Veritas Storage Foundation™ Installation Guide

Linux for IBM Power

5.0 Release Update 3

symantec™

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 RU3a

Document version: 5.0RU3.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's maintenance offerings include the following:

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

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

www.symantec.com/techsupp/

Contacting Technical Support

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

www.symantec.com/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:

https://licensing.symantec.com

Customer service

Customer service information is available at the following URL:

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

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

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

Maintenance agreement resources

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

Asia-Pacific and Japan	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

Veritas Storage Foundation and High Availability Solutions 5.0 Release Update 3 adds support for Linux for IBM Power on the following platforms:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 2 (2.6.18-92.el5 kernel) or later on ppc64
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21) on ppc64

The following 5.0 RU3 documentation is updated and available on the product disc for this 5.0 RU3 release:

- *Veritas Storage Foundation and High Availability Solutions Getting Started Guide*
- *Veritas Storage Foundation Installation Guide*
- *Veritas Storage Foundation Release Notes*
- *Veritas Cluster Server Installation Guide*
- *Veritas Cluster Server Release Notes*

The 5.0 RU3 documentation is also available online from the Symantec Veritas Storage Foundation 5.0 RU3 website:

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

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 Optionally licensed features: Veritas Volume Replicator
Storage Foundation Standard HA	Veritas File System Veritas Volume Manager Veritas Cluster Server Optionally licensed features: Veritas Volume Replicator
Storage Foundation Enterprise	Veritas File System Veritas Volume Manager Veritas Storage Checkpoint option Optionally licensed features: Veritas Volume Replicator
Storage Foundation Enterprise HA	Veritas File System Veritas Volume Manager Veritas Storage Checkpoint option Veritas Cluster Server 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 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
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 DB2 Standard	Veritas File System Veritas Volume Manager Veritas Quick I/O option Optionally licensed features: Veritas Volume Replicator
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

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

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

Follow the preinstallation instructions if you are installing one of the Veritas Storage Foundation Cluster File System for Oracle RAC products by Symantec.

Several component products are bundled with each of these SFCFS for Oracle RAC products.

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

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.

After installation, Veritas Storage Foundation commands are stored in `/opt/VRTS/bin` and HA commands are stored in `/opt/VRTSvcs/bin`. Storage Foundation HA manual pages are stored in `/opt/VRTS/man`.

Specify `/opt/VRTS/bin` and `/opt/VRTSvcs/bin` in your PATH after the path to the standard Linux commands. If you are not installing an HA product, you can omit `/opt/VRTSvcs/bin`. To invoke the VxFS-specific `df`, `fsdb`, `ncheck`, or `umount` commands, type the full path name: `/opt/VRTS/bin/`*`command`*.

To set your MANPATH environment variable to include `/opt/VRTS/man` do the following:

- If you are using a shell such as sh or bash, enter the following:

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

- If you are using a shell such as csh or tcsh, enter 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, enter the following:

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

- If you are using a shell such as csh or tcsh, enter 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.

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.

Important pre-installation information for SFCFS for Oracle RAC

Before you install SFCFS for Oracle RAC:

- Review the current hardware compatibility list to verify that your systems are compatible with the requirements:
  http://entsupport.symantec.com/docs/283161
- Review the Technical Support TechNote for the latest information on updates, patches, and software issues regarding this release:
  http://www.symantec.com/techsupp/
- Review the latest information on the product:
  http://entsupport.symantec.com/docs/281993
- Configure SSH or RSH and set up the systems so that commands between systems execute without prompting for passwords or confirmations. By default, SSH is used during installation and configuration. If you want to use RSH, use -rsh option.
- Ensure that Security Enhanced Linux (SELinux) and firewall options are disabled during Red Hat Enterprise Linux installation.

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 2-1 shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

