Veritas Storage Foundation™ Installation Guide

Linux

5.0 Maintenance Pack 3

Veritas Storage Foundation™ Installation Guide

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 MP3

Document version: 5.0MP3.0

Legal Notice

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

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/business/support/index.jsp

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

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

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

Maintenance agreement resources

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

Asia-Pacific and Japan	contractsadmin@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

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

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Chapter 1

About Storage Foundation and High-Availability Solutions

This chapter includes the following topics:

Veritas Storage Foundation product suites

The following table lists the Symantec products and optionally licensed features available with each Veritas Storage Foundation product suite.

Table 1-1 Contents of Veritas Storage Foundation products

Storage Foundation version	Products and features
Storage Foundation Basic	Veritas File System Veritas Volume Manager
Storage Foundation Standard	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Optionally licensed features: Veritas Volume Replicator

Table 1-1 Contents of Veritas Storage Foundation products *(continued)*

Storage Foundation version	Products and features
Storage Foundation Standard HA	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Cluster Server Optionally licensed features: Veritas Volume Replicator
Storage Foundation Enterprise	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Storage Checkpoint option Veritas Extension for Oracle Disk Manager option Optionally licensed features: Veritas Volume Replicator
Storage Foundation Enterprise HA	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Storage Checkpoint option Veritas Extension for Oracle Disk Manager option Veritas Cluster Server Optionally licensed features: Veritas Volume Replicator
Storage Foundation for Oracle Standard	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Optionally licensed features: Veritas Volume Replicator

Table 1-1 Contents of Veritas Storage Foundation products *(continued)*

Storage Foundation version	Products and features
Storage Foundation for Oracle Enterprise	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Veritas Database Flashsnap Optionally licensed features: Veritas Volume Replicator
Storage Foundation for Oracle Enterprise HA	Veritas File System Veritas Volume Manager Veritas Cluster Server Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Veritas Database Flashsnap Optionally licensed features: Veritas Volume Replicator
Storage Foundation for DB2 Standard	Veritas File System Veritas Volume Manager Veritas Quick I/O option Optionally licensed features: Veritas Volume Replicator

Table 1-1 Contents of Veritas Storage Foundation products *(continued)*

Storage Foundation version	Products and features
Storage Foundation for DB2 Enterprise	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Storage Checkpoint option Veritas Database Flashsnap Optionally licensed features: Veritas Volume Replicator
Storage Foundation for DB2 Enterprise HA	Veritas File System Veritas Volume Manager Veritas Cluster Server Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Veritas Database Flashsnap Optionally licensed features: Veritas Volume Replicator
Storage Foundation Cluster File System	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Optionally licensed features: Veritas Volume Replicator

Table 1-1 Contents of Veritas Storage Foundation products *(continued)*

Storage Foundation version	Products and features
Storage Foundation Cluster File System HA	Veritas File System Veritas Volume Manager Veritas Cluster Server Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Optionally licensed features: Veritas Volume Replicator
Storage Foundation Cluster File System for Oracle RAC	Veritas File System Veritas Volume Manager Veritas Cluster Server Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Optionally licensed features: Veritas Volume Replicator

About Veritas Enterprise Administrator (VEA)

The Veritas Enterprise Administrator (VEA) is the graphical administrative interface for configuring shared storage devices. VEA simplifies administrative tasks, such as mounting and unmounting file systems, creating and removing storage checkpoints, enabling and disabling change log, and many others. For basic information on running the VEA, refer to *Veritas Enterprise Administrator User's Guide*. For a complete list of administrative tasks and their instructions, see the online help that is available from within the VEA.

Chapter 2

Before you install

This chapter includes the following topics:

About planning for a Storage Foundation installation

This installation guide is designed for system administrators who already have a knowledge of basic UNIX system and network administration. Basic knowledge includes commands such as `tar`, `mkdir`, and simple shell scripting. Also required

is basic familiarity with the specific platform and operating system where Storage Foundation will be installed.

Follow the preinstallation instructions if you are installing one of the Veritas Storage Foundation products by Symantec.

The following Veritas Storage Foundation products by Symantec are installed with these instructions:

- Veritas Storage Foundation Basic, Standard, Standard High Availability (HA), Enterprise, and Enterprise High Availability (HA) Editions

Several component products are bundled with each of these Veritas Storage Foundation products.

See "Veritas Storage Foundation product suites" on page 13.

Veritas Installation Assessment Service

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

https://vias.symantec.com/

Release notes

Read the *Release Notes* for all products included with this product.

The product documentation is available on the web at the following location:

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

Accessing manual pages and documentation directories

Volume Manager manual pages are in the `VRTSvmman` package. Veritas File System manual pages are in the `VRTSfsman` package. Veritas Storage Foundation for DB2 manual pages are in the `VRTSdb2ed` package. Manual pages are installed in the `/opt/VRTS/man` directories. Set the path of your `MANPATH` environment variable to include `/opt/VRTS/man`.

If you are using a shell such as `sh` or `bash`, do the following:

```
$ MANPATH=$MANPATH:/opt/VRTS/man; export MANPATH
```

If you are using a shell such as `csh` or `tcsh`, do the following:

```
% setenv MANPATH ${MANPATH}:/opt/VRTS/man
```

On a Red Hat system, also include the `1m` manual page section in the list defined by your `MANSECT` environment variable.

If you are using a shell such as `sh` or `bash`, do the following:

```
$ MANSECT=$MANSECT:1m; export MANSECT
```

If you are using a shell such as `csh` or `tcsh`, do the following:

```
% setenv MANSECT ${MANSECT}:1m
```

If you use the `man`(1) command to access manual pages, set `LC_ALL=C` in your shell to ensure that they display correctly.

Symantec product licensing

When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased. A single key lets you install the product on the number and type of systems for which you purchased the license. A key may enable the operation of more products than are specified on the certificate; however, you are legally limited to the number of product licenses purchased.

The product installation procedure describes how to activate the key. If you encounter problems while licensing this product, visit the Symantec licensing support website.

The VRTSvlic package enables product licensing. After the VRTSvlic is installed, the following commands and their manual pages are available on the system:

vxlicinst	Installs a license key for a Symantec product
vxlicrep	Displays currently installed licenses
vxlictest	Retrieves features and their descriptions encoded in a license key

Even though other products are included on the enclosed software discs, you can only use the Symantec software products for which you have purchased a license.

Setting environment variables

Most of the commands used in the installation are in the `/sbin` or `/usr/sbin` directory. Add these directories to your `PATH` environment variable as necessary.

Additional variables may be needed to use a Veritas Storage Foundation product after installation.

Symbolic links to all Storage Foundation command executables are installed in the `/opt/VRTS/bin` directory. Add the `/opt/VRTS/bin` directory to your PATH environment variable to access the commands.

Specify `/opt/VRTS/bin` in your PATH after the path to the standard Linux commands. To invoke the VxFS-specific `df`, `fsdb`, `ncheck`, or `umount` commands, type the full path name: `/opt/VRTS/bin/`*`command`*.

If you are not installing an HA product, you can omit `/opt/VRTSvcs/bin`.

Configuring secure shell (ssh) or remote shell before installing products

Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. These utilities must run as `root` on all cluster nodes or remote systems.

Note: When installing on an RHEL5 / OEL5 system with SELinux enabled, only ssh is supported due to RedHat's SELinux policy restrictions.

This section contains an example of how to set up ssh password free communication. The example sets up ssh between a source system (system1) that contains the installation directories, and a target system (system2). This procedure also applies to multiple target systems.

Configuring and enabling ssh

The ssh program enables you to log into and execute commands on a remote system. ssh enables encrypted communications and an authentication process between two untrusted hosts over an insecure network.

In this procedure, you first create a DSA key pair. From the key pair, you append the public key from the source system to the authorized_keys file on the target systems.

Figure 2-1 illustrates this procedure.

Figure 2-1 Creating the DSA key pair and appending it to target systems

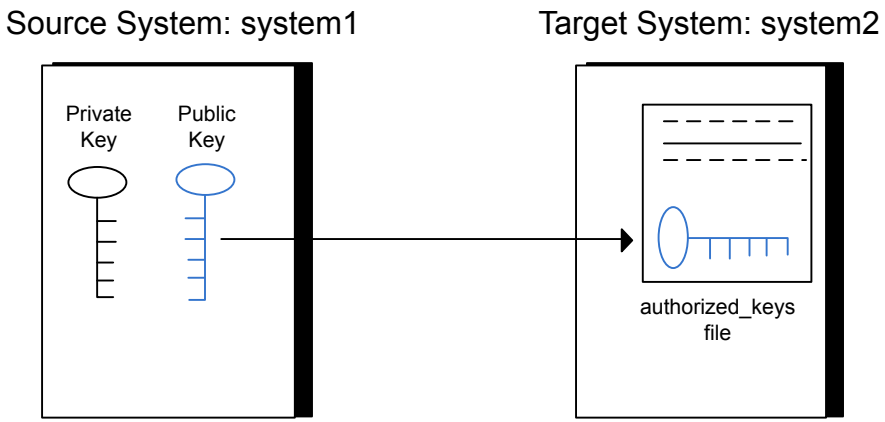

Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

Visit the OpenSSH website that is located at: http://openssh.org to access online manuals and other resources.

To create the DSA key pair

1 On the source system (system1), log in as root, and navigate to the root directory.

```
system1 # cd /
```

2 To generate a DSA key pair on the source system, type the following command:

```
system1 # ssh-keygen -t dsa
```

System output similar to the following is displayed:

```
Generating public/private dsa key pair.
Enter file in which to save the key (/root/.ssh/id_dsa):
```

3 Press Enter to accept the default location of `/root/.ssh/id_dsa`.

4 When the program asks you to enter the passphrase, press the Enter key twice.

```
Enter passphrase (empty for no passphrase):
```

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

5 Output similar to the following lines appears.

```
Your identification has been saved in /root/.ssh/id_dsa.
Your public key has been saved in /root/.ssh/id_dsa.pub.
The key fingerprint is:
1f:00:e0:c2:9b:4e:29:b4:0b:6e:08:f8:50:de:48:d2 root@system1
```

To append the public key from the source system to the authorized_keys file on the target system, using secure file transfer

1 From the source system (system1), move the public key to a temporary file on the target system (system2).

Use the secure file transfer program.

In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

Use the following command for secure file transfer:

```
system1 # sftp system2
```

If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

```
Connecting to system2 ...
The authenticity of host 'system2 (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
Are you sure you want to continue connecting (yes/no)?
```

2 Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'system2,10.182.00.00'
(DSA) to the list of known hosts.
root@system2 password:
```

3 Enter the root password of system2.

4 At the `sftp` prompt, type the following command:

```
sftp> put /root/.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /root/.ssh/id_dsa.pub to /root/id_dsa.pub
```

5 To quit the SFTP session, type the following command:

```
sftp> quit
```

6 Add the `id_dsa.pub` keys to the `authorized_keys` file on the target system. To begin the `ssh` session on the target system (system2 in this example), type the following command on system1:

```
system1 # ssh system2
```

Enter the root password of system2 at the prompt:

```
password:
```

Type the following commands on system2:

```
system2 # cat /root/id_dsa.pub >> /root/.ssh/authorized_keys
system2 # rm  /root/id_dsa.pub
```

7 When you install from a source system that is also an installation target, also add the local system `id_dsa.pub` key to the local `authorized_keys` file. The installation can fail if the installation source system is not authenticated.

To add the local system `id_dsa.pub` key to the local `authorized_keys` file, enter the following command:

```
system1 # cat /root/.ssh/id_dsa.pub >> /root/.ssh/authorized_keys
```

8 Run the following commands on the source installation system. If your ssh session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and make the key globally available for the user `root`:

```
system1 # exec /usr/bin/ssh-agent $SHELL
system1 # ssh-add

  Identity added: /root/.ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

To verify that you can connect to a target system

1 On the source system (system1), type the following command:

```
system1 # ssh -l root system2 uname -a
```

where system2 is the name of the target system.

2 The command should execute from the source system (system1) to the target system (system2) without the system requesting a passphrase or password.

3 Repeat this procedure for each target system.

Restarting ssh

After you complete this procedure, ssh can be restarted in any of the following scenarios:

- After a terminal session is closed
- After a new terminal session is opened
- After a system is restarted
- After too much time has elapsed, to refresh ssh

To restart ssh

1 On the source installation system (system1), bring the private key into the shell environment.

```
system1 # exec /usr/bin/ssh-agent $SHELL
```

2 Make the key globally available for the user `root`

```
system1 # ssh-add
```

Enabling rsh for Linux

The following section describes how to enable remote shell.

Note: When installing on an RHEL5 / OEL5 system with SELinux enabled, rsh is not supported due to RedHat's SELinux policy restrictions.

Veritas recommends configuring a secure shell environment for Veritas product installations.

See "Configuring and enabling ssh" on page 22.

See the operating system documentation for more information on configuring remote shell.

To enable rsh

1 To ensure that the `rsh` and `rsh-server` packages are installed, type the following command:

```
# rpm -qa | grep rsh
rsh-server-0.17-25.3
rsh-0.17-25.3
```

If it is not already in the file, type the following command to append the line "`rsh`" to the `/etc/securetty` file:

```
# echo "rsh" >> /etc/securetty
```

2 In the `/etc/pam.d/rsh` file, change the "`auth`" type from "`required`" to "`sufficient`" for "`pam_rhosts_auth.so`":

```
auth     sufficient         pam_rhosts_auth.so
```

3 To enable the rsh server, type the following command:

```
# chkconfig rsh on
```

4 Modify the `.rhosts` file. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system. This file also contains the name of a user having access to the local system. For example, if the root user must remotely access `system1` from `system2`, add an entry for `system2.companyname.com` to the `.rhosts` file on `system1` by typing the following command:

```
# echo "system2.companyname.com" >> $HOME/.rhosts
```

5 Install the Veritas product.

6 Remove the "`rsh`" entry in the `/etc/securetty` file.

7 Disable the rsh server by typing the following command:

```
# chkconfig rsh off
```

8 After you complete an installation procedure, delete the `.rhosts` file from each user's `$HOME` directory to ensure security:

```
# rm -f $HOME/.rhosts
```

Prerequisites for Storage Foundation Cluster File System

Each cluster node must be connected to the public network and each must have a unique host name by which it can be addressed on the public network. The local node from which you install does not have to be part of the cluster.

Provide the following information when installing the SFCFS:

- The cluster name, beginning with a letter (a-z, A-Z).
- A unique ID from 0-65535 for the cluster. Within the public subnet, a new cluster using a duplicate cluster ID can cause existing clusters to fail.
- The host names of the cluster nodes.
- The device names of the network interface cards (NICs) used for the private networks among nodes.

- Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities as root on all cluster nodes or remote systems.
- Symantec recommends configuring the cluster with I/O fencing enabled. I/O fencing requires shared devices to support SCSI-3 Persistent Reservations (PR). Enabling I/O fencing prevents data corruption caused by a split brain scenario.
  The Storage Foundation Cluster File System is supported without I/O fencing enabled. However, without I/O fencing enabled, split brain scenarios can result in data corruption.

Hardware overview and requirements for Storage Foundation Cluster File System

VxFS cluster functionality runs optimally on a Fibre Channel fabric. Fibre Channel technology provides the fastest, most reliable, and highest bandwidth connectivity currently available. By employing Fibre Channel technology, SFCFS can be used in conjunction with the latest Veritas Storage Area Network (SAN) applications to provide a complete data storage and retrieval solution.

Figure 2-2 shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

Figure 2-2 Four Node SFCFS Cluster Built on Fibre Channel Fabric

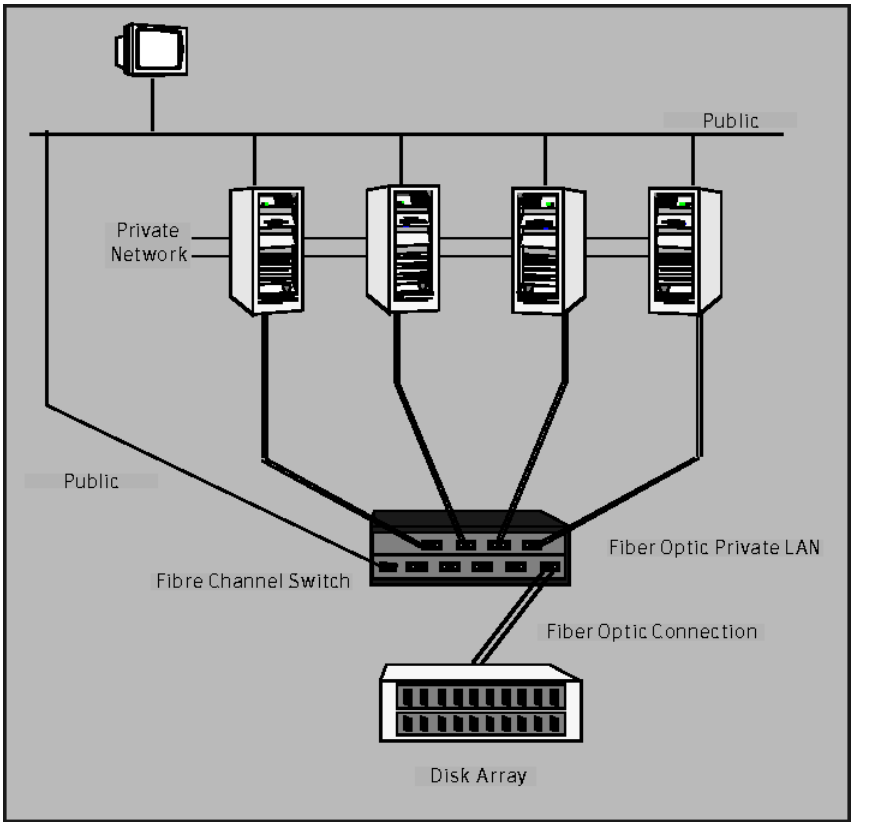

Shared storage

Shared storage can be one or more shared disks or a disk array connected either directly to the nodes of the cluster or through a Fibre Channel Switch. Nodes can also have non-shared or local devices on a local I/O channel. It is advisable to have `/`, `/usr`, `/var` and other system partitions on local devices.

Fibre Channel switch

Each node in the cluster must have a Fibre Channel I/O channel to access shared storage devices. The primary component of the Fibre Channel fabric is the Fibre Channel switch.

Cluster platforms

There are several hardware platforms that can function as nodes in a Storage Foundation Cluster File System (SF CFS) cluster.

See the *Veritas Storage Foundation Release Notes*.

Note: For a cluster to work correctly, all nodes must have the same time. If you are not running the Network Time Protocol (NTP) daemon, make sure the time on all the systems comprising your cluster is synchronized.

Preinstallation or upgrade planning for Veritas Volume Replicator

Before installing or upgrading VVR:

- Confirm that your system has enough free disk space to install VVR.
- Make sure you have root permissions. You must have root permissions to perform the install and upgrade procedures.

The following related documents are available on the documentation disc:

Veritas Volume Replicator Planning and Tuning Guide	Provides detailed explanation of VVR tunables
Veritas Volume Replicator Administrator's Guide	Describes how to change tunable values

See the *Getting Started Guide* for more information on the documentation disc.

VEA installation planning

The Veritas Enterprise Administrator (VEA) GUI consists of several packages. Follow these planning guidelines for VEA installation:

- The server packages must be installed on the hosts on which VVR is installed, not the client. These include the following:
  - To use the VVR functionality in VEA, the Veritas Volume Replicator Management Services Provider package, `VRTSvrpro`, must be installed on all hosts in the Replicated Data Set (RDS).
  - For `VRTSvrpro` to function, the Veritas Volume Manager Management Services Provider package, `VRTSvmpro`, must be installed on your system.

- To use the functionality for receiving SNMP notifications and email notifications, the Veritas Action Agent package, `VRTSaa` must be installed.

- To use the VEA client on a machine other than the machine to be administered, install the VEA client on the machine where the client will run.
- To use the VEA client on a machine, the `VRTSobgui` package must be installed.

Planning an upgrade from the previous VVR version

If you plan to upgrade VVR from the previous VVR version, you can upgrade VVR with reduced application downtime by upgrading the hosts at separate times. While the Primary is being upgraded, the application can be migrated to the Secondary, thus reducing downtime. The replication between the (upgraded) Primary and the Secondary, which have different versions of VVR, will still continue. This feature facilitates high availability even when the VVR upgrade is not complete on both the nodes. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS.

VVR supports replicating data between VVR 5.0MP3 and VVR 4.1 or later.

Replicating between versions is intended to remove the restriction of upgrading the Primary and Secondary at the same time. VVR can continue to replicate an existing RDS with RVGs on the systems that you want to upgrade. When the Primary and Secondary are at different versions, VVR does not support changing the configuration with the `vradmin` command or creating a new RDS.

Note: When replicating between versions of VVR, avoid using commands associated with new features. The earlier version may not support new features and problems could occur.

If you do not need to upgrade all the hosts in the RDS simultaneously, you can use replication between versions after you upgrade one host. You can then upgrade the other hosts in the RDS later at your convenience.

Note: If you have a cluster setup, you must upgrade all the nodes in the cluster at the same time.

See "Upgrading VVR when VCS agents are configured" on page 133.

Database requirements

Table 2-1 identifies supported database and Linux combinations if you plan to use Veritas Storage Foundation for Oracle.

In this release, Veritas Storage Foundation for Oracle is also supported on Oracle Enterprise Linux 4.5.

Table 2-1 Supported Linux platforms for Storage Foundation for Oracle

OS, Platform, Version	OEL4	OEL5	RHEL4	RHEL5	SLES9	SLES10
Oracle9i R2	X	X	Full Support	X	Full Support	X
Oracle10g R1	X	X	Full Support	X	Full Support	X
Oracle10g R2	Full Support	Full Support	Full Support	Full Support	Full Support	Full Support
Oracle11g	Full Support	Full Support	Full Support	Full Support	X	Full Support *

Note: For Oracle 11gR1, there is full support for Storage Foundation for Oracle on SLES 10 in this release.

