC or SFCFS environment, a CFSMount resource may fault even though the underlying CVMVolDg resource becomes online successfully. If the CVMVolDg resource contains layered VxVM volumes, the reason for the fault could be that the CVMVolDg resource went online before all the subvolumes of the layered volume could be started.

Workaround

In order to ensure that a CVMVolDg resource containing layered volumes becomes online only after all the subvolumes are enabled, the CVMVolume attribute in the `main.cf` file should be populated with the names of the layered volumes under that CVMVolDg resource.

See the *Veritas Cluster Server User's Guide*.

Veritas Storage Foundation Cluster File System 5.0 MP1 known issues

The following known issue exists in the 5.0 MP1 release of Veritas Storage Foundation Cluster File System.

Oracle-Linux (840486)

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

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

If you want to uninstall the product, you must uninstall the Veritas Storage Foundation for Oracle packages manually.

To uninstall the Veritas Storage Foundation for Oracle packages

1. Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
2. Stop the repository database and unmount the repository volume.

In a stand-alone configuration:	Stop the database repository:
	`# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb`
	Unmount the database repository:
	`# /opt/VRTSdbcom/config/sfua_rep_mount stop`

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

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

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

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

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

DB2-Linux (840486)

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

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

If you want to uninstall the product, you must uninstall the Veritas Storage Foundation for DB2 packages manually.

To uninstall the Veritas Storage Foundation for DB2 packages

1. Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
2. Stop the repository database and unmount the repository volume.

In a stand-alone configuration:	Stop the database repository:
	`# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb`
	Unmount the database repository:
	`# /opt/VRTSdbcom/config/sfua_rep_mount stop`

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

3. Remove the Veritas Storage Foundation for DB2 packages using the `rpm -e` command.

```
# rpm -e VRTSd2gui-common VRTSdb2ed-common VRTSdbcom-common
VRTSdbdoc
```

If Veritas Storage Foundation for DB2 and Veritas Storage Foundation Cluster File System are installed on the same machine and you are upgrading both products, use the installer to upgrade Veritas Storage Foundation Cluster File System first. Then, use the installer to upgrade Veritas Storage Foundation for DB2.

If the second upgrade fails, remove the Veritas Storage Foundation for DB2 packages as described above, then run the installer to upgrade Veritas Storage Foundation for DB2.

Veritas Storage Foundation Cluster File System for Oracle RAC known issues

Known issues in theStorage Foundation Cluster File System for Oracle RAC 5.0 release are listed in the *Storage Foundation Cluster File System for Oracle RAC 5.0 Release Notes*.

Veritas Storage Foundation Cluster File System for Oracle RAC 5.0 RU4 known issues

The following new additional issue exists in 5.0 RU4 release of Veritas Storage Foundation Cluster File System for Oracle RAC.

timezone reporting (1885725)

If the timezone of a system is changed on the fly when CVMVolDg resources are already in the ONLINE state (while the cluster is running), the CVMVolDg resources are reported as OFFLINE in the next monitor cycle. In turn, CVMVolDg and CFSMount resources enter a faulted state.

Workaround

Use the `hastop -all` to bring down the cluster before changing the timezone and then restart the cluster with the `hastart` command on all the nodes.

Workaround

Use the `hastop -all` to bring down the cluster before changing the timezone and then restart the cluster with the `hastart` command on all the nodes.

Fencing utilities vxfenswap and vxfentsthdw error (1846607)

The `chk_ssh_passwdless()` function in the `vxfenswap` script requires correction when passwordless communication is not set up on the node for which the command is executed. After proper configuration, the `vxfenswap` command works without issues.

Workaround

In order to run fencing utilities `vxfenswap` and `vxfentsthdw`, please ensure that you have ssh/rsh configured for password free logins, and provide configuration for the node on which these commands are executed.

Veritas Cluster Server known issues

The following known issues exist in the 5.0 RU4 release of Veritas Cluster Server.

Cluster state change issue for VCS 5.0 RU4 (1917243)

The system's running state changes to the initializing state for a cluster running SLES10 SPC for ppc during a primary site failure.

Configuring security using the installer results in frozen service groups

After configuring security on a cluster using the product installer (`installvcs` or `installsf`) with the `-security` option, all service groups are frozen.