Figure 2-1 Four Node SFCFS Cluster Built on Fibre Channel Fabric

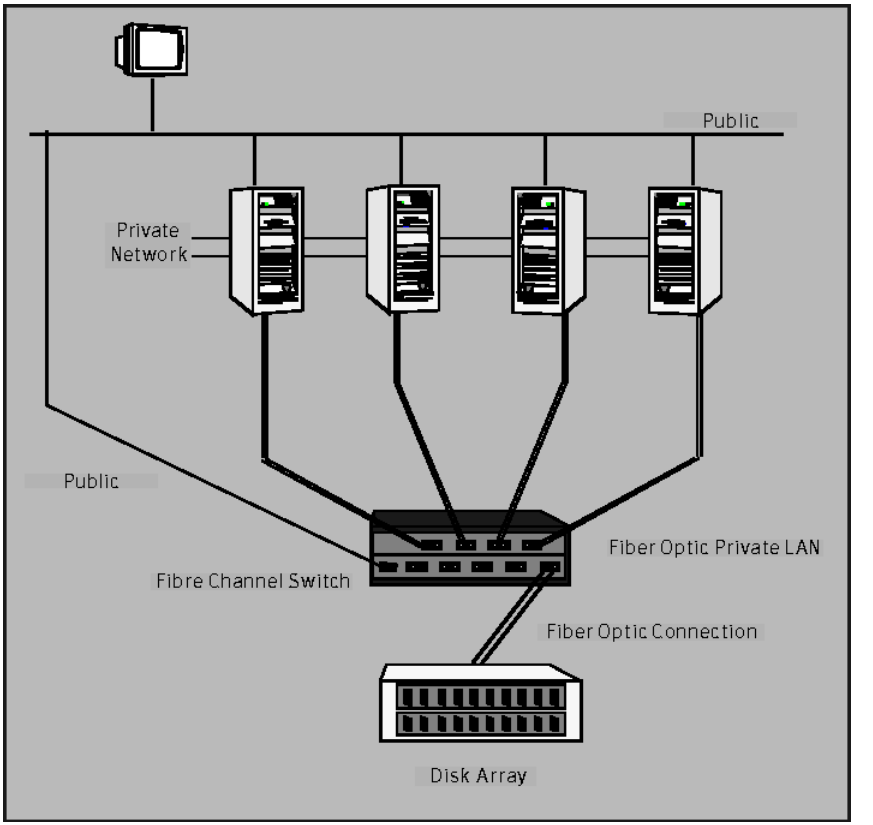

Shared storage

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 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:

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.

VEA installation planning

The Veritas Enterprise Administrator (VEA) GUI consists of several RPMs. Follow these planning guidelines to install VEA for use with VVR.

- The VEA server packages must be installed on the hosts on which VVR is installed, not the client. If you install using the product installer, these RPMs are installed when you install Storage Foundation products.

  The VEA server packages include the following:

  - The Veritas Volume Replicator Management Services Provider RPM, `VRTSvrpro`, must be installed on all hosts in the Replicated Data Set (RDS).
  - For `VRTSvrpro` to function, the Veritas Volume Manager Management Services Provider RPM, `VRTSvmpro`, must be installed on each system.

- The VEA client can be installed on a host on which VVR is installed, or a separate host that is used to administer the VVR hosts. To use the VEA client on a system, the `VRTSobgui` RPM must be installed on that system.

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 supported on Linux for IBM Power platforms.

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

OS, Platform, Version	RHEL5	SLES10
Oracle10g R2	Full Support	Full Support

Note: For Oracle 10gR2, there is full support in this release for Storage Foundation for Oracle on SLES 10 SP2 and RHEL5 r2 or later.

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	RHEL5	SLES10
DB2 9.5	Full Support	Full Support
DB2 9.7	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 27.

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 157.

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 35.