Table 2-2 identifies supported database and Linux combinations if you plan to use Veritas Storage Foundation for DB2.

Table 2-2 Supported Linux platforms for Storage Foundation for DB2

OS, Platform, Version	OEL4	OEL5	RHEL4	RHEL5	SLES9	SLES10
DB2 8.1	X	X	Full Support	Full Support	Full Support	Full Support
DB2 8.2	X	X	Full Support	Full Support	Full Support	Full Support
DB2 9.1	X	X	Full Support	Full Support	Full Support	Full Support
DB2 9.5	X	X	Full Support	Full Support	Full Support	Full Support

About centralized management

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. SFM is not available on the Storage Foundation and High Availability Solutions release and must be obtained separately.

See "Downloading Storage Foundation Manager" on page 35.

If you plan to use Storage Foundation Manager, configure the Storage Foundation products to use centralized management. Several prerequisites are necessary before you configure the system as a Storage Foundation Manager managed host. You must install and configure Storage Foundation Manager and the Authentication Broker before installing Storage Foundation.

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

If the prerequisites are met, you can set up centralized management while you are installing the Storage Foundation product using the common product installer. Select the option to enable centralized management.

If you do not plan to use centralized management, configure the system to be a stand-alone host.

Storage Foundation products can also be installed on a stand-alone host, and converted to a managed host later.

See the *Storage Foundation Manager Administrator's Guide* for more information.

Downloading the Storage Foundation and High Availability software

One method of obtaining the Storage Foundation and High Availability software is to download it to your local system from the Symantec Web site.

If you download a stand-alone Veritas product, the single product download files do not contain the general product installer. Use the installation script for the specific product to install the product.

See "About installation scripts" on page 203.

To download the software

1. Verify that you have enough space on your filesystem to store the downloaded software.

   The estimated space that is needed for download, gunzip, and tar extract is 5 GB.

   If you plan to install the software on the same system, make sure that you also have enough space for the installed software.

   See "Disk space requirements" on page 41.

2. To see the space available, you can use the `df` command with the name of the local file system where you intend to download the software.

   ```
   # df -k filesystem
   ```

3. Download the software, specifying the file system with sufficient space for the file.

Downloading Storage Foundation Manager

Storage Foundation Manager by Symantec (SF Manager) gives you a single, centralized management console for the Veritas Storage Foundation products. You can use it to monitor, visualize, and manage storage resources and generate reports about those resources.

SF Manager is a free license add-on to Veritas Storage Foundation. You can download SF Manager from the following location:

http://www.symantec.com/sfm

Chapter 3

System requirements

This chapter includes the following topics:

Software and hardware requirements

For information on hardware requirements, see the hardware compatibility list. The hardware compatibility list (HCL) is available at:

http://entsupport.symantec.com/docs/283161

For information on specific HA setup requirements, see the *Veritas Cluster Server Installation Guide*.

Supported Linux operating systems

This section lists the supported Linux operating systems for this release of Veritas Storage Foundation products, including Veritas Storage Foundation Cluster File System.

Read the Technical Support TechNote for the latest information on updates, patches, and software issues regarding this release.

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

The Veritas 5.0MP3 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 4 (RHEL 4) with Update 3 (2.6.9-34 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64).
- Red Hat Enterprise Linux 5 (RHEL 5) with Update 1 (2.6.18-53.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 9 (SLES 9) with SP3 (2.6.5-7.244 kernel or later) or SP4 (2.6.5-7.308 kernel or later) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 10 (SLES 10) with SP1 (2.6.16.46-0.12 kernel or later) or SP2 (2.6.16.60-0.21 kernel or later) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux 4.4 or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux 5.1 or later on AMD Opteron or Intel Xeon EM64T (x86_64)

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

Symantec supports only Oracle, Red Hat, and 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 Red Hat erratas and updates and SUSE service packs is available in the following TechNote. Read this TechNote before you install Symantec products.

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

The Veritas 5.0MP3 release is also supported on the Xen platform for Linux, with some restrictions.

See the Appendix, Veritas Storage Foundation 5.0 MP3 for Xen, in the *Storage Foundation Release Notes* for more information.

For information about the use of this product in a VMware Environment, refer to http://entsupport.symantec.com/docs/289033

Persistent network interface names on SUSE clusters

On SUSE systems, network interfaces can change their names following a reboot. To configure persistent network interface names, add a `PERSISTENT_NAME=eth`*`X`* entry to the `/etc/sysconfig/network/ifcfg-eth-id-`*`mac_address`* network interface configuration file for each interface on each node of the cluster, where *X* and *mac_address* correspond to the interface number and MAC address.

Storage Foundation Cluster File System node requirements

All nodes in a Cluster File System must have the same operating system version and update level.

Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Storage Foundation Cluster File System memory requirements

2 GB of memory is required.

Storage Foundation supported DB2 versions

DB2 9.1 and DB2 8.2, also referred to as DB2 8.1 with FixPak 7, or later is supported on the Linux operating systems listed above.

DB2 9.5 support for Concurrent I/O is provided in the 5.0MP3 release. There is no DB2 9.5 support for Quick I/O in this release.

You can also refer to the *5.0MP3 Veritas Storage Foundation Installation Guide* for detailed information on supported Linux platforms versions for Storage Foundation for DB2.

Storage Foundation supported Oracle versions

Oracle versions 9.2.0.6, 10g, 10gR2 and 11g are all supported on the Linux operating system.

For Oracle 11gR1, there is full support for Storage Foundation for Oracle on SLES 10 in this release.

You can also refer to the *5.0MP3 Veritas Storage Foundation Installation Guide* for detailed information on supported Linux platforms versions for Storage Foundation for Oracle.

Storage Foundation Cluster File System supported Oracle versions

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

Mandatory patch required for Oracle Bug 4130116

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

VxVM licenses

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

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

Table 3-1 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.

Disk space requirements

Before installing any of the Veritas Storage Foundation products, confirm that your system has enough free disk space.

Use the "Perform a Preinstallation Check" (P) menu or the `-precheck` option of the product installer to determine whether there is sufficient space.

```
# ./installer -precheck
```

Required disk space

Table 3-2 shows the approximate disk space used by the Storage Foundation products for all (both the required and optional) packages:

Table 3-2 Required disk space

	Minimum space required (without optional packages)	Maximum space required (including all packages)
Storage Foundation Standard or Enterprise	491MB	625MB
Storage Foundation Standard HA or Enterprise HA	610MB	798MB

Disk space requirements for Veritas Volume Replicator

Table 3-3 shows the approximate disk space used by VVR for the required and optional packages.

Table 3-3 Approximate disk space use for VVR

English	/root	/opt	/usr	/var
Required Packages	38 MB	391 MB	28 MB	0.02 MB
Optional Packages	0 MB	132 MB	0 MB	0 MB
All Packages	38 MB	523 MB	28 MB	0.02 MB

Disk space requirements for Storage Foundation Cluster File System

Table 3-4 shows the approximate disk space used by SFCFS for the required and optional packages.

Table 3-4 Approximate disk space use for SFCFS

English	/root	/opt	/usr	/var
Required Packages	58 MB	634 MB	50 MB	0.02 MB

Table 3-4 Approximate disk space use for SFCFS *(continued)*

English	/root	/opt	/usr	/var
All Packages	58 MB	689 MB	50 MB	0.68 MB

Disk space requirements for Storage Foundation for Oracle

Table 3-5 shows the approximate disk space used by SFORA for the required and optional packages.

Table 3-5 Approximate disk space use for SFORA

English	/root	/opt	/usr	/var
Required Packages	56 MB	707 MB	50 MB	0.02 MB
All Packages	57 MB	762 MB	50 MB	0.68 MB

Chapter 4

Installing Storage Foundation using the common product installer

This chapter includes the following topics:

Installation quick reference

The product installer displays a menu that simplifies the selection of installation and upgrade options. It is the recommended installation method. Select a product to install or upgrade from the menu to invoke that product's installation script.

Table 4-1 provides a quick overview of a stand-alone installation using the product installer.

Table 4-1 Installation overview

Installation task	For more information, refer to the following section:
Obtain product licenses.	See "Symantec product licensing" on page 21.
Download the software, or insert the product DVD.	See "Downloading the Storage Foundation and High Availability software" on page 34. See "Mounting a software disc" on page 46.
Set environment variables.	See "Setting environment variables" on page 22.
Configure the secure shell (SSH) on all nodes.	See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.
Verify that hardware, software, and operating system requirements are met.	See "Software and hardware requirements" on page 37.
Check that sufficient disk space is available.	See "Disk space requirements" on page 41.
Use the installer to install the products.	See "Installing and configuring Storage Foundation using the common product installer" on page 48.

Mounting a software disc

Veritas software is provided on a DVD format disc. If you have the media kit, then get the software disc from the media kit.

To mount the software disc

1. Log in as superuser.
2. Place the Veritas software disc into a DVD drive connected to your system.
3. Insert the disc and type the following command:

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

4. Change to the appropriate distribution directory and product subdirectory to view the product release notes and installation guides, or install the products.

About the common product installer

The product installer is the recommended method to license and install the product. 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.

You can use the product installer to install Veritas Storage Foundation, including the database utilities for DB2 or Oracle.

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.

See "Installation script options" on page 204.

Install Veritas products in one of the following ways:

- For new installations or upgrades from releases prior to 5.0:
  - Use the product installer, the `installer` script
    The product installer displays a menu that simplifies the selection of installation and upgrade options. It is the recommended installation method. Selecting a product to install or upgrade from the menu invokes that product's installation script.
    See "Installing and configuring Storage Foundation using the common product installer" on page 48.
  - Use the product's installation script, such as `installsf` or `installsfcfs`.
    See "About installation scripts" on page 203.
- For upgrades from 5.0 releases, including maintenance packs and rolling patches:
  Use the `installmp` script
  See "Upgrading Storage Foundation software from 5.0 to 5.0MP3 using the product installer" on page 112.

Installing and configuring Storage Foundation using the common product installer

The Veritas product installer is the recommended method to license and install Storage Foundation.

The following sample procedure is based on the installation of Storage Foundation on a single system.

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.

   See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

2. Load and mount the software disc.

   See "Mounting a software disc" on page 46.

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 and Veritas Storage Foundation for DB2 (not available on Oracle Enterprise Linux 4) can also be installed using this procedure. Select the number corresponding to one of those products, if desired.

   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
```

12 Choose whether you want to use enclosure-based naming. 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
```

13 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.

See "Setting the default disk group" on page 146.

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
```

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

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

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

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

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

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

17 The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

18 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-****
```

Installing and configuring Storage Foundation and High Availability Solutions using the common product installer

The following sample procedure is based on the installation of a Storage Foundation Enterprise High Availability (SF/HA) cluster with two nodes: "host1" and "host2."

To install Storage Foundation and High Availability products

1 To install on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

2 Load and mount the software disc.

See "Mounting a software disc" on page 46.

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 systems, if you use the ssh utilities:

```
# ./installer
```

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

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

The sample installation assumes that ssh is used.

5 Enter `I` to install and press Return.

6 When the list of available products is displayed, select Veritas Storage Foundation (SF), enter the corresponding number, and press Return.

With a Veritas Storage Foundation HA license, the high availability cluster components are also installed for this menu selection.

Veritas Storage Foundation for Oracle and Veritas Storage Foundation for DB2 (not available on Oracle Enterprise Linux 4) can also be installed using this procedure. Select the number corresponding to one of those products, if desired.

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" and "host2") 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 host2
```

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

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

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

9 Enter the product license information.

```
Enter a SF license key for
        host1: [?] XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X
  XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered
        on host1
  Do you want to enter another license key for host1?
       [y,n,q,?]  (n) n

  Enter a SF license key for
       host2:  [?] XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X
       XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered
       on host2
  Do you want to enter another license key for host2? [y,n,q,?]
       (n) n
```

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

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

10 You can choose to either install only 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 - 610 MB required
 2) Install all Veritas Storage Foundation rpms - 798 MB required

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

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

11 Configure Storage Foundation and High Availability (SF and VCS) when prompted.

```
Are you ready to configure SF? [y,n,q] (y) y
Do you want to configure VCS on these systems at this time? [y,n,q] (y) y
```

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

All systems are configured to create one cluster.

12 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
```

13 The installer discovers the network interfaces (NICs) available on the first system and reports them:

```
Discovering NICs on host1 ... discovered eth0 eth1 eth2 eth3
```

14 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`, make sure that the same NICs are available on each system; the installer may not verify this. The NICs should also be the same speed on all systems for the heartbeat links to function properly.

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.

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

16 When prompted to configure the product to use Symantec Security Services, enter `n`, unless a Root Broker has already been set up.

Warning: Before configuring a cluster to operate using Symantec Security Services, another system must already have Symantec Security Services installed and be operating as a Root Broker. Refer to the *Veritas Cluster Server Installation Guide* for more information about configuring a Symantec Product Authentication Service Root Broker.

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

17 When prompted to configure SMTP notification, enter `n` or `y` to configure. To configure SNMP notification, enter the following information. You can then confirm that it is correct, or enter it again.

```
Do you want to configure SMTP notification? [y,n,q] (y) y
Active NIC devices discovered on host1: 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
(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@mycompany.com  [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
```

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

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

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

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

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

21 Choose whether you want to use enclosure-based naming. 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
```

22 You have the option of specifying the name of a default 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.

See "Setting the default disk group" on page 146.

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
```

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

```
Will you specify one disk group name for all eligible
systems? [y,n,q,?] (y) y

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

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

25 You may be prompted to verify the fully qualified hostnames of the systems. Verify them and press Return to continue.

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

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

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

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

27 The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

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

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

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

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

29 If you installed Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, create a new repository database.

See "Creating and configuring the repository database for DB2 and Oracle" on page 82.

30 Reboot the system (or systems).

31 After the installation completes, you can view additional information about VCS.

```
The README.1st file has more information about VCS. Read it Now? [y,n,q] (y
```

Installing and configuring 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.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

2 Load and mount the software disc.

See "Mounting a software disc" on page 46.

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.

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

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.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

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 system02?
   [y,n,q] (n) n
```

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

```
Select the RPMs to be installed on all systems?
   [1-3,q,?] (3) 1
```

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

12 Configure Storage Foundation Cluster File System (SFCFS) when prompted.

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

13 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.

Enter `y` or `n` for configuring I/O Fencing in enabled mode.

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

See the *Storage Foundation Cluster File System Administrator's Guide* for more information.

14 Configure the cluster. 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
```

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

```
Discovering NICs on host1 ... discovered eth0 eth1 eth2 eth3
```

16 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.

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

18 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
```

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

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

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

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

21 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
```

22 You have the option of specifying the name of a default 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.

See "Setting the default disk group" on page 146.

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
```

23 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
```

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

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

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

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

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

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

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

See the *Storage Foundation Manager Installation Guide* for details.

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

28 The Veritas Storage Foundation Cluster File System software is installed, verified, and configured.

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

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

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

29 Reboot the system (or systems).

30 If you do not plan to use SFCFS file systems to store the Oracle database or binaries, you have completed the SFCFS installation and configuration.

31 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.

Installing and configuring Veritas Volume Replicator using the common product installer

The Veritas software disc provides a product installer, which is the recommended method to license and install Veritas Volume Replicator (VVR).

To install VVR using the product installer

1. To install on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.

   See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

2. Load and mount the software disc.

   See "Mounting a software disc" on page 46.

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 systems, if you use the ssh utilities:

   ```
   # ./installer
   ```

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

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

   The sample installation assumes that ssh is used.

5. Enter `I` to install and press Return.

6. When the list of available products is displayed, select Veritas Volume Replicator, enter the corresponding number, and press Return.

   To install Veritas Storage Foundation for Oracle, do not select the "Storage Foundation for Oracle RAC packages" option unless you have the correct license or setup.

7. 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
     install SF: system01 system02
   ```

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

If the installer hangs or asks for a login password, stop the installer and run it again after setting up ssh or rsh.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

9 Enter the product license information.

```
Enter a VVR license key for system01: [?] XXXX-XXXX-XXXX-XXXX-XXXX-XXXX
  XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered
        on system01
  Do you want to enter another license key for system01?
       [y,n,q,?]  (n) n
  Enter a VVR license key for
       system02:  [?] XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X
       XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X successfully registered
       on system02
  Do you want to enter another license key for system02? [y,n,q,?]
       (n) n
```

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

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

10 If you have multiple Veritas products, we recommend using the option to install Storage Foundation Enterprise (which includes VVR) rather than installing each product individually. This option ensures that installation steps are done in the proper order and interdependencies are met.

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:

```
Additional rpms are typically installed to simplify future upgrades.

   1)  Required Veritas Volume Replicator rpms - 574 MB required
   2)  All Veritas Volume Replicator rpms - 580 MB required
   3)  Storage Foundation Enterprise rpms - 617 MB required

Select the rpms to be installed on all systems? [1-3,q] (3) 3
```

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

11 Configure Storage Foundation and VVR when prompted. Configuring VVR also involves the configuration of Storage Foundation components.

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

12 Choose whether you want to use enclosure-based naming. 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.

If you enter `y` to the enclosure-based naming question, the script decides whether the system is eligible for enclosure-based naming. If it is eligible, confirm whether you want to set up enclosure-based naming.

See the *Veritas Volume Manager Administrator's Guide*

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

13 You have the option of specifying the name of a default 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.

See "Setting the default disk group" on page 146.

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
```

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

```
Will you specify one disk group name for all eligible
systems? [y,n,q,?] (y) y

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

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

16 The script displays the default ports for VVR. Follow the instructions on the screen if you want to change the VVR ports.

The port settings should be identical for the systems that are part of the same Replicated Data Set. They should also be identical for all the systems in a cluster.

```
Do you want to change any of the VVR ports on system01?
[y,n,q] (n) n
```

17 The VVR Statistics Collection Tool collects and maintains the statistics which are helpful in solving VVR performance issues.

Options can be set, such as the frequency for gathering the statistics, and the number of days for which the collected statistics should be preserved.

Change the frequency of online statistics collection, if needed.

```
The frequency of online stats collection on system01
is set to per 10 seconds.
Do you want to change the frequency
of online stats collection on system01 ? [y,n,q] (n) n
```

18 Change the maximum number of days that online statistics are retained, if needed.

```
The maximum number of days for which VVR statistics
can be retained is set to 3 on system01

Do you want to change the maximum number of days
for retaining VVR statistics on system01? [y,n,q] (n) n
```

19 Configure the VVR statistics options (tunables), if needed.

For more information about the VVR statistics options, refer to the *Veritas Volume Replicator Tuning and Planning Guide*.

```
Do you want to view or modify VVR tunables on
system01? [y,n,q,?] (n) n
```

20 The script displays the default ports for VVR, the Statistics Collection Tool options, and the VVR tunables on any additional systems. Follow the instructions on the screen if you want to change the VVR options on these systems.

21 Verify the fully qualified hostnames of the systems. Press Return to continue.

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

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

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

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

23 To start the VVR processes, press Return, or type `y`.

```
Do you want to start Veritas Volume Replicator
processes now? [y,n,q] (y) y
```

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

View 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-****
```

25 The installation script prompts for a reboot after configuration. Reboot the system (or systems) if the install script prompts you to do so.

Installing VVR when VxVM is already installed

If this release of Veritas Volume Manager (VxVM) is already installed on your system, you can start using VVR by installing the VVR license.

After the VVR license is installed, install VVR-specific components and configure VVR.

See "Installing and configuring Veritas Volume Replicator using the common product installer" on page 64.

If a previous version of Veritas Volume Manager (VxVM) is already installed on your system, you must upgrade to this release of VxVM. In some cases, this requires upgrading the operating system (OS) version to the latest version.

See "Upgrading Veritas Storage Foundation from 4.x to 5.0MP3" on page 114.

After VxVM is upgraded, follow the procedure to install VVR-specific components and configure VVR.

See "Installing and configuring Veritas Volume Replicator using the common product installer" on page 64.

To use the new features of VVR 5.0MP3, upgrade the version of each disk group by entering the following command:

```
# vxdg upgrade diskgroup5.0MP3
```

Installing Veritas Enterprise Administrator

This section describes the installation of VEA components.

The VEA server package, `VRTSob`, is installed by the installation script on all nodes that are to be administered. The `VRTSobgui` package contains the VEA client Graphical User Interface (GUI) program, which may also be installed on one or

more of these nodes, or on a separate system that can be used to administer both these and other nodes.

Installing the VEA client on Microsoft Windows

This package can be installed on Windows NT, Windows XP, Windows 2000, Windows 2003, Windows ME, Windows 98 and Windows 95 machines.

To install and run the VEA client, your system must conform to the following specifications:

- Windows Installer 2.0 or later must be present. For information about upgrading Windows Installer, visit:
  http://www.microsoft.com
  For Windows NT 4.0, it is also recommended that you use Windows NT 4.0 Service Pack 6.
- Java Runtime Environment 1.1 or later must be present.
- 100MHz Pentium with 256MB memory or higher specification.
- 22MB available disk space.
- Microsoft Installer is required to install the `VRTSobgui.msi` package. You can get this product from the Microsoft website if it is not already installed on your system.

If you plan to install the GUI client on Windows NT 4.0, Windows Installer must be upgraded to version 2.0. For more information about upgrading Windows Installer, visit:

http://www.microsoft.com

If you are using Windows NT 4.0, it is also recommended that you use Windows NT 4.0 Service Pack 6.

To install the VEA client on a Windows machine

1. Insert the appropriate media disc into your system's DVD-ROM drive.
2. Using Windows Explorer or a DOS Command window, go to the `windows` directory and execute the `vrtsobgui.msi` program with Windows Installer.
3. Follow the instructions presented by the `vrtsobgui.msi` program.
4. After installation is complete, ensure environment changes made during installation take effect by performing one of the following procedures:
   - For Windows NT, Windows 2000, Windows 2003 or Windows XP, log out and then log back in.
   - For Windows ME, Windows 98 or Windows 95, restart the computer.

Chapter 5

Configuring Storage Foundation and High Availability products

This chapter includes the following topics:

Configuring the Storage Foundation products

If the Storage Foundation products were installed using the common product installer, the Veritas Storage Foundation products were already configured during the product installation.

For dababases, additional configuration beyond the product installation script might be necessary.

See "Configuring Storage Foundation Database Editions" on page 82.

If the Storage Foundation products were installed with another method, they may also need to be configured. Review the configuration sections, as needed, and follow the needed procedures.

Configuring Storage Foundation

This section describes how to configure Storage Foundation with the common product installer.

To configure Storage Foundation

1 To configure Storage Foundation, enter the following command:

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

2 When the list of available products is displayed, select Veritas Storage Foundation (SF), enter the corresponding number, and press Return.