The same behavior is observed after removing security of a cluster using the product installer.

Workaround

Manually unfreeze each service group using the following steps.

To unfreeze each service group

1 Make the configuration writable

```
# haconf -makerw
```

2 Edit the Frozen attribute of each group and set it to 0

```
# hagrp -modify groupname Frozen 0
```

3 Save and make the configuration readonly

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

CFSMountAgent requirement (1890312)

The `CFSMountAgent` command requires `libstdc++.so.5`. To accommodate this requirement, be sure `compat-libstdc++ rpm (i.e. libstdc++33-32bit-3.3.3-7.8.1.x86_64.rpm)` is installed.

During a link failure, the MultiNICA agent does not failover the virtual IP to other active devices (1923602)

On SLES10 (SP3) PPC only, the virtual IP configured for a MultiNICA resource type does not fail over to other active devices. This is due to miiagent binary being unable to detect the device state correctly.

Workaround

Rename the existing `miiagent` in `/opt/VRTSvcs/bin/MultiNICA` to **miiagent1**. Create a new shell script by the name miiagent with execute permissions as a replacement. The contents should be as follows:

```
#!/bin/sh echo 99
```

With this change, the modified `miiagent` indicates that it cannot determine the state of the NIC card.

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/283839

Veritas Volume Replicator 5.0 RU4 known issues

There are no new additional known issues in this 5.0 RU4 release of Veritas Volume Replicator.

Veritas Volume Replicator 5.0 RU3 known issues

There are no new additional known issues in this 5.0 RU3 release of Veritas Volume Replicator.

Veritas Volume Replicator 5.0 MP3 known issues

The following are new additional known issues in this 5.0 MP3 release of Veritas Volume Replicator.

Mirrors are not synchronized when volumes created using init=active option (1287111)

For volumes created using `init=active` option, the mirrors (plexes) are not synchronized. The `vradmin verifydata` command could incorrectly report differences for such volumes. To rectify this situation, synchronize mirrors (plexes)

and resynchronize the secondary by doing Automatic Synchronization, Full Synchronization, or Difference-based Synchronization.

Veritas Volume Replicator 5.0 MP1 known issues

The following known issue exists in the 5.0 MP1 release of Veritas Volume Replicator.

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.

Veritas Storage Foundation for Oracle known issues

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

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

Veritas Storage Foundation for Oracle 5.0 RU4 known issues

The following known issue exists in this 5.0 RU4 release

VRTSdbms3 startup issue (1923171)

VRTSdbms3 for SFORA does not start the database server during a manual bootup if the DBED repository is not configured. The user must explicitly configure the DBED repository after a manual reboot to start using DBED commands.

Workaround

- Install and configure VRTSdbms3 for SFORA in the same session
- Configure the DBED repository after rebooting (recommended)
- Remove the repository configuration check from the `wrap.sh` file

Veritas Storage Foundation for Oracle 5.0 RU3 known issues

The following new additional known issues exist in the 5.0 RU3 release of Veritas Storage Foundation for Oracle.

Temporary files and `dbed` operations (1840689)

Temporary files must be cleaned up from `/tmp` directory after dbed operations.

Environment variable issue and `dbed` snapshot command (1840686)

The `dbed_vmsnap -o snapshot` fails if t`he ORACLE_SID` environment variable is not set.

Veritas Storage Foundation for Oracle 5.0 MP3 known issues

The following new known issues exist in the 5.0 MP3 release of Veritas Storage Foundation for Oracle.

Datafile rollback in a Dynamic Storage Tiering environment fails to rollback the files (1227083)

In a Dynamic Storage Tiering setup, when you execute the `dbed_ckptrollback` command using the `-F <datafile>` option, the operation may fail and display the following error:

$ dbed_ckptrollback -S $ORACLE_SID -H $ORACLE_HOME -c *Checkpoint_name* -F *datafile_name.dbf*

```
Rolling back files using Checkpoint_xxxxxxxxx ...
VXDBA_PRODUCT vxckptadm WARNING V-81-4565
The following files/directories do not reside on VxFS file systems:
VXDBA_PRODUCT vxckptadm  WARNING V-81-999
datafile_name.dbf
VXDBA_PRODUCT vxckptadm
ERROR V-81-4562 Storage Rollback failed.
SFORA rb.file
ERROR V-81-3046 Could not roll back datafiles.
```

This is only observed in the Dynamic Storage Tiering setup.