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
   ```

   Caution: When you select a location to download files, do not select a directory that contains Veritas products from a previous release or maintenance pack. You must download the Veritas 5.0 software and the Veritas 5.0RU3 software into separate directories.

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

Downloading Storage Foundation Manager

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

http://www.go.symantec.com/vom

Chapter 3

System requirements

This chapter includes the following topics:

Hardware and software 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

Supported Linux operating systems

This section lists the supported Linux operating systems for this 5.0 RU3 release of Veritas Storage Foundation products, including Veritas Storage Foundation Cluster File System.

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

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

The Veritas 5.0 RU3 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 2 (2.6.18-92.el5 kernel) or later on ppc64-bit Linux For IBM Power systems
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21) on ppc64-bit Linux For IBM Power systems

Information about the latest supported Red Hat erratas and updates and SuSE service packs is available in the following Late-Breaking News TechNote. Read this TechNote before you install Symantec products.

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

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=ethX` 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 versions 9.5 and 9.7 are supported on the Linux operating systems listed above for this 5.0 RU3 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 10gR2 is supported on the Linux operating system for this 5.0 RU3 release.

For Oracle 10gR2, there is full support for Storage Foundation for Oracle on SLES 10 SP 2 and RHEL 5 Update 2 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 is supported for use with Storage Foundation Cluster File System for Oracle RAC.

Installation requirements for SFCFS for Oracle RAC

Make sure that you meet the required hardware and software requirements before installing SFCFS for Oracle RAC.

Hardware requirements

Table 3-1 lists the hardware requirements for SFCFS for Oracle RAC.

Table 3-1 Hardware requirements

Item	Description
SFCFS for Oracle RAC systems	Supports a maximum of 16 systems with two or more CPUs at 2GHz or higher.
RAM	Each SFCFS for Oracle RAC system requires 2 GB or more of physical memory.
Network links	Two or more private links and one public link. Symantec recommends Gigabit Ethernet using enterprise-class switches for the private links.
DVD drive	One drive that is accessible to all nodes in the cluster.
Fibre channel or SCSI host bus adapters	SFCFS for Oracle RAC requires at least one built-in SCSI adapter per system to access the operating system disks, and at least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.
Disks	Typical SFCFS for Oracle RAC configurations require that shared disks support applications that migrate between systems in the cluster.
Disk space	■ total: 2.6 G ■ /opt: 1.4 G ■ /usr: 200 KB ■ /tmp: 512 MB ■ /var: 32 MB ■ /var/tmp: 700 MB
Swap space	Two times the main memory.

Software requirements

The supported software and the minimum required versions for SFCFS for Oracle RAC are as follows:

Oracle RAC	Oracle RAC 10g Release 2
Operating systems	See "Supported operating systems" on page 32.

Supported operating systems

This section lists the supported Linux operating systems. Within a cluster, all nodes must use the same operating system version and patch level.

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

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

The following operating systems are supported on the AMD Opteron or Intel Xeon EM64T (x86_64) architectures:

Red Hat Enterprise Linux (RHEL)	Red Hat Enterprise Linux 5 with Update 3 (2.6.18-128.el5 kernel) or later
SUSE Linux Enterprise Server (SLES)	SUSE Linux Enterprise Server 10 with SP2 (2.6.16.60-0.21 kernel) or later

No specific patches are required for SFCFS for Oracle RAC. Additionally, see the Oracle documentation for any patches required by Oracle.

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 SFCFS for Oracle RAC.

Consult the vendor documentation for more information on upgrading or reinstalling your system.

Symantec products operate on subsequent kernel and patch releases provided that 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

Oracle RAC requirements

The requirements for Oracle RAC installation are as follows:

Private IP address	Oracle Clusterware needs a private IP address per node for communication between clusters (heartbeat with all the other nodes in the cluster). While starting Oracle Clusterware, it waits till the private IP address is brought up. SFCFS for Oracle RAC does not support PrivNIC agent. Hence after starting the system, you need to manually bring up the private IP address or configure the system to automatically bring up the private IP address. Update the /etc/hosts file located on all the nodes to include an entry for the private IP address.
Public IP address for VIP resource	Oracle Clusterware needs a public IP address per node for configuring the VIP resource. Before starting the Oracle Clusterware installation, one public IP address must be set aside and registered. You may register the public IP address with DNS or place the entries in the /etc/hosts file located on all the nodes.
Private IP address for database	Oracle RAC cache-fusion traffic may be configured to use UDP. Hence additional private IP addresses must be preset before the actual installation and configuration. You may register the private IP addresses with DNS or place the entries in the /etc/hosts file located on all the nodes.
SSH/RSH	Oracle requires password-less SSH and RSH to be set up across cluster nodes for the user that performs Oracle Clusterware installation. For more information on configuring SSH/RSH, see the Oracle installation guide.
Kernel parameters	Verify that the kernel parameter values are set as recommended by Oracle. For details on the list of kernel parameters and their recommended values, see the Oracle RAC documentation.

VxVM licenses

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

Table 3-2 describes the levels of licensing in Veritas Volume Manager and supported features.

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

Table 3-3 Required disk space

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

Disk space requirements for Veritas Volume Replicator

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

Table 3-4 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 3-5 shows the approximate disk space used by SFCFS for the required and optional packages.

Table 3-5 Approximate disk space use for SFCFS

English	/root	/opt	/usr	/var
Required Packages	58 MB	634 MB	50 MB	0.02 MB
All Packages	58 MB	689 MB	50 MB	0.68 MB

Disk space requirements for Storage Foundation for Oracle

Table 3-6 shows the approximate disk space used by SFORA for the required and optional packages.

Table 3-6 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:

About the common product installer

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

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

See "About installation scripts" on page 157.

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. You can also enter `q` to quit the installer or `?` to display help information.

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 158.

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.

   Note: For platforms that do not have access to the cdrom, copy the software packages to virtual machines using ftp.

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.

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.

Note: CPI does not automatically install VCS Sybase Agent in this RU3 release. To use the VCS Sybase Agent, install the agent RPM manually.

The Veritas 5.0 RU3 release operates on the following operating system and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 2 (2.6.18-92.el5) or later on ppc64
- SuSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21-ppc64 kernels) on ppc64

To install Storage Foundation

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

   See "Mounting a software disc" on page 40.
3. Move to the top-level directory on the disc.
4. From this directory, type the following command to install on the local system only. Also use this command to install on remote systems using the secure shell (ssh) utilities:

   ```
   # ./installer
   ```

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

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

   The sample installation assumes that ssh is used.
5. Enter `I` to install and press Return.
6. When the list of available products is displayed, select Veritas Storage Foundation, enter the corresponding number, and press Return.

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

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

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

Enter **n** if you do not want to specify the name of the default disk group at this time. You can set the name of the default disk group after installation.

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

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

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

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

14 You are prompted to confirm the default disk group.

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

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

15 Verify the fully qualified hostname of the systems.

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

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

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

http:// www.go.symantec.com/vom

Press **Enter** to continue.

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

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

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

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

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

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

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.

2 Load and mount the software disc.

See "Mounting a software disc" on page 40.

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.

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)
```

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

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 secure cluster.

```
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.

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 You are prompted to confirm the default disk group.

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

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

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

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

27 This product can be configured as a Storage Foundation Manager managed host or a stand-alone host. Enabling Veritas Storage Foundation Manager management simplifies and improves management of the complex data center resources, reducing planned and unplanned down time. To enable centralized management using Storage Foundation Manager, please download Veritas Storage Foundation Manager from:

http://go.symantec.com/vom

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.

28 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
```

29 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-****
```

30 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 76.

31 Reboot the system (or systems).

32 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
```

Installing Storage Foundation Cluster File System using the common product installer

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

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

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

To install the Storage Foundation Cluster File System

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

2. Load and mount the software disc.

   See "Mounting a software disc" on page 40.

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.