```
Select a product to configure:
```

3 Enter the names of the systems on which you want to configure the software.

```
Enter the system names separated by spaces on which to
configure SF: host1
```

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

```
Checking system licensing

SF license registered on host1

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

5 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
```

6 You have 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.

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

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

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

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

```
Will you specify one disk group name for all eligible
systems? [y,n,q,?] (y) y

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

8 Verify the fully qualified hostname of the systems.

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

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

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

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

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

10 The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

11 The configuration completes automatically.

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-****
```

Configuring Storage Foundation and High Availability Solutions

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

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

Required information for configuring Storage Foundation and High Availability Solutions

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

See also the *Veritas Cluster Server Installation Guide*.

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

Veritas Storage Foundation can be configured to use Symantec Security Services.

Running Storage Foundation in Secure Mode guarantees that all inter-system communication is encrypted and that users are verified with security credentials. When running Storage Foundation in Secure Mode, NIS and system usernames and passwords are used to verify identity. Storage Foundation usernames and passwords are no longer used when a cluster is running in Secure Mode.

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

See the *Veritas Cluster Server Installation Guide* for more information on configuring a Symantec Product Authentication Service Root Broker.

The following information is required to configure SMTP notification:

- The domain-based hostname of the SMTP server
- The email address of each SMTP recipient
- A minimum severity level of messages to be sent to each recipient

The following information is required to configure SNMP notification:

- System names of SNMP consoles to receive VCS trap messages
- SNMP trap daemon port numbers for each console
- A minimum severity level of messages to be sent to each console

Configuring Veritas Storage Foundation and High Availability Solutions

After installation, you must configure the product.

Use the procedure in this section if you installed an HA version of the Storage Foundation software.

To configure Storage Foundation product on a cluster

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

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

2 When the list of available products is displayed, select Veritas Storage Foundation (SF), enter the corresponding number, and press Return.

```
Select a product to configure:
```

3 You are prompted to enter the system names (in the following example, "host1" and "host2") on which the software is to be installed. Enter the system name or names and then press Return.

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

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

The installer requires that ssh commands used between systems execute without prompting for passwords or confirmations. If the installer hangs or asks for a login password, stop the installer and run it again with the ssh configured for password free logins, or configure rsh and use the -rsh option.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

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

```
Checking system licensing

SF license registered on host1

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

6 When prompted to configure VCS, enter `y` to configure VCS on these systems.

```
Do you want to configure VCS on these systems at this time?
  [y,n,q] (y) y
```

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

All systems are configured to create one cluster.

7 Enter the unique cluster name and Cluster ID number.

```
Enter the unique cluster name: [?]cluster2
Enter the unique Cluster ID number between 0-65535: [b,?] 76
```

8 The installer discovers the network interfaces (NICs) available on the first system and reports them:

```
  Discovering NICs on host1 ... discovered eth0 eth1 eth2 eth3
```

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`, make sure that the same NICs are available on each system; the installer may not verify this. The NICs should also be the same speed on both systems for the heartbeat links to function properly.

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`, unless a Root Broker has already been set up.

Warning: Before configuring a cluster to operate using Veritas Security Services, another system must already have Veritas Security Services installed and be operating as a Root Broker. Refer to the *Veritas Cluster Server Installation Guide* for more information on configuring a VxSS Root Broker.

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

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

When prompted, set the user name and /or password for the Administrator.

Enter `n` if you want to decline. If you enter `y`, you are prompted to change the password.

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

13 You are prompted to add another user to the cluster.

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

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

14 You are prompted to enter the user information.

```
Enter the user name: [?] myuser
Enter New Password:
Enter Again:

Enter the privilege for user myuser (A=Administrator, O=Operator,
   G=Guest): [?] A
```

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

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

16 When prompted to configure SMTP notification, enter `n` or `y` to configure. To configure SNMP notification, enter the following information. You can then confirm that it is correct, or enter it again.

```
Do you want to configure SMTP notification? [y,n,q] (y) y
Active NIC devices discovered on host1: 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
(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
```

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

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

18 If you installed a valid HA/DR license, the installer prompts you to configure this cluster as a global cluster.

If you installed a HA/DR license to set up replicated data cluster or campus cluster, skip this installer option.

See *Veritas Cluster Server User's Guide* for instructions to set up VCS global clusters.

```
Do you want to configure the Global Cluster Option? [y,n,q] (y) y
```

19 If you select yes, the installer prompts you for a NIC and value for the netmask.

```
Enter the Virtual IP address for Global Cluster Option:
[b,?] (10.10.12.1)
```

20 Verify and confirm the configuration of the global cluster.

```
Global Cluster Option configuration verification:
NIC: eth0
IP: 10.10.12.1
Netmask: 255.255.240.0
Matching Cluster Management Console Virtual IP configuration
Is this information correct? [y,n,q] (y)
```

21 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
```

22 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.

See "Setting the default disk group" on page 146.

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
```

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

```
Will you specify one disk group name for all eligible
systems? [y,n,q,?] (y) y

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

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

25 You may be prompted to verify the fully qualified hostname of the systems. Press Return to continue.

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

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

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

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

27 The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

28 The configuration and startup complete automatically.

View 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-****
```

29 If you installed Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, create a new repository database.

See "Creating and configuring the repository database for DB2 and Oracle" on page 82.

Configuring Storage Foundation Database Editions

Complete the procedures in the following sections to configure Storage Foundation Database Editions.

Creating and configuring the repository database for DB2 and Oracle

After installing Veritas Storage Foundation for Oracle or Veritas Storage Foundation for DB2, you must create and configure the repository database using the `sfua_db_config` script.

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

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

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

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

- Create a separate, non-shared disk group on shared storage. Create a VxVM volume and a VxFS file system and mount the volume.
- It is recommended that you have a separate disk group for the repository volume so that any failovers are independent of other service groups.
- The mount point is used to store the repository database.
- Obtain an unique virtual IP address for public NIC interface.
- Obtain the device names for the public NIC interface for all systems in the cluster.
  For example, use these names.
  hme0
  eth0
- Obtain a subnet mask for the public NIC interface.
- Make sure VCS is not in read-write (-rw) mode. To make sure VCS is in read-only mode, use the following command:

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

Table 5-1 indicates the options available for the `sfua_db_config` script.

Table 5-1 sfua_db_config options

Option	Description
-ssh	Use this option in a high availability (HA) configuration. The option indicates that ssh and scp are to be used for communication between systems. Either ssh or rsh should be preconfigured so that you can execute the commands without being prompted for passwords or confirmations.
-o dropdb	Drops the repository database.
-o unconfig_cluster	Use this option in a high availability (HA) configuration. Unconfigures the repository database from the VCS cluster.
-o dbstatus	Verifies the status of the database and database server.
-o stopserver	Stops the database server.
-o startserver	Starts the database server.
-o serverstatus	Reports the database server status.
-o stopdb	Detaches the repository database from the database server.
-o startdb	Attaches the repository database to the database server.

To create and configure the repository database

1 Run the `sfua_db_config` script as follows:

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

2 Confirm that you are ready to configure the Veritas Storage Foundation for Oracle repository:

```
Are you ready to configure SFORA repository (y/n/q) [y]?
```

3 The mount point is displayed.

```
filesystem mount point for SFORA repository: /sfua_rep
```

4 The network interfaces (NICs) are discovered, and you are prompted to enter the NIC for the repository configuration on each host:

```
Enter the NIC for system host1 for HA Repository configuration:

[eth0]

Enter the NIC for system host2 for HA Repository  configuration:

[eth0]
```

5 Enter the Virtual IP address for repository failover.

```
Enter the Virtual IP address for repository failover:
xxx.xxx.xxx.xxx

Enter the netmask for public NIC interface: [xxx.xxx.xxx.xxx]

Following information will be used for SFORA HA configuration:

 Public IP address: xxx.xxx.xxx.xxx
 Subnet mask: xxx.xxx.xxx.xxx

Public interface:       host1 -> eth0,host2 -> eth0
```

6 The mount point information is displayed, and the script asks for confirmation. Then the repository information is added.

7 Verify that the repository was configured.

If you are installing in a high availability configuration, enter the following command:

```
# /opt/VRTS/bin/hagrp -state
Group        Attribute             System     Value
Sfua_Base    State                 guan       |ONLINE|
Sfua_Base    State                 plover     |OFFLINE|

Note: Sfua_Base group should be online on one node in the
cluster.
```

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

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

Configuring Veritas Volume Manager

Use the following procedures to configure Veritas Volume Manager.

For information on setting up VxVM disk groups and volumes after installation, see "Configuring Veritas Volume Manager" in the *Veritas Volume Manager Administrator's Guide*.

Disabling hot-relocation

The hot-relocation feature is enabled by default and it is recommended that you leave it on. However, you can disable it by preventing the `vxrelocd` daemon from starting up during system startup. This should be done after the VxVM packages have been installed.

For details, see the "Administering hot-relocation" chapter in the *Veritas Volume Manager Administrator's Guide*.

Enabling the Intelligent Storage Provisioning (ISP) feature

If you load the allocator provider package (`VRTSalloc`), enter the following commands to restart the VEA service and enable the Intelligent Storage Provisioning (ISP) feature:

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

Configuring Veritas File System

After installing Veritas File System, you can create a file system on a disk slice or Veritas Volume Manager volume with the `mkfs` command. Before you can use this file system, you must mount it with the `mount` command. You can unmount the file system later with the `umount` command. A file system can be automatically mounted at system boot time if you add an entry for it in the following file:

```
/etc/fstab
```

The Veritas-specific commands are described in the Veritas File System guides and online manual pages.

See the *Veritas File System Administrator's Guide*.

Loading and unloading the file system module

The `vxfs` file system module automatically loads on the first reference to a VxFS file system; this occurs when a user tries to mount a VxFS file system.

In some instances, you may find it efficient to load the file system module manually. For example, some larger class systems can have many dual interface I/O cards with multiple disk chains attached. The device interrogation process when such a system is rebooted can be very time consuming, so to avoid doing a reboot, use the `modprobe` command to load the `vxfs` module:

```
# modprobe vxfs ; modprobe vxportal ; modprobe fdd
```

Do not use the `insmod` command to load the `vxfs` module as `insmod` does not examine the module configuration file `/etc/modprobe.conf`.

To determine if the modules successfully loaded, use the `lsmod` command as shown here:

```
# lsmod | grep vxportal
vxportal               2952           0
vxfs                   3427960        0    fdd vxportal
# lsmod | grep fdd
fdd                    67212          0    (unused)
vxfs                   3427960        0    [fdd vxportal]
# lsmod | grep vxfs
vxfs                   3427960        0    [fdd vxportal]
```

The first field in the output is the module name. You can unload the modules by entering:

```
# rmmod fdd
# rmmod vxportal
# rmmod vxfs
```

The `rmmod` command fails if there are any mounted VxFS file systems. To determine if any VxFS file systems are mounted, enter:

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

About configuring and migrating Storage Foundation components

Review the *Storage Foundation Release Notes*.

Detailed information about configuring and upgrading Storage Foundation, Storage Foundation and High Availability Solutions, Storage Foundation Cluster File Systems, and Veritas Volume Replicator can be found in the following documents:

- *Veritas Storage Foundation Installation Guide*
- *Veritas File System Administrator's Guide*
- *Veritas Cluster File System Administrator's Guide*
- *Veritas Cluster Server Installation Guide*
- *Veritas Volume Manager Administrator's Guide*
- *Veritas Volume Replicator Administrator's Guide*
- *Veritas Cluster Server User's Guide*

Configuring Storage Foundation Cluster File System

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

Perform the procedures in the following section to configure Storage Foundation Cluster File System.

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 When the list of available products is displayed, select Veritas Storage Foundation Cluster File System (SFCFS), enter the corresponding number, and press Return.

```
Select a product to configure:
```

3 You are prompted to enter the system names (in the following example, "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
```

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

The installer requires that ssh commands used between systems execute without prompting for passwords or confirmations. If the installer hangs or asks for a login password, stop the installer and run it again with the ssh configured for password free logins, or configure rsh and use the -rsh option.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

5 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
```

6 Any running SFCFS processes are stopped. Enter Return to continue.

7 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.

Enter `y` or `n` for configuring I/O Fencing in enabled mode.

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

See the *Storage Foundation Cluster File System Administrator's Guide* for more information.

8 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
```

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

```
Discovering NICs on host1 ... discovered eth0 eth1 eth2 eth3
```

10 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.

11 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.

12 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
```

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

Enter `y` or `n` to set the username and password.

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

14 Enter `y` or `n` to add another user to the cluster.

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

15 Enter `y` if the information is correct.

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

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

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

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

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

18 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
```

19 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
```

20 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
```

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

22 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)
```

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

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

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

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

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

25 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-****
```

26 If you do not plan to use SFCFS file systems to store the Oracle database or binaries, you have completed the SFCFS installation and configuration.

If you upgraded from 5.0, 5.0 MP1 or 5.0 MP2 to 5.0 MP3 and plan to use SFCFS file systems to store the Oracle database or binaries, perform the following steps on all the nodes in the cluster:

- Install `VRTSgms` package from the 5.0 MP3 CD
- Install `VRTSodm-common` and `VRTSodm-platform` package from the 5.0 MP3 CD

27 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.

Configuring Veritas Volume Replicator

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

To configure VVR

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

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

2 When the list of available products is displayed, select Veritas Volume Replicator (VVR), enter the corresponding number, and press Return.

```
Select a product to configure:
```

3 At the prompt, enter the name of the system or systems on which you want to configure VVR.

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

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

The installer requires that ssh commands used between systems execute without prompting for passwords or confirmations. If the installer hangs or asks for a login password, stop the installer and run it again after setting it up. Set up the system with ssh configured for password free logins, or configure remote shell and use the -rsh option.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

5 The script continues the initial system check. The script confirms success by displaying information, such as the OS version, communication with the remote hosts, and whether the required VVR packages are installed. Press Return to continue.

6 The script proceeds to verify whether the required licenses are installed. If a valid license for VVR is not present, the script prompts you to enter a license. The script validates whether the current license enables VVR.

See "Symantec product licensing" on page 21.

You cannot proceed until a valid VVR license has been entered. If a valid VVR license is present on the system, the script provides the option to add additional licenses. Press Return to continue.

7 The script enables you to choose whether you want to use enclosure-based naming. 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.

If you enter `y` to the enclosure-based naming question, the script decides whether the system is eligible for enclosure-based naming. If it is eligible, confirm whether you want to set up enclosure-based naming.

See the *Veritas Volume Manager Administrator's Guide*

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

8 Specify the default name of a disk group for 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 following command on a system.

```
vxdctl defaultdg diskgroup
```

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
```

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

```
Will you specify one disk group name for all eligible
systems? [y,n,q,?] (y) y

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

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

11 The script displays the default ports for VVR. Follow the instructions on the screen if you want to change the VVR ports.

The port settings should be identical for the systems that are part of the same Replicated Data Set. They should also be identical for all the systems in a cluster.

```
Do you want to change any of the VVR ports on system01?
[y,n,q] (n) n
```

12 The VVR Statistics Collection Tool collects and maintains the statistics which are helpful in solving VVR performance issues.

Options can be set, such as the frequency for gathering the statistics, and the number of days for which the collected statistics should be preserved.

Change the frequency of online statistics collection, if needed.

```
The frequency of online stats collection on system01
is set to per 10 seconds.
Do you want to change the frequency
of online stats collection on system01 ? [y,n,q] (n) n
```

13 Change the maximum number of days that online statistics are retained, if needed.

```
The maximum number of days for which VVR statistics
can be retained is set to 3 on system01

Do you want to change the maximum number of days
for retaining VVR statistics on system01? [y,n,q] (n) n
```

14 Configure the VVR statistics options (tunables), if needed.

For more information about the VVR statistics options, refer to the *Veritas Volume Replicator Tuning and Planning Guide*.

```
Do you want to view or modify VVR tunables on
system01? [y,n,q,?] (n) n
```

15 Repeat steps 11 to 14 for all other systems.

16 Verify the fully qualified hostnames of the systems. Press Return to continue.

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

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

See the *Storage Foundation Manager Installation Guide* for details.

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

18 To start the VVR processes, press Return, or type `y`.

```
Do you want to start Veritas Volume Replicator
processes now? [y,n,q] (y) y
```

19 The configuration and startup completes automatically.

View 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-****
```

Configuring and starting Veritas Enterprise Administrator

Before using the Veritas Enterprise Administrator server or client, start them both.

Optional configuration can also be completed at this time.

Stopping and starting the VEA server

After installing the VEA packages, the VEA server may need to be stopped and restarted. The VEA service is automatically started when you reboot your system.

To start up the VEA server

1. Check the state of the VEA server.

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

2. Stop the VEA server.

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

   You can also stop the VEA server manually by killing the `vxsvc` process.

3. Start the VEA server.

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

   The VEA server is automatically started on a reboot.

Using the VEA client to administer local and remote systems on Linux

Users with appropriate privileges can use the VEA client to administer a local or remote machine. VERITAS Volume Manager and the VEA server must be installed, and the VxVM configuration daemon, `vxconfigd`, and the VEA server daemon, `vxsvc`, must be running on the machine to be administered.

To use the VEA client GUI provided with this release to administer Veritas software on other platforms, you must upgrade those systems to at least the following releases:

- AIX requires Veritas Storage Foundation 3.2.2 or later
- HP-UX requires Veritas Volume Manager 3.5 Update 2 or later
- Linux requires Veritas Storage Solutions 2.2 or later
- Solaris requires Veritas Storage Foundation 3.5 MP2 or later

To use the VEA client on Linux to administer the local system, type:

```
# /opt/VRTSob/bin/vea
```

To use the VEA client on Linux to administer a remote system, type:

```
# /opt/VRTSob/bin/vea -host remote_system -user username \
      -password password
```

To use the VEA client on Windows to administer a remote system, select Start > Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator.

VMSA and VEA co-existence

If you do not plan to use VMSA to administer other (pre-VxVM 3.5) machines, then you should uninstall VMSA before installing VEA. You can later do a client-only install if you want to run the VMSA client on your machine.

Warning: The release of VEA that ships with VxVM 5.0 is not compatible with VMSA, the previous Veritas Volume Manager GUI. You cannot run VMSA with VxVM version 5.0.

If you do not remove VMSA, the following warning appears during a reboot:

```
Veritas VM Storage Administrator Server terminated.

Stopping Veritas VM Storage Administrator Server

### Terminated
```

Configuring Veritas Enterprise Administrator for databases

You may need to configure Veritas Enterprise Administrator (VEA) for databases so that users can access the features.

Configuring Veritas Enterprise Administrator for Oracle

You may need to update Veritas Enterprise Administrator (VEA) so that users other than `root` can access features.

Adding users to the VEA Service Console Registry for Oracle

You may want to add users to the VEA service console registry to allow access to the interface to users other than `root`. You also have the option to give database administrators `root` privileges.

To add users other than root to the Veritas Enterprise Administrator Service console registry

1. Make sure that the optional GUI package was installed.
2. Stop the VEA server.

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

3 To give `root` privileges to the database administrator, use the `vxdbedusr` command as follows.

```
# /opt/VRTS/bin/vxdbedusr -a user [-A] [-f] -n user_name
```

where:

`-a user` adds a user to the registry

`-A` grants the user root access

`-f` allows the user to be a user other than the `/opt/VRTSdbed` owner.

`-n` indicates the name of the user.

For example, to add a database administrator with the name "oracle" as a user with `root` privileges, enter the following:

```
# /opt/VRTS/bin/vxdbedusr -a user -A -f -n oracle
```

4 To add a user without `root` privileges, use the `vxdbedusr` command as follows.

```
# /opt/VRTS/bin/vxdbedusr -a user -n user_name
```

where `-a` adds a user to the registry.

For example, to add "oracle" as a user, enter the following:

```
# /opt/VRTS/bin/vxdbedusr -a user -n oracle
```

5 To add a group to the console registry, use the `vxdbedusr` command as follows:

```
# /opt/VRTS/bin/vxdbedusr -a group [-A] [-f] -n group_name
```

where:

`-a user` adds a user group to the registry

`-A` grants the user group root access

`-f` allows the group access to the GUI.

For example, to add "dba" as a group, enter the following:

```
# /opt/VRTS/bin/vxdbedusr -a group -A -f -n dba
```

6 Restart the VEA Server.

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

Removing users from the VEA Service Console Registry for Oracle

You may need to restrict access to the VEA service console registry. You can remove users or user groups from the registry if they have been previously added.

You cannot remove `root` from the VEA console registry.

To remove users other than root from the Veritas Enterprise Administrator Service console registry

1. Make sure that the optional GUI package was installed.

   ```
   # rpm -q VRTSorgui
   VRTSorgui-common-5.0.00.A10-SLES9
   ```

2. Stop the VEA server.

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

3. Use the `vxdbedusr` command to remove a group or user.

   ```
   # /opt/VRTS/bin/vxdbedusr -r {user | group} \
   -n {user_name | group_name}
   ```

   where `-r` removes a user or user group from the registry.

   For example, to remove the user "oracle," enter the following:

   ```
   # /opt/VRTS/bin/vxdbedusr -r user -n oracle
   ```

4. Restart the VEA Server.

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

Setting up Veritas Enterprise Administrator for DB2

You may want to add users to the VEA Authorization Database (AZDB) to allow access to the interface to users other than `root`. You also have the option to give database administrators `root` privileges.