Workaround

You must restart the Veritas `vxdbd` daemon using the following procedure.

To restart the Veritas vxdbd daemon.

1 Check the Veritas `vxdbd` daemon status using the following command:

```
# /etc/init.d/vxdbdctrl status
Status of Veritas vxdbd
/opt/VRTSdbcom/bin/vxdbd ping SUCCESS
```

2 Stop the Veritas `vxdbd` daemon using the following command:

```
# /etc/init.d/vxdbdctrl stop
Stopping Veritas vxdbd
```

3 Start the Veritas `vxdbd` daemon using the following command:

```
# /etc/init.d/vxdbdctrl start
Starting Veritas vxdbd
/opt/VRTSdbcom/bin/vxdbd start SUCCESS
```

After restarting the Veritas `vxdbd` daemon, you may attempt the rollback operation again.

VRTSddlpr directory may exist in /opt/ after uninstalling the stack (1315258)

After uninstalling `VRTSddlpr` package, the `/opt/VRTSddlpr` directory may still exist.

Workaround

After uninstalling the `VRTSddlpr` package, you must manually clean up the `/opt/VRTSddlpr` directory.

VRTSdbcom directory may exist in /opt/ after uninstalling the stack (1315247)

After uninstalling the stack, the directory `/opt/VRTSdbcom` may still exist.

Workaround

After uninstalling the stack, you must manually clean up the `/opt/VRTSdbcom` directory .

VRTSdbms3 directory may exist in /opt/ after uninstalling the stack (1315262)

After uninstalling `VRTSdbms3` package, the `/opt/VRTSdbms3` directory may still exist.

Workaround

After uninstalling the VRTSdbms3 package, you must manually clean up the /opt/VRTSdbms3 directory.

VRTSvxms directory exists in /opt/ even after uninstalling the stack (1316534)

After uninstalling the stack, the VRTSvxms directory may still exist.

Workaround

After uninstalling the stack, you must check the VRTSvxms directory and all its subdirectories to see if they are empty.

You may manually clean up the VRTSvxms directory and all its subdirectories only if they are empty.

Reverse Resync not supported if database is created using Oracle Managed Files (1192729)

If an Oracle database is created with Oracle Managed Files (OMF), then reverse_resync operations would fail.

The following errors are displayed:

```
oracle@swlx07:~> /opt/VRTSdbed/bin/dbed_vmsnap -S $ORACLE_SID \
-f sp4 -o

reverse_resync_begin
dbed_vmsnap started at 2007-12-28 12:02:42

SFORA dbed_vmsnap WARNING V-81-5725 After reverse_resync_commit
is performed, you need to recreate the Authentication Password
File using the ORAPWD utility.
SFORA dbed_vmsnap ERROR V-81-4882 An error occurred while
reconfiguring Oracle instance 'sfora'.
SFORA dbed_vmsnap ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.12313/nomount.log.
SFORA dbed_vmsnap ERROR V-81-4918 Database sfora has not been
correctly recovered.
SFORA dbed_vmsnap ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.12313/recover.log.
```

Workaround

The reverse_resync operation for a database created with OMF is not supported in the 5.0 Release Update 4 release.

There is no workaround for this issue.

Veritas Storage Foundation for Oracle 5.0 MP2 known issues

The following are known issues in the 5.0 MP2 release of Veritas Storage Foundation for Oracle.

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_vmclonedb -o umount` does not work.

Workaround

Fix the issue that caused the failover to the other node, and then fallback to the fixed node.

Veritas Storage Foundation for Oracle 5.0 MP1 known issues

The following are known issues in the 5.0 MP1 release of Veritas Storage Foundation for Oracle.

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_vmclonedb -o umount` does not work.

Workaround

Fix the issue that caused the failover to the other node, and then fallback to the fixed node.

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

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

dbed_vmclonedb -p failed to create clonedb with modified pfile (852188)

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

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

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

Workaround

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

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

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

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.

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 will fail when attempting to clone a database using the snapshot, and the message "ORA-01195: online backup of file 1 needs more recovery to be consistent" may appear in the log file.