8 Enter the product license information.

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

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

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

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

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

```
1)Required Veritas Storage Foundation Cluster File System depots -
1566 MB required
2)All Veritas Storage Foundation Cluster File System depots -
1623 MB required
3)Storage Foundation for Oracle RAC depots - 2029 MB required

Select the depots 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.

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.

2 Load and mount the software disc.

See "Mounting a software disc" on page 40.

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.

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.

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 You are prompted to confirm the default disk group.

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

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

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.

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

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

http://go.symantec.com/vom

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.

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.

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.

After VxVM is upgraded, install VVR-specific components and configure VVR.

See "Installing and configuring Veritas Volume Replicator using the common product installer" on page 53.

To use the new features of VVR 5.0RU3, upgrade the version of each disk group.

Installing Veritas Enterprise Administrator

Veritas Enterprise Administrator (VEA) is required to access the graphical user interface (GUI) for Veritas Storage Foundation. You can use the GUI to administer disks, volumes, file systems, and database functionality on local or remote machines. This section describes the installation of VEA components.

The VEA server RPM, `VRTSob`, is installed when you install Veritas Storage Foundation products using the installation script. The VEA server package must be installed on all nodes that are to be administered.

The VEA client RPM contains the Graphical User Interface (GUI) program to administer Veritas Storage Foundation products. The VEA client may be installed on one or more of the nodes to be administered. The VEA client may also be installed on a separate system that can be used to administer Veritas Storage Foundation.

Installing the Veritas Enterprise Administrator client

Veritas Enterprise Administrator (VEA) is required to access the graphical user interface (GUI) for Veritas Storage Foundation. You can use the GUI to administer disks, volumes, file systems, and database functionality on local or remote machines.

The VEA client requires one of the following RPMs:

- Veritas Enterprise Administrator client RPM (`VRTSobgui`)
  This is the client package for UNIX.
- Veritas Enterprise Administrator for Windows (`windows\VRTSobgui.msi`)
  This is the client package for Windows.

Minimum system requirements for VEA clients

Table 4-1 shows the system minimum requirements for the GUI.

Table 4-1 VEA system minimum requirements

Operating System	System minimum requirements
Windows XP, NT, Me, 2000, 98, 2k3, Vista, 2k8	300MHz Pentium with at least 256MB of memory

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, Windows 95, Vista, and Windows 2k3 and 2k8 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.
- 100MHz Pentium with 256MB memory or higher specification.
- 100MB 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 for Databases" on page 73.

If the Storage Foundation products were installed with another method, they may also need to be configured. Review the configuration sections that are appropriate for the Storage Foundation products that were installed. 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. If you specify a name, it is used for Veritas Volume Manager commands when a disk group is not specified.

Enter **n** if you do not want to specify the name of the default disk group at this time. You can set the name of the default disk group after installation.

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

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

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 You are prompted to confirm the default disk group.

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

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

9 If a valid license for VVR is installed, the installer prompts you for the VVR configuration. If a license for VVR is not present, skip to step 14.

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
```

10 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
```

11 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
```

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 Repeat steps 9 to 12 for all other systems.

14 Verify the fully qualified hostname of the systems.

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

15 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
```

16 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 High Availability and clusters using the common product installer. Use the same procedures to configure Storage Foundation for Oracle High Availability.

Required information for configuring Storage Foundation and High Availability Solutions

To configure Storage Foundation High Availability or Storage Foundation for Oracle High Availability, 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 secure cluster.

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 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
```

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

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

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

```
Checking system licensing

SF license registered on host1

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

5 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
```

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

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

7 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.

8 A summary of the information you entered is given. When prompted, confirm that the information is correct.

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

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

9 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 secure cluster.

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

10 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
```

11 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
```

12 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
```

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

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

14 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
```

15 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
```

16 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
```

17 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)
```

18 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)
```

19 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
```

20 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.

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
```

21 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
```

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

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

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

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

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

http://go.symantec.com/vom

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.

25 The Veritas Storage Foundation or Veritas Storage Foundation for Oracle 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
```

26 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-****
```

About adding and removing nodes in a cluster

After you install Storage Foundation High Availability and create a cluster, you can add and remove nodes from the cluster. You can create a cluster of up to 32 nodes.

For information about adding and removing nodes, see the *Veritas Cluster Server Installation Guide*.

Configuring Storage Foundation for Databases

This section describes the procedure to configure Storage Foundation for Databases, using the common product installer.

You can use this procedure to configure Veritas Storage Foundation for Oracle (SFORA) or Veritas Storage Foundation for DB2 (SFDB2).

The example in this section shows a simple configuration on a single host. If you are installing Storage Foundation High Availability product or installing on multiple hosts, there are additional configuration prompts.

See "Configuring Storage Foundation and High Availability Solutions" on page 65.

Some databases may require additional configuration steps. See the following sections for details.

See "Creating and configuring the repository database for DB2 and Oracle" on page 76.

To configure Storage Foundation for Oracle

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 the number correspnding to the product you want to configure, and press Return.

You can use this procedure to configure Veritas Storage Foundation for Oracle (SFORA) or Veritas Storage Foundation for DB2 (SFDB2).

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

3 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
configure SF: host1
```

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.

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 If you are configuring Veritas Storage Foundation for Oracle or Veritas Storage Foundation for DB2, you are now prompted to configure permissions to allow database administrators (DBAs) access to the tools to support the Veritas Storage Foundation product. The default settings only allow access to the root user.

Respond **y** to change permission for a DBA or a group of DBAs to access the support tools. When prompted, enter the login account or group name.

For example, enter the following for a Veritas Storage Foundation for Oracle configuration:

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

7 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
```

8 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.

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

12 The Veritas Storage Foundation for databases software is verified and configured.

You are prompted to start the Veritas Storage Foundation product processes.

For example, when you configure Veritas Storage Foundation for Oracle, the following prompt displays:

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

13 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-****
```

14 After the configuration completes, restart the Storage Agent

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

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

16 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 76.