Adding users to Veritas Enterprise Administrator for DB2

You may want to add users to the VEA Authorization Database (AZDB) to allow access to the interface to users other than `root`. You also have the option to give database administrators `root` privileges.

To add users other than root to the Veritas Enterprise Administrator AZDB

1 Make sure that the optional GUI package was installed.

```
# rpm -q VRTSd2gui-common
```

The output should look similar to the following:

```
VRTSd2gui-common-5.0.00.A10-SLES9
```

2 Stop the VEA server.

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

3 To give `root` privileges to the database administrator, use the `vxdb2edusr` command as follows.

```
# /opt/VRTS/bin/vxdb2edusr -a {user | group} [-A] [-f] -n \
      user_name [-h fully_qualified_host_name -d domain_name \
    -t domain_type]
```

where:

`-a user` adds a user to the registry

`-A` grants the user root access

`-f` allows the user to be a user other than the `/opt/VRTSdb2ed` owner.

`-n` indicates the name of the user or group.

`-h` specifies a fully qualified host name on which you want to add a user.

`-d` specifies the domain to which the user belongs.

`-t` specifies the type of domain to which the user belongs. Valid values are `nis`, `nisplus`, `Idap`, `unixpwd`, and `gssapi`.

For example, to add a database administrator with the name `db2inst1` as a user with `root` privileges, enter the following:

```
# /opt/VRTS/bin/vxdb2edusr -a user -A -f -n db2inst1
```

4 To add a user without `root` privileges, use the `vxdbedusr` command as follows.

```
# /opt/VRTS/bin/vxdbe2dusr -a user -n user_name
```

where `-a` adds a user to the registry.

For example, to add "db2inst1" as a user, enter the following:

```
# /opt/VRTS/bin/vxdb2edusr -a user -n db2inst1
```

5 To add a group to the console registry, use the `vxdb2edusr` command as follows:

```
# /opt/VRTS/bin/vxdb2edusr -a group [-A] [-f] -n group_name
```

where:

`-a group` adds a group to the registry

`-A` grants the group root access

`-f` allows the group to be other than the `/opt/VRTSdb2ed` owner.

`-n` indicates the name of the user or group.

For example, to add `dba` as a group, enter the following:

```
# /opt/VRTS/bin/vxdb2edusr -a group -n dba
```

6 Restart the VEA Server.

```
# /opt/VRTS/bin/vxsvcctrl restart
```

Removing users from Veritas Enterprise Administrator for DB2

You may need to restrict access to the VEA Authorization Database (AZDB). You can remove users or user groups from the AZDB if they have been previously added.

You cannot remove `root` from the AZDB.

To remove users other than root from the VEA service console registry

1 Make sure that the optional GUI package was installed.

```
# rpm -q VRTSd2gui-common
```

The output should look similar to the following:

```
VRTSd2gui-common-5.0.00.A10-SLES9
```

2 Stop the VEA server.

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

3 Use the `vxdb2edusr` command to remove a group or user.

```
# /opt/VRTS/bin/vxdb2edusr -r {user | group} \
      -n {user_name | group_name} \
    [-h fully_qualified_host_name -d domain_name \
    -t domain_type]
```

where `-r` removes a user or user group from the registry.

For example, to remove the user `db2inst1`, enter the following:

```
# /opt/VRTS/bin/vxdb2edusr -r user -n db2inst1
```

4 Restart the VEA Server.

```
# /opt/VRTS/bin/vxsvcctrl restart
```

Chapter 6

Upgrading Storage Foundation

This chapter includes the following topics:

Upgrading Storage Foundation or the operating system

Perform the procedures in the following sections to upgrade Storage Foundation or your operating system, or both. You can perform an upgrade to Storage Foundation using the Veritas product installer or product installation script if you already have Storage Foundation installed.

This section describes how to upgrade Veritas Storage Foundation, Veritas Storage Foundation for DB2, and Veritas Storage Foundation for Oracle.

Caution: Make sure that supported combinations of Storage Foundation and the operating system are present on your system during the upgrades. Do not upgrade to a version of Storage Foundation that is not supported with the current operating system.

Planning the upgrade

Be sure that the administrator doing the upgrade has root access and a working knowledge of Linux operating system administration.

Complete the following tasks in advance of upgrading:

- Check the latest *Storage Foundation Release Notes* to verify that the system is running a supported Linux version.
- Schedule sufficient outage time for the upgrade, and downtime for any applications using the VxFS file systems or VxVM volumes.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. (This may not be practical, but if done, offers a failback point.)
- To upgrade on a remote host, rsh or ssh must be set up.
  See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.
- Determine if the root disk is encapsulated.
  See "Determining if the root disk is encapsulated" on page 107.
- Select the method to upgrade.
  See "Upgrade paths for Storage Foundation 5.0MP3" on page 107.

Saving system information before upgrade

Use the following procedure to save system information before an upgrade.

To save system information before an upgrade

1. Log in as superuser.
2. Before upgrading, ensure that you have made backups of all data that you want to preserve.
3. In particular, you will need the information in files such as `/boot/grub/menu.lst`, `/etc/grub.conf` or `/etc/lilo.conf` (as appropriate), and `/etc/fstab`.

4 Copy `fstab` to `fstab.orig`:

```
# cp /etc/fstab /etc/fstab.orig
```

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

6 If you are installing the HA version of the Veritas Storage Foundation 5.0MP3 software, follow the guidelines given in the *Veritas Cluster Server Installation Guide* and *Veritas Cluster Server Release Notes* for information on preserving your VCS configuration across the installation procedure.

Determining if the root disk is encapsulated

Check if the system's root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system.

If the root disk is encapsulated, follow the appropriate upgrade procedures.

Upgrade paths for Storage Foundation 5.0MP3

The only supported upgrade paths to Storage Foundation 5.0MP3 are from Storage Foundation 4.1 or 5.0, including all release packs and maintenance packs. Before upgrading to Storage Foundation 5.0MP3, upgrade the Linux system to a version that is supported for this release.

See "Supported Linux operating systems" on page 38.

Table 6-1 shows the upgrade paths for Storage Foundation on Linux.

Table 6-1 Upgrade paths for Storage Foundation on Linux

Storage Foundation version	Upgrade procedure
prior to 4.1	Upgrade to Storage Foundation 5.0MP3 is not supported. You must uninstall Storage Foundation using the procedure in the Storage Foundation Installation Guide for your version, then install Storage Foundation 5.0MP3.

Table 6-1 Upgrade paths for Storage Foundation on Linux *(continued)*

Storage Foundation version	Upgrade procedure
SF 4.1, including Maintenance Packs and Rolling Patches **Note:** To upgrade RHEL5 or SLES10 systems from 4.1MP4 to 5.0MP3, you must first upgrade 4.1MP4 to 4.1MP4RP2.	Upgrade Storage Foundation using the procedure: See "Upgrading Veritas Storage Foundation from 4.x to 5.0MP3" on page 114.
SF 5.0, including Maintenance Packs and Rolling Patches	Upgrade to 5.0MP3 using the procedure: See "Upgrading Storage Foundation from 5.0 to 5.0MP3" on page 109.

Performing pre-installation checks and configuration

Use the following procedure to prepare for the upgrade.

To prepare for the upgrade

1. Ensure that you have created a valid backup.

   See "Saving system information before upgrade" on page 106.

2. Review the *Veritas Storage Foundation Release Notes*.

3. Ensure that you have enough file system space to upgrade Veritas Storage Foundation. Also, identify where you will be copying the distribution and patch tar files. The usual place is `/patches/Veritas` when the root file system has enough space or `/var/tmp/patches` if the `/var` file system has enough space.

   Do not put the files under `/tmp`, which is erased during a system reboot. Do not put the files on a file system that is inaccessible prior to running the upgrade script.

   You may use a Veritas-supplied DVD for the upgrade as long as modifications to the upgrade script are not required. If `/usr/local` was originally created as a slice, modifications are required. See Step 8 below for details.

4. For any startup scripts in `/etc/rcS.d`, you should comment out any application commands or processes that are known to hang if their file systems are not present.

5. Make sure that all users are logged off and that all major user applications are properly shut down.

6 If required, upgrade VxFS disk layouts to a supported version.

Some previous layout versions cannot be mounted on VxFS 5.0MP3. You can upgrade these layout versions online before installing VxFS 5.0MP3, or upgrade them using `vxfsconvert` after installing VxFS 5.0MP3.

See "About upgrading disk layout versions" on page 143.

7 Upgrade arrays (if required)

See "Upgrading external ASLs and APMs" on page 109.

Upgrading external ASLs and APMs

The Storage Foundation 5.0MP3 release only supports upgrading external ASL and APM packages from release 5.0 or later.

If your system uses any existing ASLs and APMs from 4.x release, you may need to download a 5.0 version from Symantec. Check the latest array support information to determine whether an updated ASL or APM package is available for your arrays.

See the 5.0 MP3 Hardware Compatibility List for information about supported arrays.

Before upgrading a Storage Foundation product, you must manually remove any existing external ASL or APM packages which are lower than release 5.0.

After completing the upgrade, obtain the required updated ASLs or APMs to ensure the array is claimed correctly.

Upgrading Storage Foundation from 5.0 to 5.0MP3

Use this procedure to upgrade to 5.0MP3 from 5.0, or from an earlier 5.0 Maintenance Pack.

Caution: Existing data could be destroyed on any disks that are touched by upgrading the operating system. While upgrading, do not reconfigure any disks other than the root disk. To ensure the integrity of your data, it is recommended that you back it up before starting the upgrade.

The following steps are necessary to successfully upgrade Storage Foundation 5.0 to 5.0MP3:

- Save system information before the upgrade.
  See "Saving system information before upgrade" on page 106.

- Perform preinstallation checks and configuration for the upgrade.
  See "Performing pre-installation checks and configuration" on page 108.
- Upgrade Storage Foundation to 5.0MP3.
  See "Upgrading Storage Foundation software from 5.0 to 5.0MP3 using the product installer" on page 112.

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

Upgrading external 4.x ASL or APM packages from a previous 5.0 release to Storage Foundation 5.0MP3

Before upgrading to 5.0MP3, you must remove any 4.x ASLs that may be installed on the system.

After completing the upgrade, obtain the required updated ASLs or APMs to ensure the array is claimed correctly.

To remove external 4.x ASL or APM packages for 5.0 installation

1. Before you remove any packages, ensure that you are not running anything in VxVM volumes, and ensure that no volumes are mounted.

   These steps are necessary to prevent attempts to access the data in disks that were claimed by these ASLs or APMs after the packages are removed. Attempting to access the data could lead to data corruption if the disks are not claimed correctly.

2. Determine which external ASL packages are installed:

   ```
   # rpm -qf /etc/vx/lib/discovery.d | grep -v "^VRTSvxvm-common"
   ```

   This command lists the packages which installed any files in the ASL directory. The package `VRTSvxvm-common*` package is the base VxVM package. Any other packages in this directory are external ASL packages.

   Example output:

   ```
   # rpm -qf /etc/vx/lib/discovery.d | grep -v "^VRTSvxvm-common"
   VRTSIBM-DS4xxx-2.0-1.0
   ```

   The sample output shows an external ASL named `VRTSIBM-DS4xxx-2.0-1.0`.

3. If the 5.0 ASL is installed for MSA1500, remove it. The support for this array has been moved into the base VxVM package.

   Determine if the MSA1500 ASL is installed:

   ```
   # rpm -qa | grep "VRTSHP-MSA1500"
   ```

4 For the remaining ASL packages listed in step 2, determine the library file:

```
# rpm -q --filesbypkg pkg
```

where *pkg* is the ASL package name.

5 Determine if an ASL is 4.x:

```
# /etc/vx/diag.d/vxcheckasl libfile /dev/null | grep "ASL_VERSION"
```

where *libfile* is the name of ASL library found in step 4.

For example:

```
# /etc/vx/diag.d/vxcheckasl libvxhpmsa.so /dev/null | grep "ASL_VERSION
ASL_VERSION: vm-5.0-rev-2
```

If the version is 5.0, then you do not need to remove this ASL. If the version is less than 5.0, remove the ASL.

6 Remove any 4.x external ASLs which have a version less than 5.0:

```
# rpm -e VRTSIBM-DS4xxx-2.0-1.0
```

7 Determine which external APM packages are installed:

```
# rpm -qf /etc/vx/apmkey.d | grep -v "^VRTSvxvm-platform"
```

This command displays package names which installed any APM keys. Any package other than `VRTSvxvm-platform*` is an external APM package.

8 For the APM packages listed in the output of step 7, find the APM name by using the following command:

```
# rpm -q --filesbypkg pkg
```

where *pkg* is the APM package.

Look for the following entry:

```
/etc/vx/kernel/apm.ko.*
```

9 Determine if an APM is 4.x:

```
# vxdmpadm listapm apm | grep "^VxVM version:"
```

where *apm* is the name of the APM.

For example:

```
# vxdmpadm listapm dmpCLARiiON | grep "^VxVM version:"
VxVM version: 41
```

This example indicates that this is a 4.1(41) APM and must be removed.

10 Remove any 4.x APMs that are installed on the machine:

```
# rpm -e VRTSIBM-DS4xxx-2.0-1.0
```

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

Use the following procedure to upgrade Storage Foundation 5.0MP3. This procedure can be used to upgrade on a standalone system, or on the nodes of a cluster.

For an upgrade of Storage Foundation Cluster File System, refer to that upgrade section for the steps for a full or phased upgrade.

See "Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0" on page 152.

To upgrade Storage Foundation 5.0MP3

1 Load and mount the disc.

See "Mounting a software disc" on page 46.

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.

```
# shutdown -r now
```

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*.

Upgrading Veritas Storage Foundation from 4.x to 5.0MP3

If you are running an earlier release of Veritas Storage Foundation, you can upgrade to the latest version of Veritas Storage Foundation, Veritas Storage Foundation for Oracle, or Veritas Storage Foundation for DB2 using the procedures described in this chapter.

For a cluster, use the appropriate procedures to upgrade Veritas Storage Foundation, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation for DB2.

See "Upgrading the Veritas software" on page 115.

If you need to upgrade your kernel with Veritas Storage Foundation 5.0MP3 already installed, use the kernel upgrade procedure.

See the *Veritas Volume Manager Administrator's Guide* for information about upgrading the kernel.

Upgrading external ASL or APM packages from Storage Foundation 4.x to 5.0MP3

If Storage Foundation 4.x is installed on your system, remove all ASLs or APMs before upgrading from 4.x to 5.0MP3.

After completing the upgrade, obtain the required updated ASLs or APMs to ensure the array is claimed correctly.

To remove external ASL or APM packages for a 4.x installation

1 Before you remove any packages, make sure you are not running anything in VxVM volumes, make sure that no volumes are mounted etc.

These steps ensure that VxVM does not access the data in disks that were claimed by these ASLs or APMs after the packages are removed. Attempting to access the data could lead to data corruption if the disks are not claimed correctly.

2 Determine which external ASL packages are installed:

```
# rpm -qf /etc/vx/lib/discovery.d | grep -v "^VRTSvxvm-common"
```

This command lists the packages which installed any files in the ASL directory. The `VRTSvxvm-common*` package is the base VxVM package. Any other packages in this directory are external ASL packages.

Example output:

```
# rpm -qf /etc/vx/lib/discovery.d | grep -v "^VRTSvxvm-common"
VRTSIBM-DS4xxx-2.0-1.0
```

The sample output shows an external ASL named `VRTSIBM-DS4xxx-2.0-1.0`. Because the discovery command was run in a 4.x installation, any external ASL packages that are listed are version 4.x or below.

3 Remove any external ASLs:

```
# rpm -e VRTSIBM-DS4xxx-2.0-1.0
```

4 Determine which external APM packages are installed:

```
# rpm -qf /etc/vx/apmkey.d | grep -v "^VRTSvxvm-platform"
```

This command lists the packages which installed any APM keys. Any package other than `VRTSvxvm-platform*` is an external APM package.

5 Remove these APMs before upgrading Storage Foundation:

```
# rpm -e VRTSIBM-DS4xxx-2.0-1.0
```

Upgrading the Veritas software

Use this procedure to upgrade the Veritas software from 4.x to 5.0MP3 on a stand-alone system or an HA cluster.

To upgrade a Veritas Storage Foundation product

1. Log in as superuser.
2. If you are upgrading Veritas Storage Foundation for DB2, resynchronize all existing snapshots before upgrading.

   ```
   # /opt/VRTS/bin/db2ed_vmsnap -D DB2DATABASE -f SNAPPLAN \
       -o resync
   ```

3. Take all service groups offline.

   List all service groups:

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

   For each service group listed, take it offline:

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

4. If the root disk is encapsulated, unmirror and unencapsulate the root disk as described in the following steps, to be performed in the order listed:

   - 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
     ```

     Warning: 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.

     If your system is running VxVM 4.1 MP2, the following remnants of encapsulation will still be present:

     - Partition table entries for the private and public regions

- GRUB or LILO configuration entries for VxVM

5 If you are upgrading from 4.1 MP2, perform this step to correct the configuration entries.

Otherwise, proceed to step 6.

- Run the `fdisk` command on the root disk, as shown in this example:

```
# fdisk -l /dev/sda
      Disk /dev/sda: 36.3 GB, 36398825472 bytes
      255 heads, 63 sectors/track, 4425 cylinders
      Units = cylinders of 16065 * 512 = 8225280 bytes

       Device Boot    Start   End     Blocks    Id  System
      /dev/sda1       1       13      104391    83  Linux
      /dev/sda2       14      2001    15968610  83  Linux
      /dev/sda3       1       4425    35543781  7e  Unknown
      /dev/sda4       2002    4425    19470780  5   Extended
      /dev/sda5       2002    3001    8032468+  83  Linux
      /dev/sda6       3002    3003    16033+    82  Linux swap
      /dev/sda7     4425      4425    1024      7f  Unknown
```

Partitions `/dev/sda3` and `/dev/sda7` with identifiers `7f` and `7e` correspond to the private and public regions respectively.

- Run the `fdisk` command again to remove the private and public partitions, `/dev/sda3` and `/dev/sda7`.

```
# fdisk /dev/sda
      The number of cylinders for this disk is set to 4425.
      There is nothing wrong with that, but this is larger than
      1024,and could in certain setups cause problems with:
      1) software that runs at boot time (e.g., old versions of
      LILO)
      2) booting and partitioning software from other OSs
      (e.g., DOS FDISK, OS/2 FDISK)

      Command (m for help): d
      Partition number (1-7): 3

      Command (m for help): d
      Partition number (1-7): 7

      Command (m for help): w
      The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with
error 16: Device or resource busy.
The kernel still uses the old table.
The new table will be used at the next reboot.
Syncing disks.
```

- Edit the /etc/fstab file, and ensure that the entries for the root file system (/) and for other file systems on the root disk correspond to the correct disk partitions. If they do not, make any necessary changes to allow the system to boot correctly. For the example layout, no update is necessary. However, if the VxVM private region had been /dev/sda6, and other logical partitions followed it in the listing, the numbers of these partitions would decrease by 1. For example, /dev/sda7 would become /dev/sda6, /dev/sda8 would become /dev/sda7, and so on.
  Alternatively, you can copy /etc/fstab.b4vxvm back to /etc/fstab if you are certain that the entries are still valid.
- Correct the configuration of the boot loader that is used on your system. For the following loaders, correction the configuration:
- For the GRUB boot loader, edit the /boot/grub/menu.lst file.
  Remove all entries between and including the vxvm_root_default_START and vxvm_root_default_END comment markers, as shown in this example:

```
#vxvm_root_default_START ( do not remove)
# Default menu entry number has been set to vxvm_root.
#  - the vxvm_root default entry number is: 3
#  - the original  default entry number is: 0
#  - the selected  default entry number is: 0
#  - the original grub configuration is in: \
/boot/grub/menu.lst.b4vxvm
default=3
#vxvm_root_default_END ( do not remove)
```

  Remove all entries between and including the vxvm_root_START and vxvm_root_END comment markers, as shown in this example:

```
#vxvm_root_START ( do not remove)
title vxvm_root
root (hd0,0)
kernel /vmlinuz root=/dev/sda2  vga=0x314 console=tty0 \
```

```
console=ttyS0 selinux=0 resume=/dev/sda6 elevator=cfq \
showopts initrd /VxVM_initrd.img
#vxvm_root_END ( do not remove)
```

Change to the original boot kernel that was used before the root disk was encapsulated by uncommenting the line that starts `#default`, as shown in this example:

```
color white/blue black/light-gray
#default 0
timeout 8
```

which would become:

```
color white/blue black/light-gray
default 0
timeout 8
```

Save the changes to the `/boot/grub/menu.lst` file.
Alternatively, you can copy `/boot/grub/menu.lst.b4vxvm` back to `/boot/grub/menu.lst` if you are certain that the entries are still valid.

- For the LILO boot loader, edit the `/etc/lilo.conf` file.
  Remove all entries between and including the `vxvm_rootgeom_START` and `vxvm_rootgeom_END` comment markers, as shown in this example:

```
#vxvm_rootgeom_START ( do not remove )
#NOTE: Only vxvm_root entry will be able to boot the
# system, while your root disk is under Volume Manager.
# Also, running -R/lock/fallback options of LILO may
# render your system unbootable.
disk=/dev/vx/dsk/bootdg/rootvol
bios=0x80
sectors=63
heads=255
cylinders=4425
partition=/dev/vx/dsk/bootdg/bootvol
start=63
#vxvm_rootgeom_END ( do not remove )
```

  Remove all entries between and including the `vxvm_root_START` and `vxvm_root_END` comment markers, as shown in this example:

```
#vxvm_root_START ( do not remove)
image=/boot/vmlinuz
```

```
label=vxvm_root
initrd=/boot/VxVM_initrd.img
read-only
append="root=/dev/sda2 vga=0x314 console=tty0 \
console=ttyS0 selinux=0 resume=/dev/sda6 elevator=cfq \
showopts"
#vxvm_root_END ( do not remove)
```

Change to the argument to the `default=` attribute from `vxvm_root` to `Linux`, as shown in this example:

```
boot=/dev/sda
default=Linux
timeout=50
```

Save the changes to the `/etc/lilo.conf` file.