Workaround

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

Veritas Storage Foundation for DB2 known issues

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

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

Veritas Storage Foundation for DB2 5.0 RU4 known issues

The following new additional known issue exists in the 5.0 RU4 release of Veritas Storage Foundation for DB2.

VRTSdbms3 startup issue (1923171)

VRTSdbms3 for DB2 does not start the database server during a manual bootup if the DBED repository is not configured. The user must explicitly configure the DBED repository after a manual reboot to start using DBED commands.

Workaround

Configure the DBED repository after rebooting.

Veritas Storage Foundation for DB2 5.0 RU3 known issues

There are no new additional known issues in the 5.0 RU3 release of Veritas Storage Foundation for DB2.

Veritas Storage Foundation for DB2 5.0 MP3 known issues

There are no additional known issues in the 5.0 MP3 release of Veritas Storage Foundation for DB2.

Veritas Enterprise Administrator known issues

The following known issues exist in Veritas Enterprise Administrator.

Veritas Enterprise Administrator 5.0 RU4 known issues

There are no new additional known issues in the 5.0 RU4 release of Veritas Enterprise Administrator

Veritas Enterprise Administrator 5.0 RU3 known issues

The following new additional known issue exists in the 5.0 RU3 release of Veritas Enterprise Administrator.

VEA GUI and df -k command show different results (1848089)

The VEA GUI and df -k display different results. The VEA GUI **Properties** tab displays incorrect disk space used on a volume.

No VEA login to Windows 4.x (1849818)

Login to Windows 4.x stack box fails through VEA graphical user interface.The VEA interface uses IBM JRE which encounters problems during the SSL handshakes. Because the VEA interface uses anonymous cipher suites that are not supported by IBM JRE, the VEA interface running on 5.0 RU3 machines cannot be used to connect to the 4.x VEA server (running on Linux, Solaris, AIX, Windows, or HP). There are no workarounds available for this issue.

Veritas Enterprise Administrator 5.0 MP3 known issues

Veritas Enterprise Administrator includes the following additional known issues in 5.0 MP3:

Issue with SNMP settings through Tools (1209747)

This issue is specific to Rule Manager window in the VEA Client which is running on a Linux system. When `Default SNMP Settings` dialog is launched in Rule Manager using menu `Tools->Default SNMP Settings`, the `Hostname` field does not allow any characters to be entered.

Workaround

To resolve this issue

1 Click on the `Port field` in `Default SNMP Settings` dialog.

2 Enter the port number in the `Port field`.

3 Click on `Hostname` field to enter hostname details.

Veritas Storage Foundation Graphical User Interface known issues

Cannot convert mountpoint to volume set using the Java Graphical User Interface (1176531)

When you use the Java Graphical User interface to convert a mountpoint into a volume set, the operation is not successful.

Workaround

You must use the `dbdst_convert` command line interface to convert a mountpoint volume to a volume set. You must not use the Java Graphical User Interface to convert volumes.

Adding two storage classes consecutively through VEA Java Graphical User Interface fails (1231856)

When you use the VEA Java Graphical User Interface to add the first class, the operation is successful. However, when you use the VEA Java Graphical User Interface to add a second class, you may see the following error:

```
SFORA dbdst_admin V-81-6212
Do not add or remove class in a single command.
```

This issue is encountered only with Graphical User Interface and not with the command line interface.

Workaround

You must logout and disconnect from VEA. Then you must start a new session to add a second class through the Graphical User Interface.

Alternatively, you may use the `dbdst_admin` command in the command line interface to add a new class:

```
# /opt/VRTS/bin/dbdst_admin -D DB4 -o addclass=NEWCLASS:"newclass"
```

Some disks may not appear in the VEA GUI (1826607)

Whenever the VxVM disk naming scheme is modified, the VEA GUI needs to recognize the new naming scheme.

Workaround

If the VEA GUI is not showing the new disk names, use the **Actions > Refresh** menu item for proper disk updating.

Disks or volumes may not appear in the VEA GUI (1727003)

After you configure any Storage Foundation product, some disks or volumes may not appear in the VEA GUI.

