

Veritas Storage Foundation[™] Installation Guide

Solaris

5.0

Veritas Storage Foundation Installation Guide

Copyright © 2006 Symantec Corporation. All rights reserved.

Storage Foundation 5.0

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

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

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

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

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

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Solaris is a trademark of Sun Microsystems, Inc.

Licensing and registration

Veritas Volume Manager is a licensed product. See the *Veritas Storage Foundation Installation Guide* for license installation instructions.

Technical support

For technical assistance, visit <http://support.veritas.com> and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1	Preinstallation Instructions	
	Veritas Storage Foundation product suites	12
	General installation requirements	15
	Centralized management considerations	15
	Symantec product licensing	15
	Supported Solaris operating systems	16
	Software and hardware requirements	16
	Database requirements	17
	Disk space	18
	/opt directory	19
	Environment variables	19
	Prerequisites for Remote and Cluster Installation and Uninstallation	20
	Release Notes	20
	Veritas Volume Manager requirements	20
	Taking a disk inventory	20
	DMP and Alternate Pathing	22
	Array Support Libraries (ASL)	22
	Veritas Enterprise Administrator	23
	VMSA and VEA co-existence	24
	Cluster environment requirements	24
	Veritas File System requirements	25
	Converting from QuickLog to Multi-Volume Support	25
	Locating installation and upgrade instructions	26
Chapter 2	Installing a Veritas Storage Foundation product	
	Installation requirements	30
	About mapping services and performance statistics for databases	30
	About the operating system	30
	Overview of Veritas Storage Foundation installation tasks	31
	Installing a Veritas Storage Foundation product	31
	Mounting the software disc	32
	Installing using the product installer	33
	Configuring a Veritas Storage Foundation product on a cluster	35
	Completing the installation procedure	39
	Installing the Veritas Enterprise Administrator client	41

Verifying the Veritas Storage Foundation installation	43
Checking Volume Manager processes	43
Checking VxFS installation	43
Checking cluster operation	44
Loading and unloading the file system module	47
Installing the Veritas software with JumpStart	47
Unzipping the packages	47
Installation order	49
Summary of JumpStart installation tasks	52
Installing language packages	53
Installing the French and Simplified Chinese language packages	53
Installing the Japanese language packages	55

Chapter 3

Installing Veritas Storage Foundation Basic

About Veritas Storage Foundation Basic	58
Storage Foundation Basic technical support	58
Storage Foundation Basic system requirements	59
Installing Veritas Storage Foundation Basic	61
Upgrading Veritas Storage Foundation Basic	64
Planning the upgrade	64
Uninstalling Veritas Storage Foundation Basic	65

Chapter 4

Upgrading a Veritas Storage Foundation product

Preparing to upgrade Veritas Volume Manager	68
Upgrading VxVM on a system with Alternate Pathing	68
Pre-upgrade requirements	68
Determining if the root disk is encapsulated	70
Adding Sun T3+ and T3 arrays as JBODS of type A/P	70
Unsuppressing DMP for EMC PowerPath disks	71
VxVM and Solaris release support matrix	78
VxVM and Solaris operating system upgrade paths	79
Determining if the root disk is encapsulated	79
Retaining older Veritas Volume Manager releases	80
Upgrading the disk group version separately	80
Preparing to upgrade Veritas File System	81
Upgrading VxFS disk layout versions	81
Make sure the file systems are clean	83
Upgrading Veritas Storage Foundation	83
Upgrading Veritas Storage Foundation on Solaris 8, 9, or 10	84
Upgrading Veritas Storage Foundation on an encapsulated root disk	85
Upgrading Veritas Storage Foundation and earlier Solaris releases	88

Upgrading from Database Edition 3.5 for Oracle with Database FlashSnap	91
Migrating from /etc/vx/vxdba to /var/vx/vxdba for DB2 and Oracle	92
Upgrading the VEA client on a Solaris system	92
Upgrading the VEA client on a Windows system	93
Upgrading to the new repository database for DB2 and Oracle	94
Verifying the Veritas Storage Foundation upgrade	95
Checking Volume Manager processes	95
Checking cluster operation	95
Upgrading from SUNWvxxvm	97
Upgrading language packages	99

Chapter 5 Configuring the Veritas Storage Foundation software

Database configuration requirements	101
Creating and configuring the repository database for DB2 and Oracle	102
Configuring Veritas Storage Foundation	105
Setting administrative permissions	105
Veritas Enterprise Administrator	106
vxtunefs command permissions and Cached Quick I/O	109
Configuring Veritas Volume Manager	110
Using vxinstall to configure Veritas Volume Manager	110
Using Dynamic Multipathing with EMC Symmetrix Disk Arrays	112
Preventing multipathing/suppress devices from VxVM's view	112
About hot-relocation	114
Starting and enabling the configuration daemon	114
Starting the volume I/O daemon	115
Starting the VEA server	116
Starting the VEA client	116
Enabling cluster support in VxVM (optional)	117
Configuring shared disks	118
Reserving shared disks	121
Adding new array support	121
About hot-relocation	121
About placing disks in another disk group	122
Protecting your system and data	122
Configuring Veritas File System	123

Chapter 6 Uninstalling the Veritas Storage Foundation software

Uninstallation requirements	126
Remote uninstallation	126
Veritas Volume Manager	126
Veritas File System	132

- Dropping the repository database for DB2 and Oracle 133
- Shutting down cluster operations 134
- Uninstalling Veritas Storage Foundation 134
- Uninstalling the Veritas Enterprise Administrator client 136
- Uninstalling language packages 136
 - Uninstalling language packages using the pkgadd command 136