Alternatively, you can copy `/etc/lilo.conf.b4vxvm` back to `/etc/lilo.conf` if you are certain that the entries are still valid.

Run the following command after updating the `/etc/lilo.conf` file:

```
# /sbin/lilo
```

- Reboot the system.

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 Verify that all file systems have been cleanly unmounted:

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

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

Perform the following steps in the order listed:

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

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

  This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.
  There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

  ```
  file system device busy
  ```

  You know for certain that an extended operation is pending if the following message is generated on the console:

  ```
  Storage Checkpoint asynchronous operation on file_system
  file system still in progress.
  ```

- If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
- Repeat this step to verify that the unclean file system is now clean.

9 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.

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

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

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

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

11 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.

12 Perform any necessary preinstallation checks.

See "About planning for a Storage Foundation installation" on page 19.

13 Upgrade Veritas Storage Foundation 5.0MP3 or Veritas Storage Foundation 5.0MP3 for DB2 and any additional required packages on the nodes by running the `installer` script.

See "About the common product installer" on page 47.

14 Shut down and reboot each of the upgraded nodes. After the nodes come back up, application failover capability is available for that group.

15 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*.

16 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node that you recorded in step 11.

17 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.

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

```
# haconf -makerw
```

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

```
# hasys -unfreeze -persistent nodename
```

20 Make the configuration read-only:

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

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

```
# hagrp -online cvm -sys nodename
```

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

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

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

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

24 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

25 Repeat step 4 through step 24 for the second group of nodes.

26 You can perform the followting optional configuration steps:

- If you want to use features of Veritas Storage Foundation 5.0MP3 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.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
  See "Upgrading VxFS disk layout versions" on page 143.
  See "Upgrading VxVM disk group versions" on page 146.
- If you upgraded to the current version of Veritas Storage Foundation for DB2, complete the DB2 configuration.
  See "Upgrading to the new repository database for DB2 and Oracle" on page 142.

Upgrading Veritas Enterprise Administrator clients

This section provides information about upgrading VEA clients.

Upgrading the VEA client on a Microsoft Windows system

To upgrade the VEA client on a Windows system, first uninstall the existing VEA client package as follows:

1 Log in as administrator.

2 Select **Start > Settings > Control Panel**.

3 Double-click **Add/Remove Programs** to display a list of installed products.

4 Select **Veritas Enterprise Administrator** from the list, and click the **Remove** button.

5 Click **Yes** when a dialog box appears asking you to confirm the removal.

6 After removing the existing package, install the new VEA client package by following the instructions for a new installation.

See "Installing Veritas Enterprise Administrator" on page 69.

Upgrading Veritas Volume Replicator

Supported upgrade methods for Veritas Volume Replicator

This section explains how to upgrade Veritas Volume Replicator (VVR) from an earlier version to VVR 5.0MP3.

See "Upgrading Veritas Volume Replicator using the Veritas product installer" on page 124.

To upgrade VVR from 5.0 (or a 5.0 Maintenance Pack) to 5.0MP3, use the `installmp` procedure.

See "Upgrading Storage Foundation from 5.0 to 5.0MP3" on page 109.

Upgrading Veritas Volume Replicator using the Veritas product installer

This section describes using the Veritas product installer. Use this method to upgrade VVR, unless you are upgrading in a scenario that is not supported by the Veritas product installer.

Note: We recommend using the Veritas product installer for upgrading VVR when possible.

See "Supported upgrade methods for Veritas Volume Replicator" on page 124.

To upgrade VVR only, use the Veritas product installer and select the Veritas Volume Replicator option. You can also use the `installvvr` script.

If you have multiple Veritas products, select the option for the appropriate Veritas product suite, and refer to the corresponding installation guide for more details. For example, if you have Veritas Storage Foundation installed, select Veritas Storage Foundation in the Veritas product installer, or use the `installsf` script.

Refer to the *Veritas Storage Foundation Installation Guide*.

Be sure to see a complete list of Veritas products, including the associated installation script names and where to find documentation about installation.

Refer to the *Getting Started Guide*.

You may also be required to configure VVR after the upgrade.

See "Configuring Veritas Volume Replicator" on page 92.

Note: If you have multiple Veritas products, we strongly recommend using the option to upgrade the entire product suite rather than upgrading each product individually. This ensures that upgrade steps are done in the proper order and product interdependencies are met.

To upgrade VVR, perform the following steps in the order presented:

- "Preparing to upgrade using the product installer" on page 125.
- "Upgrading Veritas Volume Replicator using the product installer" on page 126.
- "Restoring the original configuration using the product installer" on page 127.

Preparing to upgrade using the product installer

This section describes upgrade preparation prior to using the product installer.

Note: Use a different set of instructtructions to upgrade an installation that uses VCS Agents for VVR.

See "Preparing for the upgrade when VCS agents are configured" on page 136.

To prepare the upgrade through the product installer

1. Make sure that the disk groups that contain RVGs are at least at disk group version 110.

   ```
   # vxdg list diskgroup
   ```

2. Make sure the size of the SRL volume is greater than 110 MB.

   Refer to the *Veritas Volume Replicator Administrator's Guide*.

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

4. Verify that all the Primary RLINKs are up-to-date on all the hosts.

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

   Note: Do not continue until the Primary RLINKs are up-to-date.

Upgrading Veritas Volume Replicator using the product installer

This section describes how to upgrade using the product installer.

To upgrade the Veritas Volume Replicator with the product installer

1. Start the product installer:

```
# cd disc_path

# ./installer
```

2. Select Install/Upgrade a Product.
3. Select the appropriate product name:
   - If you are upgrading VVR only, use the Veritas Volume Replicator option.
   - If you are upgrading multiple Veritas products, select the appropriate option in the product installer to update all the Veritas products at the same time. Refer to the appropriate installation guide for detailed instructions.

Note: If you have multiple Veritas products, we strongly recommend using the option to upgrade the entire product suite rather than upgrading each product individually. This ensures that upgrade steps are done in the proper order and interdependencies are met.

---

4. The script detects that an existing installation of VVR is present, and handles upgrade tasks.
5. Follow the prompts.

   The script displays the name of the directory used for the upgrade. The upgrade directory is created in `/var/tmp` on the host from which the upgrade procedure was begun. The upgrade directory has the name *`vvr_upgrade_hostname$timestamp`* where the *`hostname`* is the machine being upgraded, and *`$timestamp`* is the same digit sequence as the suffix of the log file created for the current session.

   For example, the directory `/var/tmp/vvr_upgrade_system01126061743` contains the upgrade files for the host system01.

Note: We strongly recommend you back up the upgrade directory created here, because it is used to restore the configuration.

---

6 The script displays the location of the log files.

When the script completes, it displays messages similar to the following:

```
CPI WARNING V-9-11-2246 You have completed upgrading VxVM
on some or all of the systems. Reboot your systems at this time.
```

7 Prior to rebooting, copy the VVRTypes.cf from `/etc/VRTSvcs/conf` to:

```
/etc/VRTSvcs/conf/config.
```

8 When the upgrade completes, the hosts that are being upgraded must be rebooted.

See "Restoring the original configuration using the product installer" on page 127.

Note: If you are upgrading an installation that uses VCS Agents for VVR, do not configure VVR until after you reboot the machine, unfreeze the service groups and restore the original configuration.

See "Unfreezing the service groups" on page 138.

See "Restoring the original configuration when VCS agents are configured" on page 139.

Restoring the original configuration using the product installer

You must configure VVR to restore the original configuration and complete the upgrade. Configuring VVR also starts the VVR processes.

To restore the original configuration through the product installer

1 On all Secondary hosts, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup shrinkto volume_name volume_length
```

where *volume_length* is the length of the volume on the Primary.

2 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

3 Restore the original configuration on each host that has been upgraded, starting with the Secondary hosts. To restore the configuration, configure VVR using one of the following methods:

- Use the Veritas product installer, select Configure an Installed Product, and then select Veritas Volume Replicator.
- Use the installation script `installvvr` with the `-configure` option.

The configuration is restored from the configuration files and scripts that were saved in the upgrade directory during the upgrade session.

See "Configuring Veritas Volume Replicator" on page 92.

4 Restart the applications that were stopped.

After the configuration is restored, the current step can be retried.

If the upgrade fails

If the upgrade fails during the configuration phase, after displaying the VVR upgrade directory, the configuration needs to be restored before the next attempt. Run the scripts in the upgrade directory in the following order to restore the configuration:

```
# restoresrl
# adddcm
# srlprot
# attrlink
# start.rvg
```

After the configuration is restored, the current step can be retried.

Upgrading using VVR upgrade scripts

This section describes the procedure to upgrade to VVR 5.0MP3 using upgrade scripts. Use this method only if you are upgrading in a scenario that is not supported by the Veritas product installer.

Note: We recommend using the Veritas product installer for upgrading VVR when possible.

See Supported upgrade methods for Veritas Volume Replicator.

Upgrading from VVR 4.0 requires a complete reinstallation of the operating system. You must ensure that you have made backups of all data that you want to preserve, before reinstalling the operating system.

The upgrade procedure retains the existing VVR configuration. After upgrading, you can use the existing VVR configuration, without running the `vxinstall` command.

To upgrade VVR, perform the following tasks in the order presented:

- Preparing to upgrade using upgrade scripts
- Upgrading Veritas packages using upgrade scripts
- Restoring the original configuration using upgrade scripts

Preparing to upgrade using upgrade scripts

This section describes how to prepare your configuration for the upgrade using the upgrade scripts:

To prepare your upgrade thorugh the upgrade scripts

1. Make sure that the disk groups that contain RVGs are at least at disk group version 110.

   ```
   # vxdg list diskgroup
   ```

2. Make sure the size of the SRL volume is greater than 110 MB.

   Refer to the *Veritas Volume Replicator Administrator's Guide*.

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

4. Verify that all the Primary RLINKs are up-to-date on all the hosts.

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

   Caution: Do not continue until the Primary RLINKs are up-to-date.

5. Run the vvr_upgrade_start script on all hosts to save the original VVR configuration using the following command:

   ```
   # /disc_path/volume_replicator/scripts/vvr_upgrade_start
   ```

Note: If the vvr_upgrade_start script finds that the SRL size is less than 110 MB, then the script fails and reverts back to the original configuration. It stops with a message that prompts you to modify the SRL size.

Refer to the *Veritas Volume Replicator Administrator's Guide*.

Upgrading Veritas packages using upgrade scripts

This section describes how to upgrade the packages through the upgrade script.

On all hosts on which the upgrade is to be performed

1 If necessary, upgrade the operating system.

Upgrade VxVM from the product disc to overwrite the previous version with version 5.0MP3.

Refer to the *Veritas Storage Foundation Installation Guide*.

2 If you have not rebooted the system, reboot it now using the following command:

```
# shutdown -y -i6 -g0
```

During the reboot process, ignore the following error messages that appear on the Primary console:

```
VxVM VVR vxrlink ERROR V-5-1-3371 Can not recover rlink_name.
rvg_name is in PASSTHRU mode
```

```
VxVM VVR vxrlink ERROR V-5-1-3473 Log header I/O error
```

Also ignore the following error message that appears on the Secondary console:

```
WARNING: VxVM VVR vxio V-5-0-278 Rlink rlink_name is stale and
not replicating
```

3 Upgrade the required and optional packages for VVR. Perform the following tasks in the ordered indicated:

- Remove the old packages.
  See "Removing the VVR packages" on page 200.
- Copy the packages from the Veritas software disk to a temporary directory.

  ```
  # cd /disc_path
  ```

  ```
  # cp -r volume_replicator/pkgs/* /tmp_dir
  ```

- Unzip the package files.

  ```
  # gunzip VRTS*.gz
  ```

- Decompress and extract each package.

  ```
  # tar xf package_name.tar
  ```

- Use the following command to display the list of VVR packages. The packages must be installed in the order shown.

```
# ./installvvr -installpkgs
```

4 Install the new packages using the `rpm` command.

Some configurations may require upgrading an installation with VCS Agents for VVR configured.

See "Unfreezing the service groups" on page 138.

Note: If you have additional Veritas products to upgrade, refer to the product installation guide for a list of packages to upgrade.

Restoring the original configuration using upgrade scripts

This section describes how to complete the upgrade and restore the original configuration using the upgrade scripts.

To restore the original configuration through the upgrade script

1 On all Secondary hosts, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup shrinkto volume_name volume_length
```

where volume_length is the length of the volume on the Primary

2 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

3 Issue the following command on all the hosts to complete the upgrade. If a host contains only Secondary RVGs, we recommend that you first run the following command on that host:

```
# /disc_path/volume_replicator/scripts/vvr_upgrade_finish
```

The vvr_upgrade_finish script upgrades only the SRL, after which, the RVG cannot work with the earlier versions of VxVM or VVR.

4 Restart the applications that were stopped.

Upgrading VVR without disrupting replication

This section describes the upgrade procedure from an earlier version of VVR to the current version of VVR when replication is in progress, assuming that you do not need to upgrade all the hosts in the RDS simultaneously.

You may also need to set up replication between versions.

See "Planning an upgrade from the previous VVR version" on page 32.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade can be performed either on the Primary or on the Secondary. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS. This section includes separate sets of steps, for the Primary upgrade and for the Secondary upgrade.

Note: If you have a cluster setup, you must upgrade all the nodes in the cluster at the same time.

See "Upgrading VVR when VCS agents are configured" on page 133.

Upgrading on the Secondary

Follow these instructions to upgrade the Secondary hosts.

To upgrade the Secondary

1. Stop replication to the Secondary host by initiating a Primary pause using the following command:

   ```
   # vradmin -g diskgroup pauserep local_rvgname
   ```

2. Upgrade from VVR 4.1 or later to VVR 5.0MP3 on the Secondary.

   See "Upgrading Veritas Volume Replicator using the Veritas product installer" on page 124.

3. Resume the replication from the Primary using the following command:

   ```
   # vradmin -g diskgroup resumerep local_rvgname sec_hostname
   ```

Upgrading on the Primary

After you upgrade the Secondary, use the Veritas product installer to upgrade the Primary.

See "Upgrading Veritas Volume Replicator using the Veritas product installer" on page 124.

Note: Reduce application downtime while upgrading by planning your upgrade.

See "Planning an upgrade from the previous VVR version" on page 32.

Upgrading VVR when VCS agents are configured

This section details the procedure for upgrading VVR when VCS agents for VVR are configured:

- Use the Veritas product installer for upgrading VVR, unless you are upgrading in a scenario that is not supported by the Veritas product installer.

  Note: We recommend using the Veritas product installer for upgrading VVR when possible.

  See "Supported upgrade methods for Veritas Volume Replicator" on page 124.

- Use the VVR upgrade scripts only if you are upgrading in a scenario that is not supported by the Veritas product installer.

Prerequisites

The following lists the VVR upgrade prerequisites with VCS agents:

- Make sure the size of the SRL volume is greater than 110 MB.
  Refer to the *Veritas Volume Replicator Administrator's Guide*.

To upgrade VVR when VCS agents for VVR are configured, perform the following tasks in the order presented:

- Freezing the service groups and stopping all the applications
- Preparing for the upgrade when VCS agents are configured
- Upgrading Veritas packages when VCS is present
- Unfreezing the service groups
- Restoring the original configuration when VCS agents are configured

Freezing the service groups and stopping all the applications

This section describes how to freeze the service groups and stop all applications.

To freeze the service groups and stop applications

Perform the following steps for the Primary and Secondary clusters:

1. Log in as the superuser.
2. Make sure that `/opt/VRTS/bin` is in your PATH so that you can execute all the product commands.
3. Because the upgrade requires a reboot, cleanly shut down all applications as follows:
   - OFFLINE all application service groups that do not contain RVG or RVGShared resources, except the ClusterService, cvm and RVGLogowner groups.
   - If the application resources are part of the same service group as an RVG or RVGShared resource, then OFFLINE only the application resources.

   Note: You must also stop any remaining applications not managed by VCS.
4. On any node in the cluster, make the VCS configuration writable:

   ```
   # haconf -makerw
   ```
5. On any node in the cluster, list the groups in your configuration:

   ```
   # hagrp -list
   ```
6. On any node in the cluster, freeze all service groups except the ClusterService group by typing the following command for each group name displayed in the output from step 5.

   ```
   # hagrp -freeze group_name -persistent
   ```

   Note: Write down the list of frozen service groups for future use.
7. On any node in the cluster, save the configuration file (`main.cf`) with the groups frozen:

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

   Note: Continue only after you have performed steps 3 to 7 for each cluster.

8 Display the list of service groups that have RVG resources and the nodes on which each service group is online by typing the following command on any node in the cluster:

```
# hares -display -type RVG -attribute State
 Resource        Attribute         System     Value
 VVRGrp           State            system02   ONLINE
 ORAGrp           State            system02   ONLINE
```

Note: For the resources that are ONLINE, write down the nodes displayed in the System column of the output.

9 Repeat step 1 for each cluster.

10 For private disk groups, determine and note down the hosts on which the disk groups are imported.

See Determining the nodes on which disk groups are online.

11 For shared disk groups, run the following command on any node in the CVM cluster:

```
# vxdctl -c mode
```

Note the master and record it for future use.

Determining the nodes on which disk groups are online

For private disk groups, determine and note down the hosts on which the disk groups containing RVG resources are imported. This information is required for restoring the configuration after the upgrade.

To determine the online disk groups

1. On any node in the cluster, list the disk groups in your configuration, and note down the disk group names listed in the output for future use:

   ```
   # hares -display -type RVG -attribute DiskGroup
   ```

   Note: Write down the list of the disk groups that are under VCS control.

2. For each disk group listed in the output in step 1, list its corresponding disk group resource name:

   ```
   # hares -list DiskGroup=diskgroup Type=DiskGroup
   ```

3. For each disk group resource name listed in the output in step 2, get and note down the node on which the disk group is imported by typing the following command:

   ```
   # hares -display dg_resname -attribute State
   ```

   The output displays the disk groups that are under VCS control and nodes on which the disk groups are imported.

Preparing for the upgrade when VCS agents are configured

If you have configured the VCS agents, it is recommended that you take backups of the configuration files, such as `main.cf` and `types.cf`, which are present in the `/etc/VRTSvcs/conf/config` directory.

To prepare a configuration with VCS agents for an upgrade

1 List the disk groups on each of the nodes by typing the following command on each node:

```
# vxdisk -o alldgs list
```

The output displays a list of the disk groups that are under VCS control and the disk groups that are not under VCS control.

Note: The disk groups that are not locally imported are displayed in parentheses.

2 If any of the disk groups have not been imported on any node, import them. For disk groups in your VCS configuration, you can import them on any node. For disk groups that are not under VCS control, choose an appropriate node on which to import the disk group. Enter the following command on the appropriate node:

```
# vxdg -t import diskgroup
```

3 If a disk group is already imported, then recover the disk group by typing the following command on the node on which it is imported:

```
# vxrecover -bs
```

4 Verify that all the Primary RLINKs are up to date.

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

Note: Do not continue until the Primary RLINKs are up-to-date.

Upgrading Veritas packages when VCS is present

After you have performed the steps in the preceeding sections, upgrade your Veritas products in one of the following ways:

- Using the Veritas product installer
- Using the VVR upgrade scripts

Using the Veritas product installer

Upgrade your Veritas products by selecting the option for the product suite.

See "Upgrading Veritas Volume Replicator using the product installer" on page 126.

Note: You must upgrade VVR on all nodes for the Primary and Secondary cluster, after preparing for your VCS agent configuration.

See "Preparing for the upgrade when VCS agents are configured" on page 136.

Using the VVR upgrade scripts

Use this option to upgrade only if you are upgrading in a scenario that is not supported by the Veritas product installer.

Note: We recommend using the Veritas product installer for upgrading VVR when possible.

See "Supported upgrade methods for Veritas Volume Replicator" on page 124.

Note: To preserve your configuration, you must run the `vvr_upgrade_start` script before upgrading your installation.

To upgrade VVR and preserve the original configuration

1 To preserve the original VVR configuration, run the `vvr_upgrade_start` script on the nodes that are to be upgraded, before upgrading your installation.

```
# /disc_path/foundation_suite/volume_replicator/scripts \
/vvr_upgrade_start
```

2 Upgrade your Veritas products.

See "Using the Veritas product installer" on page 137.

If you do not want to preserve the original VVR configuration, see Preparing to upgrade using the product installer.

Unfreezing the service groups

This section describes how to unfreeze services groups and bring them online.

To unfreeze the service groups

1 On any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

2 Unfreeze all service groups that were frozen in step 6 of the section Preparing for the upgrade when VCS agents are configured by typing the following command on any node in the cluster:

```
# hagrp -unfreeze service_group -persistent
```

3 Save the configuration on any node in the cluster.

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

4 If you are upgrading in a shared disk group environment, bring online the RVGShared groups with the following commands:

```
# hagrp -online RVGShared -sys masterhost
```

5 Bring the respective IP resources online on each node.

See "Preparing for the upgrade when VCS agents are configured" on page 136.

Type the following command on any node in the cluster.

```
# hares -online ip_name -sys system
```

This IP is the virtual IP that is used for replication within the cluster.

6 In shared disk group environment, online the virtual IP resource on the master node that you noted in step 11.

Restoring the original configuration when VCS agents are configured

This section describes how to restore a configuration with VCS configured agents.

Note: Restore the original configuration only after you have upgraded VVR on all nodes for the Primary and Secondary cluster.