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.
  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: 10.209.87.240

Enter the netmask for public NIC interface: [255.255.252.0]

Following information will be used for SFORA HA configuration:

        Public IP address:              10.209.87.240
        Subnet mask:                    255.255.252.0
        Public interface:               host1 -> lan0,host2 -> lan0
        Mount point:                    /sfua_rep
        Volume Name for mount point:    repvol
        Diskgroup for mount point:      repdg

Is this correct (y/n/q) [y]?

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. If you have installed and configured VxVM using the product installer, you do not need to complete the procedures in this section.

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 restart
```

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
```

Configuring Storage Foundation Cluster File System

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

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

To configure Storage Foundation Cluster File System

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

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

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

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

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

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

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

```
Checking system licensing

SFCFS license registered on system01

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

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

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

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

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

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

All systems are configured to create one cluster.

Enter the unique cluster name and Cluster ID number.

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

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

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

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

10 A summary of the information you entered is given. When prompted, confirm that the information is correct.

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

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

11 When prompted to configure the product to use Veritas Security Services, enter `n`, or enter `y` to configure.

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

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

12 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)
```

13 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)
```

14 Enter `y` if the information is correct.

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

15
```
Do you want to configure SMTP notification? [y,n,q] (y)
```

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

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

Enter the domain-based hostname of the SMTP server (for example: smtp.yourcompany.com):

[b,?] **smtp.mycompany.com**

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

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

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

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

```
Active NIC devices discovered on host1: eth0 Enter the NIC
for the SF Notifier to use on host1: [b,?] (eth0) eth0
Is bge0 to be the public NIC used by all systems [y,n,q,b,?] (y) y
```

Enter the SNMP trap daemon port: [b,?] (162)

```
162
```

Enter the SNMP console system name: [b,?] **host1**

```
Enter the minimum severity of events for which SNMP traps should
be sent to host1 [I=Information, W=Warning, E=Error, S=SevereError]:
[b,?] E

Would you like to add another SNMP console? [y,n,q,b] (n) n
```

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

See the *Veritas Volume Manager Administrator's Guide*

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

18 You are 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
```

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

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

Specify a default disk group name for all systems. [?] diskgroup001
```

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

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

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

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

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

22 After the configuration completes, restart the Storage Agent.

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

23 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.

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

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

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

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

Configuring Veritas Volume Replicator

This section describes configuring Veritas Volume Replicator using the Veritas product installer. If you configured Veritas Volume Replicator during the installation process, you do not need to perform the procedure in this section.

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.

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 20.

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 Planning and Tuning 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 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
```

18 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.

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

2 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 -h host_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.

`-h` specifies the hostname.

For example, to add a database administrator with the name "oracle" as a user with `root` privileges on the host "host1", enter the following:

```
# /opt/VRTS/bin/vxdbedusr -a user -A -f -n oracle -h host1
```

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

```
# /opt/VRTS/bin/vxdbedusr -a user -n user_name -h host_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 -h host1
```

4 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 -h hostname
```

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.

`-h` specifies the hostname.

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

```
# /opt/VRTS/bin/vxdbedusr -a group -A -f -n dba -h host1
```

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

```
# /opt/VRTS/bin/vxdbedusr -r {user | group} \
-n {user_name | group_name} -h host_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 -h host1
```

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
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
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

Installing and configuring SFCFS for Oracle RAC

This chapter includes the following topics:

About installing and configuring SFCFS for Oracle RAC

You can install SFCFS for Oracle RAC on clusters of up to 16 nodes.

The following packages are installed on each cluster node:

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

You can configure the following components for SFCFS for Oracle RAC:

- Veritas Cluster Server (VCS)

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

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

Installing SFCFS for Oracle RAC

The following procedure describes the installation of an SFCFS for Oracle RAC cluster with two nodes: 'galaxy' and 'nebula'.

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

The product installer 'installsfcfsrac' is the recommended program to license and install SFCFS for Oracle RAC.

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

To install SFCFS for Oracle RAC

1. Insert the product disc with the SFCFS for Oracle RAC software into a drive connected to the system.
2. Navigate to the directory containing the installation program.

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

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

   Using `installer` program:

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

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

   Using `installsfcfsrac` program:

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

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

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

5 Enter the product license information.

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

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

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

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

The installer displays the option to install RPMs.

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

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

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

8 Enter `n` to configure SFCFS for Oracle RAC by running the `installsfcfsrac` script with the `-configure` option.

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

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

For instructions:

See "Configuring SFCFS for Oracle RAC" on page 102.

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

Configuring SFCFS for Oracle RAC

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

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

To configure SFCFS for Oracle RAC

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

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

2 Run the following command to configure SFCFS for Oracle RAC:

- If you are using SSH:
  ```
  # ./installsfcfsrac -configure
  ```
- If you are using RSH:
  ```
  # ./installsfcfsrac -rsh -configure
  ```

3 Enter the names of the systems on which you want to configure SFCFS for Oracle RAC. Press Return.

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

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

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

4 Enter additional licenses, if required.

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

If there are any SFCFS for Oracle RAC processes running, these processes are stopped. Enter Return to continue.

5 Press Return to continue.

All systems are configured to create one cluster.

Enter the unique cluster name and Cluster ID number.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

---

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

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

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

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

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

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

Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,?] smtp.mycompany.com
Enter the full email address of the SNMP recipient
(example: user@yourcompany.com): [b,?] user@mycompany.com
Enter the minimum severity of events for which mail should be sent
to user@mycompany.com [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
```

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

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

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

```
Do you want to configure SNMP notification? [y,n,q] (y) y
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the SF Notifier to use on galaxy: [b,?] (eth0) eth0
Is eth0 to be the public NIC used by all systems [y,n,q,b,?] (y) y
Enter the SNMP trap daemon port: [b,?] (162) 162
Enter the SNMP console system name: [b,?] galaxy
Enter the minimum severity of events for which SNMP traps should
be sent to host1 [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
Would you like to add another SNMP console? [y,n,q,b] (n) n
```

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

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

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

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

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

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

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

Enter `n` if you do not want to specify the name of the default disk group at this time. You can set the name of the default disk group after installation by running the `vxdctl defaultdg diskgroup` command on a system.