Workaround

Restart the Storage Agent using the following command:

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

Chapter 6

Documentation

This chapter includes the following topics:

Documentation

Symantec's Veritas Storage Foundation 5.0 Release Update 4 documentation set is available on the Symantec website at the following URL:

http://www.symantec.com/business/support/overview.jsp?pid=15107

Storage Foundation guides

The following manuals, along with the online help, comprise the Veritas Storage Foundation 5.0 MP3 documentation set:

Table 6-1 describes the guides in the Veritas Storage Foundation 5.0 MP3 documentation set.

Table 6-1 Guides in Veritas Storage Foundation documentation set

Guide Title	Filename
Third-party Legal Notices	`3rdpartyattributions.pdf`
Veritas Storage Foundation and High Availability Getting Started Guide	`getting_started.pdf`
Read me end user license agreement documentation	`README_EULA`

Table 6-1 Guides in Veritas Storage Foundation documentation set *(continued)*

Guide Title	Filename
Veritas Storage Foundation read me first	`readme_first.txt`
Veritas Storage Foundation Release Notes (this document)	`sfha_readfirst.pdf`
Veritas Storage Foundation Installation Guide	`sf_install.pdf`
Veritas Enterprise Administrator User's Guide	`vea_users.pdf`
Veritas File System Administrator's Guide	`vxfs_admin.pdf`
Veritas File System Programmer's Reference Guide	`vxfs_ref.pdf`
Veritas Volume Manager Administrator's Guide	`vxvm_admin.pdf`
Veritas Storage Foundation Cluster File System Administrator's Guide	`sfcfs_admin.pdf`

Veritas Volume Replicator documentation

The following Veritas Volume Replicator documentation is available with the Veritas Volume Replicator option:

Table 6-2 describes the Veritas Volume Replicator documentation set.

Table 6-2 Guides in Veritas Volume Replicator documentation set

Guide Title	Filename
Veritas Volume Replicator Administrator's Guide	`vvr_admin.pdf`
Veritas Volume Replicator Planning and Tuning Guide	`vvr_planning.pdf`
Veritas Volume Replicator Web Console Administrator's Guide	`vvr_web_admin.pdf`
Veritas Volume Replicator Advisor User's Guide	`vvr_advisor_users.pdf`
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	`vvr_agents_config.pdf`

Veritas Storage Foundation for Databases documentation Changes

Support for database products for this RU4 release will be announced as soon as possible in the *Late Breaking News Technote* on the Symantec Technical Support website:

http://entsupport.symantec.com/docs/285834

The following documentation is available with the Veritas Storage Foundation for Databases options:

Table 6-3 describes the Veritas Storage Foundation for databases documentation set.

Table 6-3 Guides in Veritas databases documentation set

Guide Title	Filename
Veritas Storage Foundation for DB2 Administrator's Guide	sf_db2_admin.pdf
Veritas Storage Foundation for DB2 Graphical User Interface Guide	sf_db2_gui.pdf
Veritas Storage Foundation for Oracle Administrator's Guide	sf_ora_admin.pdf
Veritas Storage Foundation for Oracle Graphical User Interface Guide	sf_ora_gui.pdf
Veritas Storage Foundation for Sybase Administrator's Guide	sf_syb_admin.pdf

The 5.0 MP3 release introduced changes to the documentation for Veritas Storage Foundation for Databases. The `VRTSordoc` and `VRTSd2doc` packages no longer exist. All database guides are now in one package, `VRTSdbdoc`. New guides have been created for the GUI documentation and the PDF names have changed for existing guides.

Veritas Volume Replicator documentation

The following Veritas Volume Replicator documentation is available with the Veritas Volume Replicator option:

Table 6-2 describes the Veritas Volume Replicator documentation set.

Table 6-4 Guides in Veritas Volume Replicator documentation set

Guide Title	Filename
Veritas Volume Replicator Administrator's Guide	vvr_admin.pdf
Veritas Volume Replicator Planning and Tuning Guide	vvr_planning.pdf
Veritas Volume Replicator Web Console Administrator's Guide	vvr_web_admin.pdf
Veritas Volume Replicator Advisor User's Guide	vvr_advisor_users.pdf
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	vvr_agents_config.pdf

Veritas Cluster Server documentation

The following Veritas Cluster Server documentation is available with all Veritas Storage Foundation HA product suites:

Table 6-5 describes the Veritas Cluster Server documentation set.