To restore the original configuration

1 Import all the disk groups in your VVR configuration.

```
# vxdg -t import diskgroup
```

Each disk group should be imported onto the same node on which it was online when the upgrade was performed. The reboot after the upgrade could result in another node being online; for example, because of the order of the nodes in the AutoStartList. In this case, switch the VCS group containing the disk groups to the node on which the disk group was online while preparing for the upgrade.

```
# hagrp -switch grpname -to system
```

2 Recover all the disk groups by typing the following command on the node on which the disk group was imported in step 1.

```
# vxrecover -bs
```

3 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

4 On all nodes that are Secondary hosts of VVR, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup  shrinkto volume_name volume_length
```

where *volume_length* is the length of the volume on the Primary.

Note: Do not continue until you complete this step on all the nodes in the Primary and Secondary clusters on which VVR is upgraded.

---

5 Restore the configuration according to the method you used for upgrade:

If you upgraded with the VVR upgrade scripts

Complete the upgrade by running the `vvr_upgrade_finish` script on all the nodes on which VVR was upgraded. We recommend that you first run the `vvr_upgrade_finish` script on each node that is a Secondary host of VVR.

Perform the following tasks in the order indicated:

- To run the `vvr_upgrade_finish` script, type the following command:

```
# /disc_path/volume_replicator/scripts/
vvr_upgrade_finish
```

where *disc_path* is the location where the Veritas software disc is mounted.

- Attach the RLINKs on the nodes on which the messages were displayed:

```
# vxrlink -g diskgroup -f att rlink_name
```

If you upgraded with the product installer

Use the Veritas product installer and select Configure an Installed Product. Or use the installation script with the `-configure` option.

See "Restoring the original configuration using the product installer" on page 127.

6 Bring online the RVGLogowner group on the master:

```
# hagrp -online RVGLogownerGrp -sys masterhost
```

7 Start and bring online the cvm group on the remaining host:

```
# hagrp -online cvm -sys slave_host
```

8 Restart the applications that were stopped.

Post-upgrade tasks

The tasks in the following sections must be performed after upgrade, to restore the previous configurations and set up SF 5.0MP3 correctly. Perform the tasks required for the products and features that are relevant to your installation.

Optional configuration steps for Linux

After the upgrade is complete, additional tasks may need performed.

You can perform the following optional configuration steps:

- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.
- If you want to use features of Veritas Storage Foundation 5.0MP3 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.

- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
  See "Upgrading VxFS disk layout versions" on page 143.
  See "Upgrading VxVM disk group versions" on page 146.
- If you upgraded to the current version of Veritas Storage Foundation for DB2, refer to the DB2 configuration procedure.
  See "Configuring Storage Foundation Database Editions" on page 82.

Upgrading to the new repository database for DB2 and Oracle

If you are installing or upgrading Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, you need to either create a new repository database or migrate your old repository database to the new one. To use the `db2ed_update` or `dbed_update` command, you must be the instance owner or database administrator.

To upgrade your repository database

1. Create and configure the new repository database with the sfua_db_config command.

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

2. Migrate your old repository information into the new repository database.

3 If you are upgrading Veritas Storage Foundation for DB2 in a single-host environment, run the `db2ed_update` command.

```
# /opt/VRTS/bin/db2ed_update -D DB2DATABASE
```

If you are upgrading Veritas Storage Foundation for DB2 in a high availability (HA) environment, run the `db2ed_update` command with the `-G` option.

```
# /opt/VRTS/bin/db2ed_update -D DB2DATABASE -G service_group
```

If you are upgrading Veritas Storage Foundation for Oracle in a single-host environment, run the `dbed_update` command.

```
# /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

If you are upgrading Veritas Storage Foundation for Oracle in a high availability (HA) environment, run the `dbed_update` command with the `-G` option.

```
# /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME \
  -G service_group
```

4 After the upgrade, the old repository database will be marked with a hidden file name, such as `/etc/vx/vxdba/.instance_name`, to prevent further updates. If you need to perform an additional upgrade, the file must be removed.

About upgrading disk layout versions

You must upgrade your older disk layout versions to make use of the extended features available in theVeritas File System 5.0MP3 release.

See the *Veritas Storage Foundation Release Notes 5.0MP3* for information on new features.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 7 as described in the following sections.

Upgrading VxFS disk layout versions

Veritas File System 5.0MP3 allows Version 4, 6 and 7 file system disk layouts to be mounted. Disk layout Versions 1, 2, and 5 are not supported by VxFS 5.0MP3 on the Linux platform. All file systems created on VxFS 5.0MP3 use disk layout Version 7 by default.

To determine the disk layout version of a VxFS file system, run the `fstyp_vxfs` command on the file system physical device, for example:

```
# /opt/VRTS/bin/fstyp -v /dev/vx/dsk/rootdg/volname | grep \
    version
    magic a501fcf5  version 7 ctime Thu Jul 31 11:29:31 2004
```

When to use vxfsconvert

You can use the `vxfsconvert` command to convert an unmounted ext2 or ext3 file system to a Veritas file system with disk layout Version 7.

```
# vxfsconvert /device_name
```

See the `vxfsconvert`(1M) and `fsadm_vxfs`(1M) manual pages.

When to use vxupgrade

You can use the `vxupgrade` command to upgrade older VxFS disk layouts to disk layout Version 7 while the file system remains mounted.

```
# vxupgrade -n 7 /mount_point
```

See the `vxupgrade`(1M) and `fsadm_vxfs`(1M) manual pages.

Warning: The contents of intent logs created on a previous disk layout version cannot be used after the disk layout version is upgraded.

Requirements for upgrading to disk layout Version 7

Converting a previous disk layout to a Version 7 disk layout requires adequate free space. The space and time required to complete the upgrade increases with the number of files, extended attributes, and hard links in the file system. Typical maximum space is at least two additional inodes with one block for every inode. Allow at least ten minutes to upgrade for every million inodes in the file system.

Migrating from /etc/vx/vxdba to /var/vx/vxdba for DB2 and Oracle

If you are upgrading Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, you can migrate to `/var/vx/vxdba` to save space under the root partition. Migrating to `/var/vx/vxdba` is optional. However, if you do not perform this migration, you cannot remove any file or directory from `/etc/vx/vxdba` to ensure proper operation. This procedure can be done at any time.

To migrate from /etc/vx/vxdba to /var/vx/vxdba

1 Copy the `/etc/vx/vxdba` directory and contents to `/var/vx/vxdba`.

```
# cp -rp /etc/vx/vxdba /var/vx/vxdba
```

2 Remove /etc/vx/vxdba.

```
# rm -rf /etc/vx/vxdba
```

3 Link the two directories.

```
# ln -s /var/vx/vxdba /etc/vx/vxdba
```

Upgrading CVM protocol and disk group version

If you want to take advantage of the new features in this release, you must upgrade the Veritas Cluster Volume Manager (CVM) protocol version (70), and upgrade to the latest disk group version (140).

Downgrading disk group versions is not supported. If a Veritas cluster is used, the disk group version should match the lowest Volume Manager version installed in the cluster.

After upgrading from Storage Foundation 4.x to 5.0MP3, you must upgrade the version for any existing disk groups which are organized by Intelligent Storage Provisioning (ISP). Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations do not function correctly.

To upgrade CVM protocol and the disk group version

1 To upgrade the CVM protocol version to version 70, enter the following command on the master node:

```
# vxdctl upgrade
```

2 To upgrade the disk group version to 140, enter the command:

```
# vxdg -T 140 upgrade dgname
```

3 After upgrading CVM in a VCS environment, you should run the command:

```
# vxcvmconfig upgrade
```

If this command is not run, you will see a warning in the engine log file, `/opt/VRTSvcs/log/engine_A.log`.

Upgrading VxVM disk group versions

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk groups.

After upgrading from Storage Foundation 4.x to 5.0MP3, you must upgrade any existing disk groups which are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

For 5.0MP3, the Veritas Volume Manager disk group version is the same as it was for the VxVM 5.0 release. Upgrading the disk group version is only required if you upgraded from a version earlier than 5.0.

Use the following command to find the version of a disk group:

```
# vxdg list diskgroup
```

To upgrade a disk group to the current disk group version, use the following command:

```
# vxdg upgrade diskgroup
```

For more information about disk group versions, see the *Veritas Volume Manager Administrator's Guide*.

Updating variables

In `/etc/profile`, update the `PATH` and `MANPATH` variables as needed.

MANPATH could include `/opt/VRTS/man` and PATH `/opt/VRTS/bin`.

Setting the default disk group

In releases prior to Volume Manager 4.0, the default disk group was `rootdg` (the root disk group). For Volume Manager to function, the `rootdg` disk group had to exist and it had to contain at least one disk.

This requirement no longer exists; however, you may find it convenient to create a system-wide default disk group. The main benefit of creating a default disk group is that VxVM commands default to the default disk group. You do not need to use the `-g` option.

You can set the name of the default disk group after installation by running the following command on a system:

```
# vxdctl defaultdg diskgroup
```

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

If you want to confirm that the root disk is encapsulated, enter the command:

```
# vxdg bootdg
```

Upgrading the Array Support Library

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software packages.

You can obtain ASL packages from the following locations:

- The VxVM release package
- The disk array provided by the vendor
- The Symantec Technical Support site
  See "Contacting Technical Support" on page 4.

Changing permissions for Storage Foundation for Databases

After installing the Veritas Storage Foundation 5.0MP3 patches, follow these post-installation steps to ensure Veritas Storage Foundation for Oracle and Veritas Storage Foundation for DB2 commands work correctly.

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

---

To change permissions

1. Change permissions for the following directory, depending on which product you have installed:

   For Veritas Storage Foundation for Oracle:

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

   For Veritas Storage Foundation for DB2:

   ```
   # chmod 550 /opt/VRTSdb2ed
   ```

2. Reset owner and group settings to the appropriate owner and group for the database administrators on your system.

   For example, in Veritas Storage Foundation for Oracle, to change owner to the user oracle and the group dba, run the following command:

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

   In Veritas Storage Foundation for DB2, for example, to change owner to the user db2 and the group db2grp, run the following command:

   ```
   # chown db2:db2grp /opt/VRTSdb2ed
   ```

3. Upgrade the repository.
4. In a standalone instance, run `sfua_db_config` once:

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

   This step completes the upgrade of the repository in a standalone configuration.
5. In a cluster environment, complete the remaining steps.
6. Unconfigure the SFUA repository from the VCS configuration.

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

7. Mount the repository file system manually.
8. Run the repository upgrade command again with no options.

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

Verifying the Veritas Storage Foundation upgrade

Refer to the section about verifying the installation to verify the upgrade.

See "Verifying that the products were installed" on page 177.

Chapter 7

Upgrading Storage Foundation Cluster File System

This chapter includes the following topics:

Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0

This section contains procedures for the Veritas Storage Foundation Cluster File System upgrade.

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

Planning the upgrade

Complete the following tasks in advance of upgrading:

- Review the *Veritas Storage Foundation Release Notes* for any late-breaking information on upgrading your system.
- Be sure that the administrator doing the upgrade has root access and a working knowledge of system administration.
- Schedule sufficient outage time for the upgrade.
- Make sure you have upgraded all the file systems to disk layout Version 6, before you upgrade SFCFS to 5.0MP3. Disk layout Version 7 is the recommended version for SFCFS 5.0MP3.
  See `vxupgrade`(1M), `vxconvert`(1M), and `fsadm`(1M) manual pages.
  See the *Veritas File System Administrator's Guide*.
- Verify all the file systems are working fine and data is intact.
  See the `cfsmount`(1M) manual page.

Preparing the system and backing up files before upgrading

Before upgrading an installed Veritas Storage Foundation Cluster File System, 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.
- Back up the configuration files.

```
# cp /etc/VRTSvcs/conf/ backupdirectory/
```

Upgrade paths for Veritas Storage Foundation Cluster File System 5.0MP3

The only supported upgrade paths to Storage Foundation Cluster File System 5.0MP3 are from Storage Foundation Cluster File System 4.1 or 5.0, including all release packs and maintenance packs. Before upgrading to Storage Foundation Cluster File System 5.0MP3, upgrade the Linux system to a version that is supported for this release.

See "Supported Linux operating systems" on page 38.

Table 7-1 shows the upgrade paths for Storage Foundation on Linux.

Table 7-1 Upgrade paths for Storage Foundation Cluster File System on Linux

Storage Foundation Cluster File System version	Upgrade procedure
prior to 4.1	Upgrade to Storage Foundation Cluster File System 5.0MP3 is not supported. You must uninstall Storage Foundation Cluster File System using the procedure in the Storage Foundation Cluster File System Installation Guide for your version, then install Storage Foundation Cluster File System5.0MP3.
SFCFS 4.1, including Maintenance Packs and Rolling Patches **Note:** To upgrade RHEL5 or SLES10 systems from 4.1MP4 to 5.0MP3, you must first upgrade 4.1MP4 to 4.1MP4RP2.	Upgrade Storage Foundation using the procedure: See "Overview of procedures" on page 154.
SFCFS 5.0, including Maintenance Packs and Rolling Patches	Upgrade to 5.0MP3 using the procedure: See "Overview of procedures" on page 154.

Overview of procedures

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

The upgrade procedures apply to both the phased and full upgrade procedures unless otherwise noted. Occasionally, steps differ between the two procedures. Screen output is also common between both procedures unless otherwise noted.

Note: Both procedures automatically uninstall the previous version of the software.

Phased upgrade

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.

The stages of the phased upgrade procedure are:

- Select two or more nodes to upgrade.
- Install the new version.
- Shut down VCS on remaining non-upgraded nodes and ensuring the file systems are clean.
- Reboot the upgraded nodes.
- Install the new version on each remaining node and reboot them.

Note: A phased upgrade should not be performed from one of the nodes in the cluster.

Full upgrade

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:

- Ensuring the file systems are clean.
- Install the new version.

Ensuring the file systems are clean (full only)

Before upgrading to SFCFS 5.0MP3, ensure that the file systems are clean. To ensure that the logs have been replayed and the file systems are marked clean:

To ensure the file systems are clean

1. Log in as superuser onto any node in the cluster.
2. Offline the group on each node of the cluster:

```
# hagrp -offline group -sys system01
# hagrp -offline group -sys system02
# hagrp -offline group -sys system03
# hagrp -offline group -sys system04
```

   where *group* is the VCS service group that has the CVMVolDg and CFSMount resource.

   Repeat this step for each SFCFS service group.
3. Check and repair each file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

   Repeat this step for each file system.

Performing the upgrade (phased or full)

This section describes how to upgrade for phased or full.

If you are performing a phased upgrade, select one or more nodes to upgrade.

To perform the upgrade

1. Log in as superuser.
2. Insert the appropriate media disc per your distribution and architecture 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 Verify there are no VxFS file systems mounted on the nodes being upgraded:

```
# mount -t vxfs
```

If any VxFS file systems are mounted, offline the group on each node of the cluster:

```
# hagrp -offline group -sys system01
# hagrp -offline group -sys system02
# hagrp -offline group -sys system03
# hagrp -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFSMount resource.

Repeat this step for each SFCFS service group.

6 If performing a phased upgrade, start the installation procedure from the node selected in the cluster. In the following example the phased upgrade is performed on one node in a four-node cluster.

If upgrading SFCFS for Oracle RAC, replace `installsfcfs` with `installsfcfsrac`.

Ensure that the HAD daemon of VCS is running on all nodes selected for the upgrade. Enter the following command, and then press **Return**.

```
# ./installsfcfs system01 system02
```

If performing a full upgrade, start the installation from any node in the cluster. Enter the following command, and then press **y** to upgrade the cluster configuration.

```
# ./installsfcfs
```

7 Press **Return** to begin installing infrastructure packages.

8 Press **Return** to begin license verification.

9 Press **Return** to begin the uninstall. The uninstall utility checks the system's uninstallation requirements and identifies packages, patches and dependencies.

The output can be extensive and exceed 100 lines, depending on your configuration.

10 If VCS is running you are prompted to upgrade and reset the VCS password. To ensure security, passwords are not written to the install log.

```
installsfcfs must now make configuration updates and stop
    the cluster before upgrading VCS packages.

    Are you ready to begin the Cluster Server upgrade at this
    time? [y,n,q]
```

11 Enter **y**.

12 At the prompt, enter your new password.

13 Reenter your new password.

14 Output shows information that Cluster Server must be started on a running system. Enter **y** to continue.

15 Press **Return** to begin removing the previous packages and installing the new.

16 Press **Return** again for summary information about logs and reboots.

Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.

17 If performing a phased upgrade, proceed to shut down VCS.

See "Shutting down VCS (phased only)" on page 157.

If performing a full upgrade, proceed to updating the configuration.

See "Updating the configuration and confirm startup (phased or full)" on page 159.

Shutting down VCS (phased only)

Shutdown VCS on remaining nodes that are not being upgraded to preventing them from rejoining the cluster.

To shut down the cluster

1 Separate the nodes that are not being upgraded from those that are.

2 Check to see if there are frozen CVM and SFCFS groups, enter

```
# /opt/VRTSvcs/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

- Make configuration read/write:

  ```
  # /opt/VRTSvcs/bin/haconf -makerw
  ```

- Unfreeze the group:

  ```
  # /opt/VRTSvcs/bin/hagrp -unfreeze group -persistent
  ```

- Save the configuration:

  ```
  # /opt/VRTSvcs/bin/haconf -dump -makero
  ```

3 Offline the group on each of the remaining nodes of the cluster:

```
# hagrp -offline group -sys system01
# hagrp -offline group -sys system02
# hagrp -offline group -sys system03
# hagrp -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFSMount resource.

Repeat this step for each SFCFS service group.

4 Check and repair each file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

Repeat this step for each file system.

5 On each node that is not being upgraded, shutdown VCS, enter:

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

6 On each node that is not being upgraded, stop ODM and GMS if you are running SFCFS for Oracle RAC, enter:

```
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
```

7 On each node that is not being upgraded, shutdown CFS, enter

```
# /etc/init.d/vxfen stop
# /opt/VRTS/bin/fsclustadm cfsdeinit
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

8 Proceed to updating the configuration.

See "Updating the configuration and confirm startup (phased or full)" on page 159.

Updating the configuration and confirm startup (phased or full)

Perform the following steps on each upgraded node.

To update the configuration and confirm startup

1 Remove the `/etc/VRTSvcs/conf/config/.stale` file, if it exists.

```
# rm -f /etc/VRTSvcs/conf/config/.stale
```

2 Reboot the upgraded nodes.

```
# reboot
```

3 After the nodes reboot, verify that LLT is running:

```
# lltconfig
  LLT is running
```

4 Verify GAB is configured:

```
# gabconfig -l | grep 'Driver.state' | \
  grep Configured
  Driver state : Configured
```

5 Verify VxVM daemon is started and enabled:

```
# /opt/VRTS/bin/vxdctl mode
  mode: enabled
```

6 Confirm all upgraded nodes are in a running state.

```
# /opt/VRTSvcs/bin/hasys -state | grep RUNNING | \
  /usr/bin/wc -1
  1
```

7 After the configuration is complete, the CVM and SFCFS groups may come up frozen. To find out the frozen CVM and SFCFS groups, enter the following command:

```
# /opt/VRTS/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

- Make the configuration read/write.

  ```
  # /opt/VRTS/bin/haconf -makerw
  ```

- Unfreeze the group.

  ```
  # /opt/VRTS/bin/hagrp -unfreeze group -persistent
  ```

- Save the configuration.

  ```
  # /opt/VRTS/bin/haconf -dump -makero
  ```

8 If the CVM and SFCFS groups are offline, bring them online on all systems.