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

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

```
Will you specify one disk group name for all eligible
systems? [y,n,q,?] (y) y
Specify a default disk group name for all systems. [?] diskgroup001
```

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

18 Review the fully qualified host name information and enter `y` to confirm.

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

19 Enter `y` to accept the fully qualified host name information, otherwise enter `n`.

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

20 Enable or disable Storage Foundation Management Server, as required.

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

Enabling Veritas Storage Foundation Manager management simplifies and improves management of the complex data center resources, reducing planned and unplanned down time. To enable centralized management using Storage Foundation Manager, please download Veritas Storage Foundation Manager from:

http://go.symantec.com/vom

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.

The installer starts the necessary processes for SFCFS for Oracle RAC.

The SFCFS for Oracle RAC configuration process is complete.

You can verify that the configuration process completed successfully by verifying the log file in the following directory: /opt/VRTS/install/logs.

```
Configuration log files, summary file, and response file are saved at:
/opt/VRTS/install/logs/installer-****
```

Verifying installation

After installing and configuring SFCFS for Oracle RAC, verify that the following 7 ports are active and shown in GAB port membership:

```
# gabconfig -a

GAB Port Memberships
===================================
Port a gen   79e807 membership 01
Port b gen   79e82c membership 01
Port d gen   79e82c membership 01
```

```
Port h gen   79e82c membership 01
Port f gen   79e82c membership 01
Port v gen   79e82c membership 01
Port w gen   79e82c membership 01
```

Note: GAB Port 'b' indicates that Symantec I/O fencing is up. In SFCFS for Oracle RAC, Symantec I/O fencing runs in a disabled mode.

Chapter 7

Installing Oracle RAC in an SFCFS for Oracle RAC environment

This chapter includes the following topics:

Preparing to install Oracle RAC

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

Recommendations before installing Oracle RAC software

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

About the location of ORACLE_HOME

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

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

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

About the location of CRS_HOME

CRS_HOME is the location where Oracle Clusterware binaries are installed.

Creating operating system groups and users

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

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

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

Creating the Oracle user and groups

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

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

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

1 Create the 'oinstall' group on each system:

```
# groupadd -g 1000 oinstall

# groupadd -g 1001 dba
```

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

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

Creating CRS_HOME

The following procedure provides instructions on creating `CRS_HOME`.

To create CRS_HOME on each system

1 Log in as root user on a system.

```
# su - root
```

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

```
# vxdg init crsbindg sdc
```

3 Create the volume in the group for `CRS_HOME`:

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

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

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

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

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

5 Create the mount point for `CRS_HOME`:

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

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

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

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

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

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

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

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

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

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

See "Sample main.cf file for configuring a volume and file system under VCS" on page 121.

Creating ORACLE_HOME

The following procedure provides instructions on creating ORACLE_HOME:

To create ORACLE_HOME on each system

1 Log in as root user on a system.

```
# su - root
```

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

```
# vxdg init orabindg sde
```

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

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

3 Create the volume in the group for ORACLE_HOME:

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

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

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

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

5 Create the mount point for ORACLE_HOME:

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

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

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

For each shared ORACLE_HOME, run the following command:

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

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

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

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

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

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

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

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

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

See "Sample main.cf file for configuring a volume and file system under VCS" on page 121.

Configuring private IP addresses

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

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

Obtaining public virtual IP addresses for use by Oracle

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

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

Creating OCR and Vote-disk volumes

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

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

To create the OCR and Vote-disk shared volumes

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

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

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

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

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

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

   where `ocrvol` is the OCR volume.

   where `vdvol` is the Vote-disk volume.

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

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

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

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

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

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

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

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

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

where `oinstall` is the group name.

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

where `ocrvol` is the OCR volume.

where `vdvol` is the Vote-disk volume.

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

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

See "Sample main.cf file for configuring a volume and file system under VCS" on page 121.

Verifying the OCR and Vote-disk shared volumes

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

To verify the OCR and Vote-disk shared volumes

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

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

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

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

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

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

Installing Oracle Clusterware and database software

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

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

Completing the post-installation tasks

Perform the following tasks after installing Oracle RAC:

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

Relinking with ODM

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

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

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

To configure Veritas Extension for Oracle Disk Manager

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

   For Oracle RAC 10g:

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

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

- Back up `libodm10.so` file.

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

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

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

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

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

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

   The alert log location depends on the Oracle version used.

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

Creating Oracle databases

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

Before you create database tablespaces:

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

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

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

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

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

Creating database tablespaces on shared raw VxVM volumes

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

To create database tablespace on shared raw VxVM volumes

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

```
# vxdctl -c mode

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

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

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

```
# vxdisk -o alldgs list

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

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

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

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

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

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

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

For example, enter:

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

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

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

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

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

For example:

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

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

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

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

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

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

See "Sample main.cf file for configuring a volume and file system under VCS" on page 121.