Table 6-5 Guides in Veritas Cluster Server documentation set

Guide Title	Filename
Veritas Cluster Server Release Notes	vcs_notes.pdf
Veritas Cluster Server Installation Guide	vcs_install.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents.pdf
VCS Enterprise Agent for Oracle Installation and Configuration Guide	vcs_oracle_install.pdf
VCS Enterprise Agent for DB2 Installation and Configuration Guide	vcs_db2_install.pdf
VCS Enterprise Agent for Sybase Installation and Configuration Guide	vcs_sybase_install.pdf

Manual Pages

The Veritas online manual pages are installed in the `/opt/VRTS/man` directory. This directory can be added to the `MANPATH` environment variable.

If the `MANPATH` environment variable does not include `/opt/VRTS/man`, you can view the desired manual page by entering the following command:

```
# man -M /opt/VRTS/man manual_page_name
```

Note: Installing documentation and online manual pages is optional.

Appendix A

Veritas Storage Foundation 5.0 RU4 for Xen

This appendix includes the following topics:

Supported operating systems for Xen

This release provides support for Veritas Storage Foundation 5.0 Release Update 4 on the Xen platform for Linux.

Supported operating systems are:

- SuSE Linux Enterprise Server 10 (SLES 10) SP3 on x86_64

Supported features

In this release, only Veritas Volume Manager is supported on Xen Dom0.

The following features are provided in this release in the Xen environment:

- Configuration of Veritas Volume Manager volumes in Dom0.
- Export of VxVM volumes to DomU as raw partition devices.

- Support for both SCSI devices and IDE disk devices (`/dev/sdx` and `/dev/hdx`).
- Relayout of a volume in Dom0 that has been exported to DomU is possible. Some reduction in I/O throughput may be experienced during the operation.
- Resizing a volume in Dom0 is not seen in DomU until the domain is rebooted. The file system must be resized independently of the volume. This limitation is imposed by Xen.
- Array Support Libraries (ASLs) and Array Policy Modules (APMs) that are supported in 5.0 Release Update 4.
- Root encapsulation.
- SELinux in permissive or enforcing mode.

Unsupported features

The following features of Veritas Storage Foundation Enterprise product are not supported in this release:

- Veritas File System (VxFS).
- Cluster Volume Manager (CVM).
- Cluster File System (CFS).
- Veritas Cluster Server (VCS).
- The VEA server in DomU.
- Multi-volume file systems, volume sets, and Quality of Storage Service (QoSS)
- Enhanced I/O support in DomU that would usually be obtained by configuring a VxFS file system on a VxVM volume.
- Resizing a VxVM volume in DomU.

Installing the Veritas Software in Dom0

If you plan to install the Veritas Storage Foundation 5.0 Release Update 4 software for the first time and not upgrade an existing system, review the preinstallation instructions.

See the *Veritas Storage Foundation 5.0 Release Update 4 Installation Guide* .

For important release information, review the *Veritas Storage Foundation Release Notes*, and all documents in the `release_notes` directory.

To install the Veritas software in Dom0

1. Obtain the Veritas Storage Foundation 5.0 Release Update 4 software from the software disc or downloaded image.
2. Use the Veritas Volume Manager installation script (`installvm`) instead of the generic Storage Foundation `installer` script to install the packages.

   Enter the following command from the *`platform`*`/volume_manager` directory of the mounted DVD-ROM or downloaded image:

   ```
   # ./installvm -installonly
   ```

   The `-installonly` option is required to perform the installation without configuring the software. Do not create any VxVM volumes before you install the packages that are required to support the Xen environment.
3. To configure the software, run the installation script again, this time specifying the `-configure` option.

   ```
   # ./installvm -configure
   ```

   You are prompted to supply a valid Veritas Storage Foundation product license. Presently, there is no separate product license required for enabling Xen support.

   See the *Veritas Storage Foundation Installation Guide* for help on how to respond to the installation prompts.
4. If the `installvm` script did not reboot the Xen kernel, then manually reboot the Xen kernel now.

Verifying Software Versions

To list the Veritas packages that are installed on your system, enter the following command from the appropriate domain:

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

Uninstalling the Veritas Software from Dom0

To uninstall the Veritas software from Dom0

1. Log on as superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.

3. 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.
4. Stop all VxVM volumes by entering the following command for each disk group:

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

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

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

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