```
# /opt/VRTS/bin/hagrp -online group -sys system01
# /opt/VRTS/bin/hagrp -online group -sys system02
```

If the SFCFS service groups do not come online then your file system could be dirty.

See "Making the file systems clean" on page 162.

9 If performing a phased upgrade, upgrade the remaining nodes.

See "Upgrading remaining nodes (phased only)" on page 160.

If performing a full upgrade, proceed to upgrade the remaining nodes.

See "Upgrading remaining nodes (full only)" on page 162.

Upgrading remaining nodes (phased only)

This section describes how to upgrade the remaining nodes.

To upgrade remaining nodes

1. Log in as superuser.
2. Insert the appropriate media disc per your distribution and architecture into your system's DVD-ROM drive.
3. Change to the `storage_foundation_cluster_file_system` directory:

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

4. Install SFCFS on the remaining nodes.

   If upgrading SFCFS for Oracle RAC, replace `installsfcfs` with `installsfcfsrac`.

   ```
   # ./installsfcfs system03 system04
   ```

   When upgrading the remaining nodes, you may be prompted that Cluster Server should have been running during the initial upgrade.
5. Press **y** to continue and follow all the prompts.
6. Check cluster status. Type:

   ```
   # hastatus -summary
   ```

7. If you are configuring SFCFS for a fenced environment.

   See the *Veritas Cluster Server Administrator's Guide*.

   See the *Veritas Storage Foundation Cluster File System Administrator's Guide*.
8. To verify the cluster protocol version, enter the following command that displays the same on a running node:

   ```
   # /opt/VRTS/bin/vxdctl protocolversion
   ```

   If the cluster protocol version is less than 70, then it needs to be upgraded to 70 for SFCFS 5.0MP3.

   The cluster protocol version can only be updated on the master node.

   Upgrade the entire cluster using the following command on the master node.

   ```
   # /opt/VRTS/bin/vxdctl upgrade
   ```

9. Type the following command on one upgraded node to enable membership:

   ```
   # gabconfig -xc
   ```

Upgrading remaining nodes (full only)

This section describes how to upgrade the remaining nodes.

To upgrade remaining nodes

1. If you are configuring SFCFS for a fenced environment.

   See the *Veritas Cluster Server Administrator's Guide*.

   See the *Veritas Storage Foundation Cluster File System Administrator's Guide*.

2. To verify the cluster protocol version, enter the following command that displays the same on a running node:

   ```
   # /opt/VRTS/bin/vxdctl protocolversion
   ```

   If the cluster protocol version is less than 70, then it needs to be upgraded to 70 for SFCFS 5.0MP3.

   The cluster protocol version can only be updated on the master node.

   Upgrade the entire cluster using the following command on the master node.

   ```
   # /opt/VRTS/bin/vxdctl upgrade
   ```

3. Type the following command on one upgraded node to enable membership:

   ```
   # gabconfig -xc
   ```

Making the file systems clean

If you upgrade to SFCFS 5.0MP3 and the file systems are dirty, you have to deport the shared disk group and import it as non-shared. After the import, run `fsck`. `fsck` should succeed. Then deport the disk group and import it back as shared.

To make the file systems clean

1 Log in as superuser onto the CVM master node.

2 If performing a full upgrade, offline the group on all the nodes in the cluster:

```
# hagrp -offline group -sys system01
# hagrp -offline group -sys system02
# hagrp -offline group -sys system03
# hagrp -offline group -sys system04
```

If performing a phased upgrade, offline the group:

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

where *group* is the VCS service group that has the CVMVolDg and CFSMount resource.

3 Deport the disk group:

```
# vxdg deport diskgroup
```

where *diskgroup* is the SFCFS disk group.

4 Import the disk group:

```
# vxdg -C import diskgroup
```

5 Start all the volumes in the disk group:

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

6 Check and repair each file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

Repeat this step for each file system.

7 Deport the disk group:

```
# vxdg deport diskgroup
```

8 Import the disk group:

```
# vxdg -s import diskgroup
```

9 Start all the volumes in the disk group:

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

10 If performing a full upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01
# hares -clear resource -sys system02
# hares -clear resource -sys system03
# hares -clear resource -sys system04
```

If performing a phased upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01
# hares -clear resource -sys system02
```

11 If performing a full upgrade, online the group on all the nodes in the cluster:

```
# hagrp -online group -sys system01
# hagrp -online group -sys system02
# hagrp -online group -sys system03
# hagrp -online group -sys system04
```

If performing a phased upgrade, online the group:

```
# hagrp -online group -sys system01
# hagrp -online group -sys system02
```

Upgrading Storage Foundation Cluster File System to 5.0MP3 on a Storage Foundation 5.0 system

Use this procedure to upgrade to 5.0MP3 from 5.0, or from an earlier 5.0 Maintenance Pack.

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 a Maintenance Pack" on page 165.

See "Full upgrade for a Maintenance Pack" on page 172.

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.

You should also review the *Veritas Storage Foundation Release Notes* for important release information.

Caution: A phased upgrade procedure results in a system PANIC on configurations where LLT is configured over UDP. This issue is fixed in 5.0 MP1. This issue is specific to configurations where LLT is configured over UDP and not present in usual LLT Ethernet configurations. The full upgrade procedure should be used for upgrading from SFCFS 5.0 or SFCFS 5.0 RP1 to SFCFS 5.0 MP1 on configurations where LLT is configured over UDP.

Preparing to upgrade to the Maintenance Pack

If you are upgrading an installed Veritas Storage Foundation 5.0 version or from an earlier 5.0 Maintenance Pack, 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 a Maintenance Pack

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 a Maintenance Pack

This section describes how to perform a phased upgrade for a Maintenance Pack.

To freeze service group operations and stop cluster failover

1. Log in as superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so that you can execute all product commands.
3. From any node in the cluster, make the cluster configuration writable.

   ```
   # haconf -makerw
   ```

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

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

5. Make the configuration read-only

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

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 a Maintenance Pack

After the selected group of nodes is offline, the Veritas Storage Foundation Cluster File System software can be upgraded, using `installmp` for the Maintenance Pack.

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

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

   ```
   # hastop -local
   ```

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

   ```
   # df -v /
   ```

3. 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.

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

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

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

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

6 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.

7 Check if the VEA service is running:

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

8 If the VEA service is running, stop it:

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

9 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.

10 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

11 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 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. Change to the top-level directory on the disc:

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

4. 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
     ```

5. After the initial system checks are complete, press **Return** to start the requirement checks.
6. 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.
7. When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
8. (Optional) If you are going to upgrade your operating system, then upgrade your operating system, and patch it to a kernel version.
9. Reboot the system (or systems).

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 to restart cluster failover services.

Upgrade the second group of nodes.

To take the second group of nodes offline

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

```
# hastop -local
```

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 170.

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 software on the second group of nodes.

   See "Upgrading the Veritas Storage Foundation Cluster File System software to a Maintenance Pack" on page 167.

2. Then bring the second group of nodes online.

   See "Bringing the upgraded group of nodes online" on page 170.

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

```
# haconf -makerw
```

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

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

6 Make the configuration read-only:

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

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

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

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

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

9 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
```

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

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

11 Check if the VEA service was restarted:

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

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

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

Full upgrade for a Maintenance Pack

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:

- Freeze service group operations and stop cluster failover operations.
- 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 a Maintenance Pack

This section describes how to perform a full upgrade to a Maintenance Pack.

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 a Maintenance Pack

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 Change to the top-level directory on the disc:

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

4 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
  ```

5. After the initial system checks are complete, press **Return** to start the requirement checks.
6. 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.
7. When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
8. (Optional) If you are going to upgrade your operating system, then upgrade your operating system, and patch it to a kernel version.
9. Shut down and reboot the system.

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
```

Chapter 8

Verifying the Storage Foundation installation

This chapter includes the following topics:

Verifying that the products were installed

Verify that the Veritas Storage Foundation products are installed.

Use the following sections to further verify the product installation.

Installation log files

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

- Installation log file
- Response file
- Summary file

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

Using the installation log file

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

Using the response file

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

Using the summary file

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

Checking Volume Manager processes

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

To confirm that key Volume Manager processes are running

- Type the following command:

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

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

Verifying the configuration files for Storage Foundation Cluster File System

You can inspect the contents of the configuration files that were installed and modified after a successful installation process. These files reflect the configuration based on the information you supplied.

To verify the configuration files

1. Log in as superuser to any system in the cluster.
2. Set up your environment `PATH` variable.

```
# export PATH=$PATH:/sbin:/usr/sbin:/opt/VRTS/bin
```

Low Latency Transport configuration files

The following files are required by the VCS communication services for Low Latency Transport (LLT).

/etc/llthosts

The file `llthosts`(4M) is a database, containing one entry per system, that links the LLT system ID (in the first column) with the LLT host name. This file is identical on each system in the cluster.

For example, the file `/etc/llthosts` contains entries that resemble:

```
0    system01
1    system02
```

/etc/lltab

The file `lltab`(4M) contains information that is derived during installation and used by the utility `lltconfig`(1M). After installation, this file lists the network links that correspond to the specific system.

For example, the file `/etc/lltab` contains entries that resemble:

```
set-node system01
set-cluster 100

link lan1 eth:1 - ether - -
link lan2 eth:2 - ether - -
```

The first line identifies the local system name. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines, beginning with the `link` command, identify the two network cards used by the LLT protocol.

See the `lltab`(4M) manual page.

Checking Low Latency Transport operation

Use the `lltstat` command to verify that links are active for LLT. This command returns information about the links for LLT for the system on which it is typed. See the `lltstat`(1M) manual page.

In the following example, `lltstat -n` is typed on each system in the cluster.

To check LLT operation

1 Log into system01.

```
# lltstat -n
```

Output resembles:

```
   LLT  node  information:
    Node                          State             Links
*   0    system01                 OPEN                  2
    1    system02                 OPEN                  2
```

2 Log into system02.

```
  # lltstat -n
```

Output resembles:

```
   LLT  node  information:
    Node                          State             Links
    0    system01                 OPEN                  2
*   1    system02                 OPEN                  2
```

Each system has two links and that each system is in the OPEN state. An asterisk (*) denotes the system on which the command is typed.

With LLT configured correctly, the output of `lltstat -n` shows all of the systems in the cluster and two links for each system. If the output shows otherwise, you can use the verbose option of `lltstat`. For example, type `lltstat -nvv | more` on a system to view additional information about LLT. In the following example, `lltstat -nvv | more` is typed on a system in a two-node cluster.

3 Log into system01.

```
# lltstat -nvv | more
```

Output resembles:

```
Node        State       Link       Status      Address

*0 system01 OPEN      lan1       UP          08:00:20:93:0E:34
                       lan2       UP          08:00:20:93:0E:34
 1 system02 OPEN       lan1       UP          08:00:20:8F:D1:F2
                       lan2       DOWN        08:00:20:8F:D1:F2
 2          CONNWAIT
                       lan1       DOWN
                       lan2       DOWN
 .
 .
 .
 31         CONNWAIT
                       lan1       DOWN
                       lan2       DOWN
```

The output lists 32 nodes. It reports on the two cluster nodes, system01 and system02, plus non-existent nodes. For each correctly configured system, the information shows a state of OPEN, a status for each link of UP, and an address for each link. However, in the example above, the output shows that for node system02, the private network may have failed earlier, or the information in `/etc/lltab` may be incorrect.

To obtain information about the ports open for LLT, type `lltstat -p` on any system. In the following example, `lltstat -p` is typed on one system in the cluster.

4 Log into system01.

```
# lltstat -p
```

Output resembles:

```
LLT port information:
 Port      Usage      Cookie
 0         gab        0x0
           opens:           0 1 3 4 5 6 7 8 9 10 11 12 13...
           connects:        0 1
```

The two systems with node ID's 0 and 1 are connected.

See "/etc/llthosts" on page 179.

Group Membership and Atomic Broadcast configuration files

The following files are required by the VCS communication services for Group Membership and Atomic Broadcast (GAB).

/etc/gabtab

After installation, the file `/etc/gabtab` contains a `gabconfig`(1M) command that configures the GAB driver for use.

The file `/etc/gabtab` contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

where the `-c` option configures the driver for use and `-nN` specifies that the cluster will not be formed until at least *N* systems are ready to form the cluster. *N* is the number of systems in the cluster.

Checking Group Membership and Atomic Broadcast operation

This section describes how to check GAB operation.

To check GAB operation

- Enter the following command on each node in the cluster.

  ```
  # gabconfig -a
  ```

  If GAB is operational, the following output displays with GAB port membership information:

  ```
  GAB Port Memberships
  ===============================================================
  Port a gen 1bbf01 membership 01
  Port b gen 1bbf06 membership 01
  Port f gen 1bbf0f membership 01
  Port h gen 1bbf03 membership 01
  Port v gen 1bbf0b membership 01
  Port w gen 1bbf0d membership 01
  ```

  If GAB is not operational, the following output display with no GAB port membership information:

  ```
  GAB Port Memberships
  ===============================================================
  ```

  See the *Veritas Cluster Server User's Guide*.

Checking cluster operation

This section describes how to check cluster operation.

To check cluster operation

1 Enter the following command on any system:

```
# hastatus -summary
```

The output for an SFCFS HA installation resembles:

```
-- SYSTEM STATE
-- System               State                Frozen

A  system01             RUNNING              0
A  system02             RUNNING              0

-- GROUP STATE
-- Group       System      Probed AutoDisabled   State

B  cvm         system01    Y      N              ONLINE
B  cvm         system02    Y      N              OFFLINE
```

If the State value is running, VCS is successfully installed and running on that node. The group state lists the cvm group, which is online on system01 and offline on system02.

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

See the *Veritas Cluster Server User's Guide*.

2 Enter the following command on any systems:

```
# hasys -display
```

See the *Veritas Cluster Server User's Guide*.

For more information on the `hasys -display` command, see the `hasys`(1M) manual page.

The example shows the output of system01. The list continues with similar information for system02 (not shown) and any other systems in the cluster. The output should be similar on each system.

```
#System    Attribute           Value
system01    AgentsStopped      0
system01    AvailableCapacity  100
system01    CPUBinding         BindTo    NONE    CPUNumber       0
system01    CPUUsage           0
system01    CPUUsageMonitoring Enabled   0       ActionThreshold 0       Action
system01    Capacity           100
system01    ConfigBlockCount   172
```

```
system01    ConfigCheckSum     18170
system01    ConfigDiskState    CURRENT
system01    ConfigFile         /etc/VRTSvcs/conf/config
system01    ConfigInfoCnt      0
system01    ConfigModDate      Wed Mar 19 16:07:22 2008
system01    ConnectorState     Down
system01    CurrentLimits
system01    DiskHbStatus
system01    DynamicLoad        0
system01    EngineRestarted    0
system01    EngineVersion      5.0.30.0
system01    Frozen             0
system01    GUIIPAddr
system01    LLTNodeId          0
system01    LicenseType        PERMANENT_SITE
system01    Limits
system01    LinkHbStatus       ge0       UP      ge1     UP
system01    LoadTimeCounter    0
system01    LoadTimeThreshold  600
system01    LoadWarningLevel   80
system01    NoAutoDisable      0
system01    NodeId             0
system01    OnGrpCnt           1
system01    ShutdownTimeout    120
system01    SourceFile         ./main.cf
system01    SysInfo            Solaris:system01,Generic_127111-08,5.10,sun4
system01    SysName            system01
system01    SysState           RUNNING
system01    SystemLocation
system01    SystemOwner
system01    TFrozen            0
system01    TRSE               0
system01    UpDownState        Up
system01    UserInt            0
system01    UserStr
system01    VCSFeatures        NONE
system01    VCSMode            VCS_CFS_VRTS
```

Verifying agent configuration for Storage Foundation Cluster File System

This section describes how to verify the agent configuration.

To verify the agent configuration

- Enter the cluster status command from any node in the cluster:

```
# cfscluster status
```

Output resembles:

```
Node             :  system01
Cluster Manager  :  running
CVM state        :  running
No mount point registered with cluster configuration

Node             :  system02
Cluster Manager  :  running
CVM state        :  running
No mount point registered with cluster configuration
```

Synchronizing time on Cluster File Systems

SFCFS requires that the system clocks on all nodes are synchronized using some external component such as the Network Time Protocol (NTP) daemon. If the nodes are not in sync, timestamps for change (`ctime`) and modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

Configuring VCS for Storage Foundation Cluster File System

Configuring VCS means conveying to the VCS engine the definitions of the cluster, service groups, resources, and resource dependencies. VCS uses two configuration files in a default configuration:

- The `main.cf` file defines the entire cluster.
- The `types.cf` file defines the resource types.

By default, both files reside in the directory `/etc/VRTSvcs/conf/config`. Additional files similar to `types.cf` may be present if agents have been added, such as `Oracletypes.cf`.

In a VCS cluster, the first system to be brought online reads the configuration file and creates an internal (in-memory) representation of the configuration. Systems brought online after the first system derive their information from systems running in the cluster. You must stop the cluster while you are modifying the files

from the command line. Changes made by editing the configuration files take effect when the cluster is restarted. The node on which the changes were made should be the first node to be brought back online.

main.cf file

The VCS configuration file `main.cf` is created during the installation procedure. After installation, the `main.cf` file contains the base definitions of the cluster and its nodes. Additionally, the file `types.cf` listed in the include statement defines the bundled agents for VCS resources.

See the *Veritas Cluster Server User's Guide*.

A typical VCS configuration file for SFCFS file resembles:

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"

cluster sfcfs_1 (
        HacliUserLevel = COMMANDROOT
        )

system thor150 (
        )

system thor151 (
        )

group cvm (
        SystemList = { thor150 = 0, thor151 = 1 }
        AutoFailOver = 0
        Parallel = 1
        AutoStartList = { thor150, thor151 }
        )

        CVMCluster cvm_clus (
                CVMClustName = sfcfs_1
                CVMNodeId = { thor150 = 0, thor151 = 1 }
                CVMTransport = gab
                CVMTimeout = 200
                )

        CVMVxconfigd cvm_vxconfigd (
```

```
        Critical = 0
        CVMVxconfigdArgs = { syslog }
        )

cvm_clus requires cvm_vxconfigd


// resource dependency tree
//
//      group cvm
//      {
//      CVMCluster cvm_clus
//          {
//          CVMVxconfigd cvm_vxconfigd
//          }
//      }
```

Storage Foundation Cluster File System HA Only

If you configured VCS Cluster Manager (Web Console), a service group, "ClusterService," was created that includes IP, Process, and Notifier resources. These resources were configured according to information you provided during the installation procedure. A resource dependency was also created.

Veritas Cluster Server application failover services

If you installed SFCFS HA, you can begin implementing the application monitoring failover services provided by the Veritas Cluster Server. Information about setting up VCS services is beyond the scope of this document.

See the *Veritas Cluster Server* documentation.

Chapter 9

Uninstalling Storage Foundation

This chapter includes the following topics:

About removing Veritas Storage Foundation

This section covers steps to uninstall Veritas Storage Foundation, Veritas Storage Foundation for DB2, and Veritas Storage Foundation for Oracle.

Only users with superuser privileges can uninstall Veritas Storage Foundation, Veritas Storage Foundation for DB2, or Veritas Storage Foundation for Oracle.

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

Dropping the repository database for DB2 and Oracle

When uninstalling Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, drop the repository database. If you want to recreate the repository database, you can drop the existing repository database using these steps.

To drop the repository database in a stand-alone configuration

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

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

```
# /opt/VRTSdbcom/config/sfua_rep_mount start
```

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

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

To drop the repository database in a DB2 or Oracle cluster or Oracle RAC configuration

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

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

2 Import the repository database disk group.

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

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

```
# /opt/VRTSdbcom/config/sfua_rep_mount start
```

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

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

Shutting down cluster operations

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

To take all service groups offline and shutdown VCS

- Use the `hastop` command as follows:

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

Warning: Do not use the `-force` option when executing `hastop`. This will leave all service groups online and shut down VCS, causing undesired results during uninstallation of the packages.

Removing VxFS file systems

The VxFS package cannot be removed if there are any mounted VxFS file systems. Unmount all VxFS file systems before removing the package. After you remove the VxFS package, VxFS file systems are not mountable or accessible until another VxFS package is installed. It is advisable to back up VxFS file systems before installing a new VxFS package. If VxFS will not be installed again, all VxFS file systems must be converted to a new file system type.

To remove VxFS file systems

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

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

2 Make backups of all data on the file systems that you wish to preserve, or recreate them as non-VxFS file systems on non-VxVM volumes or partitions.

3 Unmount all Storage Checkpoints and file systems:

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

4 Comment out or remove any VxFS file system entries from the `/etc/fstab` file.

Removing rootability

Perform this procedure if you configured rootability by encapsulating the root disk.

To remove rootability

1 Check if the system's root disk is under VxVM control by running this command:

```
# df -v /
```

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

2 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
```

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

---

3 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.

Moving volumes to disk partitions

All volumes must be moved to disk partitions.

This can be done using one of the following procedures:

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

Moving volumes onto disk partitions using VxVM

Use the following procedure to move volumes onto disk partitions.

To move volumes onto disk partitions

1. Evacuate disks using the `vxdiskadm` program, VEA, or the `vxevac` script. You should consider the amount of target disk space required for this before you begin.

   Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.

2. Remove the evacuated disks from VxVM control using the following commands:

   ```
   # vxdg -g diskgroup rmdisk disk _media_name
   # vxdisk rm disk_access_name
   ```

3. Decide which volume to move first. If the volume to be moved is mounted, unmount it.
4. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume is synced.
5. Create a partition on free disk space of the same size as the volume. If there is not enough free space for the partition, a new disk must be added to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.

6 Copy the data on the volume onto the newly created disk partition using a command similar to the following:

```
# dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/sdb2
```

where sdb is the disk outside of VxVM and 2 is the newly created partition on that disk.

7 Replace the entry for that volume (if present) in /etc/fstab with an entry for the newly created partition.

8 Mount the disk partition if the corresponding volume was previously mounted.

9 Stop the volume and remove it from VxVM using the following commands:

```
# vxvol -g diskgroup -f stop volume_name
# vxedit -g diskgroup -rf rm volume_name
```

10 Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

```
# vxprint -F "%sdnum" disk_media_name
```

11 If the output is not 0, there are still some subdisks on this disk that must be subsequently removed. If the output is 0, remove the disk from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk _media_name
# vxdisk rm disk_access_name
```

12 The free space now created can be used for adding the data in the next volume to be removed.

13 After all volumes have been converted into disk partitions successfully, reboot the system. After the reboot, none of the volumes should be open. To verify that none of the volumes are open, use the following command:

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

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

Shutting down Veritas Volume Manager

Use the following procedure to shut down Veritas Volume Manager.

To shut down Veritas Volume Manager

- Enter the `vxdctl` and `vxiod` commands as follows:

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

Uninstalling Veritas Storage Foundation packages

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

Not all packages may be installed on your system depending on the choices that you made when you installed the software.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

To shut down and remove the installed Veritas Storage Foundation packages

1. In a stand-alone configuration, if you are uninstalling Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, stop the repository database and unmount the database repository volume.

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

2. In an HA configuration, stop VCS processes on either the local system or all systems.

   To stop VCS processes on the the local system:

   ```
   # hastop -local
   ```

   To stop VCS processes on all systems:

   ```
   # hastop -all
   ```

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

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

   For Veritas Storage Foundation

   ```
   # ./uninstallsf
   ```

   For Veritas Storage Foundation for DB2

   ```
   # ./uninstallsfdb2
   ```

   For Veritas Storage Foundation for Oracle

   ```
   # ./uninstallsfora
   ```

Uninstalling Storage Foundation Cluster File System

If you need to uninstall SFCFS software. Use the `uninstallsfcfs` script.

To uninstall SFCFS HA

1. Log in as superuser.

   Do not use the `hastop -force` command to stop VCS.

2. Change directory to `/opt/VRTS/install`:

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

3 Run the `uninstallsfcfs` command to uninstall SFCFS. The `uninstallsfcfs` script uses ssh to communicate with remote nodes as default:

```
# ./uninstallsfcfs
```

If you want to use rsh you must specify on the command line:

```
# ./uninstallsfcfs -rsh
```

4 Enter the system names to uninstall SFCFS.

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

5 Enter **y** to uninstall SFCFS.

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

Uninstalling the VCS agents for VVR

To uninstall the VCS Agents for VVR, you must first disable the agents.

If VCS Agents for VVR are not installed on your system, go to Uninstalling Veritas Volume Replicator (VVR).

Disabling the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

To disable the agents

1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
# hagrp -state service_group -sys system_name
```

If none of the service groups is online, skip to 3.