7 Create the database using the Oracle documentation.

Creating database tablespaces on CFS

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

To create database tablespaces on CFS

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

```
# vxdctl -c mode

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

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

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

```
# vxdisk -o alldgs list

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

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

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

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

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

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

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

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

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

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

6 Create a VxFS file system in this volume:

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

7 Create a mount point for the shared file system:

```
# mkdir /oradata
```

8 From the same node, mount the file system:

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

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

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

See "Sample main.cf file for configuring a volume and file system under VCS" on page 121.

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

```
# chown oracle:oinstall /oradata

# chmod 775 /oradata
```

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

11 Create the Oracle database using the Oracle documentation.

Sample main.cf file for configuring a volume and file system under VCS

The main.cf file is located in the folder /etc/VRTSvcs/conf/config.

To configure a volume and file system under VCS, update the VCS configuration file, main.cf, given below:

```
// Sample main.cf assumes the following configuration:
// CRS_HOME on local VxFS file system
// ORACLE_HOME on CFS
// OCR and Vote-Disk on CVM
// Oracle database tablespaces on CVM or CFS

include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
```

```
cluster rac_cluster101 (
        UserNames = { admin = bopHo}
        Administrators = { admin }
        )

system galaxy (
        )

system nebula (
        )

// CRS_HOME on galaxy on local VxFS file system
group crshome_grp_galaxy (
        SystemList = { galaxy  = 0 }
        AutoFailOver = 0
        AutoStartList = { galaxy }
        )

        DiskGroup crshome_voldg_galaxy (
                DiskGroup = crsbindg_galaxy
                )

        Mount crshome_mnt_galaxy (
                MountPoint = "/oracle/crsbin"
                BlockDevice = "/dev/vx/dsk/crsbindg_galaxy/crsbinvol"
                FSType = vxfs
                MountOpt = rw
                FsckOpt = "-y"
                )

        Volume crshome_vol_galaxy (
                DiskGroup = crsbindg_galaxy
                Volume = crsbinvol
                )

        crshome_mnt_galaxy requires crshome_vol_galaxy
        crshome_vol_galaxy requires crshome_voldg_galaxy

// CRS_HOME on nebula on local VxFS file system
group crshome_grp_nebula (
        SystemList = { nebula  = 0 }
        AutoFailOver = 0
```

```
        AutoStartList = { nebula }
        )

        DiskGroup crshome_voldg_nebula (
                DiskGroup = crsbindg_nebula
                )

        Mount crshome_mnt_nebula (
                MountPoint = "/oracle/crsbin"
                BlockDevice = "/dev/vx/dsk/crsbindg_nebula/crsbinvol"
                FSType = vxfs
                MountOpt = rw
                FsckOpt = "-y"
                )

        Volume crshome_vol_nebula (
                DiskGroup = crsbindg_nebula
                Volume = crsbinvol
                )

        crshome_mnt_nebula requires crshome_vol_nebula
        crshome_vol_nebula requires crshome_voldg_nebula

// CVM group for:
//      OCR and Vote-Disk on CVM
//      Oracle database tablespaces on CVM or CFS
group cvm (
        SystemList = { galaxy = 0, nebula = 1 }
        AutoFailOver = 0
        Parallel = 1
        AutoStartList = { galaxy, nebula }
        )

// OCR and Vote-disk on CVM
CVMVolDg ocrvote_voldg (
        Critical = 0
        CVMDiskGroup = ocrvotedg
        CVMVolume = { ocrvol, vdvol }
        CVMActivation = sw
        )

// Oracle database tablespaces on CFS (Not to be used for
Oracle database on CVM)
```

```
CFSMount oradata_mnt (
        Critical = 0
        MountPoint = "/oradata"
        BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
        )

// Oracle database tablespaces on CFS or CVM
CVMVolDg oradata_voldg (
        Critical = 0
        CVMDiskGroup = oradatadg
        CVMVolume = { oradatavol }
        CVMActivation = sw
        )

        requires group orahome_grp online local firm

        oradata_mnt requires oradata_voldg

// ORACLE_HOME group
group orahome_grp (
        SystemList = { galaxy = 0, nebula = 1 }
        AutoFailOver = 0
        Parallel = 1
        AutoStartList = { galaxy, nebula }
        )

// ORACLE_HOME on CFS
CFSMount orabin_mnt (
        Critical = 0
        MountPoint = "/oracle/orabin"
        BlockDevice = "/dev/vx/dsk/orabindg/orabinvol"
        )

CVMVolDg orabin_voldg (
        Critical = 0
        CVMDiskGroup = orabindg
        CVMVolume = { orabinvol }
        CVMActivation = sw
        )

CFSfsckd vxfsckd (
        )
```

```
CVMCluster cvm_clus (
        CVMClustName = rac_cluster101
        CVMNodeId = { galaxy = 1, nebula = 2 }
        CVMTransport = gab
        CVMTimeout = 200
        )

CVMVxconfigd cvm_vxconfigd (
        Critical = 0
        CVMVxconfigdArgs = { syslog }
        )

        orabin_mnt requires orabin_voldg
        orabin_mnt requires vxfsckd
        orabin_voldg requires cvm_clus
        vxfsckd requires cvm_clus
        cvm_clus requires cvm_vxconfigd
```

Increasing the peer inactivity timeout of LLT

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

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

For instructions, refer to the following technote:

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

Configuring LLT to use bonded network interfaces (optional)

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

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

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

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

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

For more information, see the following technote:

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

For more information, see the following documents:

Veritas Volume Manager Administrator's Guide

Veritas Storage Foundation for Cluster File System Administrator's Guide

Veritas Storage Foundation Installation Guide

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

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

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

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

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

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

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 the status (success or failure) of each package. 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`, `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 129.

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              ONLINE
```

If the State value is running, VCS is successfully installed and running on that node.

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"
include "OracleASMTypes.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.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

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

About removing Veritas Storage Foundation

This section covers uninstallation requirements and steps to uninstall the Veritas software.

Only users with superuser privileges can uninstall Veritas Storage Foundation.

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.

To shut down and remove the installed Veritas Storage Foundation packages

1. 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
   ```

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

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

   For Veritas Storage Foundation

   ```
   # ./uninstallsf
   ```

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

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

4. The uninstall script prompts you to confirm the uninstall. If you respond yes, the processes are stopped and the packages are uninstalled.

   The uninstall script creates log files and displays the location of the log files.

5. Most packages have kernel components. In order to ensure complete removal, a system reboot is recommended after all packages have been removed.

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.
2. Stop the cluster:

   ```
   # hastop -all
   ```

   Do not use the `hastop -force` command to stop VCS.

3 Change directory to `/opt/VRTS/install`:

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

4 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
```

5 Enter the system names to uninstall SFCFS.

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

6 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.

Uninstalling Veritas Volume Replicator (VVR) involves performing the following tasks in the order indicated:

- Removing the Replicated Data Set
- Removing the VVR packages

For more information about VVR commands, 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.

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
```

Removing the VVR packages

Use the uninstall program to remove the VVR software 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
```

If you do not have any other Veritas products installed on the system, you can remove the `/etc/vx` directory, the `/usr/lib/vxvm` directory, and the `/opt/VRTS*` directories.

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.

Chapter 10

Uninstalling SFCFS for Oracle RAC from a cluster

This chapter includes the following topics:

Preparing to uninstall SFCFS for Oracle RAC from a cluster

Perform the steps in the following procedure before you uninstall SFCFS for Oracle RAC from a cluster.

To prepare to uninstall SFCFS for Oracle RAC from a cluster

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

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

   ```
   /opt/VRTSvcs/bin
   ```
3. On all the nodes, stop the CFS-dependant applications that are not under VCS control using application specific commands.

   For example, to stop Oracle Clusterware:

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

4 Stop VCS:

```
# hastop -all
```

5 Verify that port h is not open:

```
# gabconfig -a
```

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

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

7 Unmount all Storage Checkpoints and file systems:

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

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

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

To verify that no volumes are open:

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

Uninstalling SFCFS for Oracle RAC from a cluster

Perform the steps in the following procedure to uninstall SFCFS for Oracle RAC from a cluster.

To uninstall SFCFS for Oracle RAC from a cluster

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

2 Navigate to the directory that contains the uninstallation program:

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

3 Start the uninstallation program:

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

4 Back up the following configuration files:

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

5 Reboot the nodes:

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

Appendix A

Installation scripts

This appendix includes the following topics:

About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.0RU3 provides several installation scripts.

To install a fresh installation on a system, or to upgrade from Veritas Storage Foundation and High Availability Solutions version prior to 5.0RU3, 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 39.

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 Storage Foundation (SF)	`installsf`
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`

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, except where otherwise noted.

See "About installation scripts" on page 157.

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.
-hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
–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.
–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.

Table A-1 Available command line options *(continued)*

Command Line Option	Function
–precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
–responsefile *response_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.
–rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
-security	Enable or disable Symantec Product Authentication Service in a VCS cluster that is running. Install and configure Root Broker for Symantec Product Authentication Service. For more information about Symantec Product Authentication Service in a VCS cluster, see the *Veritas Cluster Server Installation Guide.*
-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.
–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.

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
VRTSddlpr	Veritas Device Discovery Layer Services Provider Provides the necessary management backend required to administer Veritas Volume Manager (VxVM) Dynamic Multipathing (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 for administering VxFS and other platform file systems 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 (SDK), which includes headers, libraries, and sample code. The SDK 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.
VRTScscm	Veritas Cluster Server Cluster Manager	Required Depends on VRTSvcs and VRTSjre15.
VRTScscw	Veritas Cluster Server configuration wizards	Required Depends on VRTSvcsag and VRTSjre15.
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.
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.

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
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.
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
Other RPMs		
VRTSdbms3	Veritas Shared DBMS	Required
VRTSdsa	Veritas Datacenter Storage Agent	Required
VRTSmapro-common	Veritas Storage Foundation GUI for Mapping	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
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
VRTSgcspr	Veritas SAN Global Configuration Server Object Bus provider	Required

Table B-1 Storage Foundation RPMs *(continued)*

RPM	Contents	Required/Optional
windows/ vrtsobgui.msi	Veritas Enterprise Administrator for Windows	Optional

Obsolete RPMs in Veritas Storage Foundation Cluster File System for Oracle RAC

The following RPMs were included in previous releases of Veritas Storage Foundation Cluster File System for Oracle RAC 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.

Suggested solution: You need to set up the systems to allow remote access using `ssh` or `rsh`.

Note: Remove remote shell permissions after completing the SFCFS for Oracle RAC 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.

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" 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