2 If the service group is online, take it offline.

To take the service group offline without bringing it online on any other system in the cluster, enter:

```
# hagrp -offline service_group -sys system_name
```

3 Stop the agent on the system by entering:

```
# haagent -stop agent_name -sys system_name
```

When you get the message `Please look for messages in the log file`, check the file `/var/VRTSvcs/log/engine_A.log` for a message confirming that each agent has stopped.

You can also use the `ps` command to confirm that the agent is stopped.

4 Remove the system from the `SystemList` of the service group. If you disable the agent on all the systems in the `SystemList`, you can also remove the service groups and resource types from the VCS configuration.

Read information on administering VCS from the command line.

Refer to the *Veritas Cluster Server User's Guide*.

Uninstalling Veritas Volume Replicator (VVR)

This section describes how to uninstall Volume Replicator.

Note: If you are upgrading Veritas Volume Replicator, do not remove the Replicated Data Set, but only remove the VVR packages.

See Removing the VVR packages.

Uninstalling Veritas Volume Replicator (VVR) involves performing the following tasks in the order indicated:

- Removing the Replicated Data Set
- Removing the VVR packages

Refer to the *Veritas Volume Replicator Administrator's Guide*.

Removing the Replicated Data Set

This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

To remove the Replicated Data Set

1 Verify that all RLINKs are up-to-date:

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

If the Secondary is not required to be up-to-date, proceed to 2 and stop replication using the `-f` option with the `vradmin stoprep` command.

2 Stop replication to the Secondary by issuing the following command on any host in the RDS:

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

```
# vradmin -g diskgroup stoprep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

3 Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

```
# vradmin -g diskgroup delsec local_rvgname sec_hostname
```

The argument local_rvgname is the name of the RVG on the local host and represents its RDS.

The argument sec_hostname is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

4 Remove the Primary from the RDS by issuing the following command on the Primary:

```
# vradmin -g diskgroup delpri local_rvgname
```

When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

The RDS is removed. Go on to uninstalling Volume Manager to uninstall VVR.

5 If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

```
# vxedit -r -g diskgroup rm srl_name
```

6 Uninstall the VVR packages.

See Removing the VVR packages.

Removing the VVR packages

To remove the VVR packages

1 Insert the software disc, mount it, and enter the following commands:

```
# cd /disc_path/pkgs

# ./installer
```

2 Select Uninstall from the menu.

3 Select VVR.

The program prompts you to confirm whether you want to remove the packages that are being used by other Veritas products.

4 Answer the set of questions depending on your requirements. Note that if you uninstall the `VRTSvxvm` package you will not be able to use the Veritas Volume Manager functionality.

The program asks you to confirm that you want to remove VVR and then removes all the packages except the infrastructure packages. If open volumes exist, the program prompts you to stop the open volumes and unmount the file systems.

The output is similar to the following:

```
uninstallvvr is now ready to uninstall VVR packages.
All VVR processes that are currently running will be stopped.
Are you sure you want to uninstall VVR packages? [y,n,q] (y)
```

5 Press Return to continue.

6 Confirm the rpms have been removed.

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

Removing license files (Optional)

Optionally, you can remove the license files.

To remove the VERITAS license files

1 To see what license key files you have installed on a system, enter:

```
# /sbin/vxlicrep
```

The output lists the license keys and information about their respective products.

2 Go to the directory containing the license key files and list them:

```
# cd /etc/vx/licenses/lic
# ls -a
```

3 Using the output from step 1, identify and delete unwanted key files listed in step 2. Unwanted keys may be deleted by removing the license key file.

Removing the Veritas Enterprise Administrator client

You should also remove the client software from any machines you used to access the Veritas software.

To remove the VEA client for Linux

1 Stop the VEA service.

```
# /opt/VRTSob/bin/vxsvc -k
```

2 Remove the installed VEA client and server packages.

```
# rpm -ev VRTSobgui VRTSob
```

To remove the VEA client from a Windows system

1 Log in as the database administrator.

2 Select **Start > Settings > Control Panel**.

3 Double-click **Add/Remove Programs** to display a list of installed products.

4 Select **Veritas Enterprise Administrator** from the list, and click the **Remove** button.

5 Click **Yes** when a dialog box appears asking you to confirm the removal.

Appendix A

Installation scripts

This appendix includes the following topics:

About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.0MP3 provides several installation scripts.

To install the Veritas Storage Foundation products 5.0MP3 on a system that already has Veritas Storage Foundation 5.0, 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 version prior to 5.0, the recommended installation method is to use the common product installer. To use the common product installer, run the `installer` command.

See "About the common product installer" on page 47.

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 Cluster File System (SFCFS)	`installsfcfs`
Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC)	`installsfcfsrac`
Symantec Product Authentication Service (AT)	`installat`
Veritas Volume Manager.	`installvm`
Veritas File System.	`installfs`

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

Installation script options

Table A-1 shows command line options for the product installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts.

See "About installation scripts" on page 203.

Table A-1 Available command line options

Command Line Option	Function
system1 system2...	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
–configure	Configures the product after installation.
–enckeyfile *encryption_key_file*	Specifies the location of a file containing the key to decrypt encrypted passwords stored in response files. See the `-responsefile` and the `-encrypt` options.

Table A-1 Available command line options *(continued)*

Command Line Option	Function
–encrypt *password*	Encrypts *password* using the encryption key provided with the `-enckeyfile` option so that the encrypted password can be stored in response files.
-hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
–installpkgs	Displays all product packages in correct installation order. Output can be used to create scripts for command line installs, or for installations over a network. See the `requiredpkgs` option.
–installonly	Installs packages, but does not configure the product.
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-i ssh_key_file` to every SSH invocation.
–license	Registers or updates product licenses on the specified systems.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
–noextrapkgs	Additional packages can be installed so that you can upgrade to another Symantec product simply by installing a new license. The `noextrapkgs` option bypasses installation of extra product packages to simplify future maintenance updates.
-nohapkgs	Limits the list of Storage Foundation packages to exclude the Veritas Cluster Server packages. This option only applies to the `installsf` script when one of the following options is specified: ■ -installpkgs ■ -requiredpkgs ■ -jumpstart.

Table A-1 Available command line options *(continued)*

Command Line Option	Function
–nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
–nooptionalpkgs	Bypasses installation of optional product packages such as manual pages.
–nostart	Bypasses startup of the product following installation and configuration.
–patchpath *patch_path*	Designates the path of a directory that contains all patches to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
–pkgpath *package_path*	Designates the path of a directory that contains all packages to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
–precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
–requiredpkgs	Displays all required product packages in correct installation order. Optional packages are not listed. Output can be used to create scripts for command line installs, or for installations over a network. See `installpkgs` option.
–responsefile *response_file* [–enckeyfile *encryption_key_file*]	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file. The `-enckeyfile` option and *encryption_key_file* name are required with the `-responsefile` option when the response file contains encrypted passwords.

Table A-1 Available command line options *(continued)*

Command Line Option	Function
–rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-timeout *timeout_value*	Specifies the timeout (in seconds) that the installer uses for each command it issues during the installation. The default timeout is set to 600 secs. Use the -timeout option to override the default value.
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
-verbose	Displays details during installation of product RPMs. By default, the installation displays only a progress bar .

Appendix B

Storage Foundation and High Availability components

This appendix includes the following topics:

Veritas Storage Foundation installation RPMs

Table B-1 shows the RPM name and contents for each English language RPM for Veritas Storage Foundation, Veritas Storage Foundation High Availability, Veritas Storage Foundation Cluster File System, and Veritas Storage Foundation for databases.

Table B-1 Storage Foundation RPMs

RPM	Contents	Required/Optional
Veritas Volume Manager		
VRTSalloc	Veritas Volume Manager Veritas Intelligent Storage Provisioning Provides the volume tagging features, which is required for dynamic storage tiering (DST).	Required

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSdcli	Veritas Distributed Command Line Interface	Required
VRTSddlpr	Veritas Device Discovery Layer Services Provider Provides the necessary management backend required to administer VxVM DMP features and objects like enclosures, controllers, and paths from the GUI.	Required
VRTSlvmconv	Veritas Linux LVM to VxVM Converter	Optional
VRTSvdid	Veritas Device Identification API	Required
VRTSvmman	Veritas Volume Manager Manual Pages	Optional
VRTSvmpro	Veritas Volume Manager Management Services Provider Provides the necessary management backend required to administer VxVM from the GUI.	Required
VRTSvxvm-common	Veritas Volume Manager common files	Required
VRTSvxvm-platform	Veritas Volume Manager platform-specific files	Required
Veritas File System		
VRTSfsman	Veritas File System Manual Pages	Optional
VRTSfsmnd	Veritas File System Software Developer Kit Manual Pages	Optional

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSfspro	Veritas File System Management Services Provider Provides the necessary management that is required to administer VxFS and other platform filesystems to manage from the GUI. Also, provides dynamic storage tiering (DST) capability that allows users to do policy based control for data placement.	Required
VRTSfssdk	Veritas File System Software Developer Kit For VxFS APIs, the package contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.	Required
VRTSvxfs-common	Veritas File System common files Required for VxFS file system support.	Required
VRTSvxfs-platform	Veritas File System platform-specific files Required for VxFS file system support.	Required
Storage Foundation Cluster File System		
VRTScavf	Veritas Cluster Server Agents for Storage Foundation Cluster File System	Required
VRTSglm	Veritas Group Lock Manager for Storage Foundation Cluster File System	Required

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSgms	Veritas Group Messaging Services for Storage Foundation Cluster File System	Required
Databases		
VRTSd2gui-common	Veritas Storage Foundation for DB2 Graphical User Interface	Required (for Storage Foundation for DB2)
VRTSdb2ed-common	Veritas Storage Foundation for DB2	Required (for Storage Foundation for DB2)
VRTSdbcom-common	Veritas Storage Foundation Common Utilities for Databases	Required (for all database products)
VRTSdbed-common	Veritas Storage Foundation for Oracle	Required (for Storage Foundation for Oracle)
VRTSodm-common	ODM Driver for VxFS Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager allows Oracle9i and 10g to improve performance and manage system bandwidth.	Required (for Storage Foundation for Oracle)
VRTSodm-platform	ODM Driver for VxFS Platform-specific Veritas Extension for Oracle Disk Manager.	Required (for Storage Foundation for Oracle)
VRTSorgui-common	Veritas Storage Foundation for Oracle Graphical User Interface	Required (for Storage Foundation for Oracle)
VRTSvxmsa	Veritas Mapping Service, Application Libraries	Required (for DB2 and Oracle products)
Veritas Enterprise Administrator		
VRTSaa	Veritas Enterprise Administrator Action Agent	Required

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSccg	Veritas Enterprise Administrator Central Control Grid	Required
VRTSob	Veritas Enterprise Administrator	Required
VRTSobc33	Veritas Enterprise Administrator Core	Required
VRTSobgui	Veritas Enterprise Administrator	Optional
Infrastructure		
VRTSatClient	Symantec Product Authentication Service client Installs the Symantec Product Authentication Service, which provides authentication services to other Symantec products. This package contains a server and client component. The client allows Symantec products to communicate with the brokers.	Required
VRTSatServer	Symantec Product Authentication Service Server Installs the Symantec Product Authentication Service, which provides authentication services to other Symantec products. This package contains a server and client component. The server provides services for a root broker, authentication broker, or both.	Required
VRTSicsco	Symantec Infrastructure Core Services Common	Required
High Availability	Note: some of these RPMs are also required for Storage Foundation Cluster File System.	

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSacclib	Veritas Application Competency Center Library VRTSacclib is a set of Perl modules that many cluster server agents use.	Required Depends on VRTSvcs.
VRTScmccc	Veritas Cluster Server Management Console Cluster Connector	Optional
VRTScmcs	Veritas Cluster Management Console for single cluster environments	Optional
VRTScscm	Veritas Cluster Server Cluster Manager	Required Depends on VRTSvcs and VRTSjre15.
VRTScscw	Veritas Cluster Server configuration wizards	Required Depends on VRTSvcsag and VRTSjre15.
VRTScsocw	Veritas agent for Oracle and SF Oracle RAC configuration wizards.	Optional for VCS. Required to use VCS with the high availability agent for Oracle.
VRTScssim	Veritas Cluster Server Simulator	Optional
VRTScutil	Veritas Cluster Server Utilities	Required Depends on VRTSvcs.
VRTSgab	Veritas Cluster Server group membership and atomic broadcast services	Required Depends on VRTSllt.
VRTSjre	Veritas Java Runtime Environment Redistribution	Required

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSjre15	Veritas Java Runtime Environment Redistribution This package installs the Java Runtime Environment for all Symantec products that require Java.	Required
VRTSllt	Veritas Cluster Server low-latency transport	Required
VRTSvcs	Veritas Cluster Server	Required Depends on VRTSut, VRTSperl, VRTSvxfen, VRTSgab, and VRTSllt.
VRTSvcsag	Veritas Cluster Server Bundled Agents	Required Depends on VRTSvcs. Depends on VRTSvcsdr.
VRTSvcsdb	Veritas High Availability Agent for DB2	Optional for VCS. Required to use VCS with the high availability agent for DB2. Depends on VRTSvcs.
VRTSvcsdr	Veritas Cluster Server disk reservation	Required
VRTSvcsmg	Veritas Cluster Server English message catalogs	Required Depends on VRTSvcs. Depends on VRTSvcsdr.
VRTSvcsmn	Manual Pages for Veritas Cluster Server	Optional
VRTSvcsor	Veritas High Availability Agent for Oracle	Optional for VCS. Required to use VCS with the high availability agent for Oracle. Depends on VRTSvcs.

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSvxfen	Veritas I/O Fencing	Required Depends on VRTSgab.
VRTSweb	Symantec Web Server	Required
Veritas Volume Replicator		
VRTSvcsvr	Veritas Cluster Server Agents for VVR	Required
VRTSvrpro	Veritas Volume Replicator Client Extension and Proivder for Veritas Enterprise Administrator	Required
VRTSvrw	Veritas Volume Replicator Web Console	Required
Other RPMs		
SYMClma	Symantec License Inventory Agent	Required
VRTSdbms3	Veritas Shared DBMS	Required
VRTSdsa	Veritas Datacenter Storage Agent	Required
VRTSmapro-common	Veritas Storage Foundation GUI for Mapping	Required
VRTSmh	Veritas Centralized Management for Storage Foundation Managed Host	Required

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSpbx	Symantec Private Branch Exchange This package installs the Symantec Private Branch Exchange, which allows other Symantec products to share a common well-known port for publishing services and communicating.	Required If VRTSpbx is removed, Symantec products that use it are unable to communicate, which can cause the products to stop working. If VRTSat is configured to work with VRTSpbx, and VRTSpbx is removed, VRTSat continues to work. However, the Symantec Product Authentication Service remote administration functionality are not available. Removing VRTSat can affect Symantec products that use the Symantec Product Authentication Service remote administration feature, such as VEA.
VRTSperl	Perl 5.8.8 for Veritas	Required
VRTSsmf	Symantec Service Management Framework	Required
VRTSspt	Veritas Software Support Tools	Required
VRTSvlic	Veritas License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Required If VRTSvlic is removed, the Storage Foundation products may not be able to access their license information, The products may fail to start or fail to work properly.
VRTScweb	Symantec Web Server	Required
VRTSdcp	Veritas Disk Correlator Provider	Required
VRTSdsm	Veritas Datacenter Storage Manager	Required
VRTSgcsha	Veritas GCS High Availability Agents	Required

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
VRTSgcspr	Veritas SAN Global Configuration Server Object Bus provider	Required
windows/vrtsobgui.msi	Veritas Enterprise Administrator for Windows	Optional

Obsolete RPMs in Storage Foundation

The following RPMs were included in previous releases of Storage Foundation but are now obsolete:

```
SYMClma
VRTSsmf
VRTScmcm
VRTSjre
VRTSvsvc
VRTSfsdoc
VRTSvmdoc
VRTSvrdoc
VRTSvcsdc
VRTSdbdoc
VRTScsdoc
VRTScfsdc
```

Appendix C

Troubleshooting information

This appendix includes the following topics:

Troubleshooting information

The `VRTSspt` package provides a group of tools for troubleshooting a system and collecting information on its configuration. The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. The tools are not required for operation of any Veritas product, and they may adversely impact system performance if not used correctly. Veritas provides these tools to analyze systems if you suspect that there are performance problems. The tools should be used only under the direction of a Veritas Technical Support Engineer.

Storage Foundation Cluster File System installation issues

If you encounter any issues installing SFCFS, refer to the following paragraphs for typical problems and their solutions.

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Checking ssh communication with system01 ......... permission denied
installer requires that ssh commands used between systems execute without
prompting for passwords or confirmations.  Please run installer again with
the ssh configured for password free logins, or configure rsh and use the
-rsh option.
```

Suggested solution: You need to set up the systems to allow remote access using `ssh` or `rsh`.

See "Configuring secure shell (ssh) or remote shell before installing products" on page 22.

Note: Remove remote shell permissions after completing the SFCFS installation and configuration.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
Checking communication with system01 ................ FAILED
  System not accessible : system01
```

Suggested solution: Verify that you entered the system name correctly; use the `ping`(1M) command to verify the accessibility of the host.

Storage Foundation Cluster File System problems

If there is a device failure or controller failure to a device, the file system may become disabled cluster-wide. To address the problem, unmount file system on all the nodes, then run a full `fsck`. When the file system check completes, mount all nodes again.

Unmount failures

The `umount` command can fail if a reference is being held by an NFS server. Unshare the mount point and try the unmount again.

Mount failures

Mounting a file system can fail for the following reasons:

- The file system is not using disk layout Version 6 or 7.
- The mount options do not match the options of already mounted nodes.
- A cluster file system is mounted by default with the `qio` option enabled if the node has a Quick I/O for Databases license installed, even if the `qio` mount option was not explicitly specified. If the Quick I/O license is not installed, a cluster file system is mounted without the `qio` option enabled. So if some nodes in the cluster have a Quick I/O license installed and others do not, a cluster mount can succeed on some nodes and fail on others due to different mount options. To avoid this situation, ensure that Quick I/O licensing is uniformly applied, or be careful to mount the cluster file system with the `qio/noqio` option appropriately specified on each node of the cluster.
  See the `mount`(1M) manual page.
- A shared CVM volume was not specified.
- The device is still mounted as a local file system somewhere on the cluster. Unmount the device.
- The `fsck` or `mkfs` command is being run on the same volume from another node, or the volume is mounted in non-cluster mode from another node.
- The `vxfsckd` daemon is not running. This typically happens only if the `CFSfsckd` agent was not started correctly.
- If `mount` fails with an error message:

  ```
  vxfs mount: cannot open mnttab
  ```

  `/etc/mnttab` is missing or you do not have `root` privileges.
- If `mount` fails with an error message:

  ```
  vxfs mount: device already mounted, ...
  ```

  The device is in use by `mount`, `mkfs` or `fsck` on the same node. This error cannot be generated from another node in the cluster.
- If this error message displays:

  ```
  mount: slow
  ```

  The node may be in the process of joining the cluster.
- If you try to mount a file system that is already mounted without `-o cluster` option (that is, not in shared mode) on another cluster node,

```
# mount -t vxfs /dev/vx/dsk/share/vol01 /vol01
```

The following error message displays:

```
vxfs mount: /dev/vx/dsk/share/vol01 is already mounted,
  /vol01 is busy, allowable number of mount points exceeded,
  or cluster reservation failed for the volume
```

Command failures

This section describes command failures.

- Manual pages not accessible with the `man` command. Set the MANPATH environment variable appropriately.
  See "Accessing manual pages and documentation directories" on page 20.
- The `mount`, `fsck`, and `mkfs` utilities reserve a shared volume. They fail on volumes that are in use. Be careful when accessing shared volumes with other utilities such as `dd`, it is possible for these commands to destroy data on the disk.
- Running some commands, such as `vxupgrade -n 7/vol02`, can generate the following error message:

  ```
  vxfs vxupgrade: ERROR: not primary in a cluster file system
  ```

  This means that you can run this command only on the primary, that is, the system that mounted this file system first.

Performance issues

Quick I/O File system performance is adversely affected if a cluster file system is mounted with the `qio` option enabled, but the file system is not used for Quick I/O files. Because `qio` is enabled by default, if you do not intend to use a shared file system for Quick I/O, explicitly specify the `noqio` option when mounting.

High availability issues

This section describes high availability issues.

Network partition/jeopardy

Network partition (or split brain) is a condition where a network failure can be misinterpreted as a failure of one or more nodes in a cluster. If one system in the cluster incorrectly assumes that another system failed, it may restart applications

already running on the other system, thereby corrupting data. CFS tries to prevent this by having redundant heartbeat links.

At least one link must be active to maintain the integrity of the cluster. If all the links go down, after the last network link is broken, the node can no longer communicate with other nodes in the cluster. Thus the cluster is in one of two possible states. Either the last network link is broken (called a network partition condition), or the last network link is okay, but the node crashed, in which case it is not a network partition problem. It is not possible to identify whether it is the first or second state, so a kernel message is issued to indicate that a network partition may exist and there is a possibility of data corruption.

Jeopardy is a condition where a node in the cluster has a problem connecting to other nodes. In this situation, the link or disk heartbeat may be down, so a jeopardy warning may be displayed. Specifically, this message appears when a node has only one remaining link to the cluster and that link is a network link. This is considered a critical event because the node may lose its only remaining connection to the network.

Warning: Do not remove the communication links while shared storage is still connected.

Low memory

Under heavy loads, software that manages heartbeat communication links may not be able to allocate kernel memory. If this occurs, a node halts to avoid any chance of network partitioning. Reduce the load on the node if this happens frequently.

A similar situation may occur if the values in the `/etc/lltab` files on all cluster nodes are not correct or identical.

Index