Chapter 7 Installing, upgrading and uninstalling VxVM and VxFS

- Installing Veritas Volume Manager 140
 - Installing VxVM using the product installer 140
 - Installing VxVM using the pkgadd command 142
 - Installing VxVM using JumpStart 144
- Installing Veritas File System 145
 - Installing VxFS using the product installer 146
 - Installing VxFS using the pkgadd command 147
- Upgrading Veritas Volume Manager 148
 - Upgrading VxVM only 148
 - Migrating a standalone host to a managed host 154
 - Upgrading VxVM using upgrade scripts 154
 - Upgrading VxVM and the Solaris OS 154
 - Upgrading the Solaris OS only 156
 - Upgrading VxVM and/or Solaris using Live Upgrade 158
- Upgrading Veritas File System 166
 - Upgrading VxFS using the product installer 166
 - Upgrading VxFS using Live Upgrade 167
- Uninstalling Veritas Volume Manager 168
 - Uninstalling VxVM using the uninstallation script 168
 - Uninstalling VxVM using the pkgrm command 169
- Uninstalling Veritas File System 169
 - Uninstalling VxFS using the uninstallation script 169
 - Uninstalling VxFS using the pkgrm command 170

Appendix A Installation script options

- Installation Script Options 174

Appendix B Veritas Storage Foundation install packages

- Veritas Storage Foundation install packages 177
 - Japanese language packages 182
 - Chinese language packages 185
 - French language packages 185
- Package installation for manual installations 187

Unzipping the packages	187
Installing the packages using the pkgadd command	188
Veritas Storage Foundation	189
Veritas Storage Foundation for DB2	189
Veritas Storage Foundation for Oracle	190
Veritas Storage Foundation for Sybase	190
Veritas Volume Manager	191
Veritas File System	191

Appendix C Configuring the Symantec License Inventory Agent

About the Symantec License Inventory Manager	194
When the Symantec License Inventory Agent is installed	195
When the server and access points are installed	195
What you can do with the agent after it is installed	195
How to remove the agent	196
How to order the Symantec License Inventory Manager license and media kit	197

Preinstallation Instructions

Topics covered in this chapter include:

- [“Veritas Storage Foundation product suites”](#) on page 12
- [“General installation requirements”](#) on page 15
- [“Veritas Volume Manager requirements”](#) on page 20
- [“Veritas File System requirements”](#) on page 25
- [“Locating installation and upgrade instructions”](#) on page 26

Follow the preinstallation instructions outlined in this chapter if you are installing one of the following Veritas Storage Foundation products by Symantec:

- Veritas Storage Foundation Standard, Enterprise, and Enterprise High Availability (HA) Editions
- Veritas Storage Foundation for DB2 Standard, Enterprise, and Enterprise High Availability (HA) Editions
- Veritas Storage Foundation for Oracle Standard, Enterprise, and Enterprise High Availability (HA) Editions
- Veritas Storage Foundation for Sybase Standard, Enterprise, and Enterprise High Availability (HA) Editions
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)

Note: If you are installing Veritas Volume Replicator, see the *Veritas Volume Replicator Installation Guide* on the product disc. The *Veritas Volume Replicator Installation Guide* explains how to install the product and directs you to the Veritas Volume Replicator documentation.

After reviewing the information in this chapter, see “[Installing a Veritas Storage Foundation product](#)” on page 31 for information on installing the software for the first time. See “[Upgrading Veritas Storage Foundation](#)” on page 83 if you are upgrading an existing product.

Veritas Storage Foundation product suites

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

Table 1-1 Contents of Veritas Storage Foundation products

Storage Foundation version	Products and features
Storage Foundation Basic	Veritas File System Veritas Volume Manager
Storage Foundation Standard	Veritas File System Veritas Volume Manager 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 Optionally licensed features: Veritas Volume Replicator

Table 1-1 Contents of Veritas Storage Foundation products

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

Table 1-1 Contents of Veritas Storage Foundation products

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

Note: If you are installing Veritas Volume Replicator, see the *Veritas Volume Replicator Installation Guide* on the product disc. The *Veritas Volume Replicator Installation Guide* explains how to install the product and directs you to the Veritas Volume Replicator documentation.

General installation requirements

Before installing a Veritas Storage Foundation product, make sure you understand and comply with the basic requirements of the software.

Centralized management considerations

Veritas Storage Foundation Management Server by Symantec ties together the Storage Foundation product offerings to ensure that the hosts in your data center use storage as efficiently as possible. You can use it to centrally monitor, visualize, and manage Storage Foundation hosts and generate reports about the hosts and the storage resources they consume.

The central console seamlessly integrates a wide range of management tasks like monitoring and reporting.

SF Management Server also offers customizable policy-based management that helps you automate:

- notification
- recovery
- other user-definable actions

SF Management Server is not available on the Storage Foundation and High Availability Solutions release and must be obtained separately. For information on ordering SF Management Server, visit:

<http://www.symantec.com/enterprise/sfms>

Symantec product licensing

This product includes 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 at:

<http://www.veritas.com/buy/vLicense/vLicenseHome.jhtml>

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 install the Symantec software products for which you have purchased a license.

Supported Solaris operating systems

Veritas Storage Foundation can only be installed on a system running Solaris 8 (32- or 64-bit), 9 (32- or 64-bit), or 10 (64-bit). Installing this product on any other Solaris release will fail. If necessary, upgrade Solaris before you install the Veritas products.

See “[Upgrading Veritas Storage Foundation and earlier Solaris releases](#)” on page 88.

Software and hardware requirements

For information on hardware requirements, see the *Veritas Storage Foundation Release Notes*.

Database requirements

The table identifies supported database and Solaris combinations if you plan to use Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, or Veritas Storage Foundation for Sybase.

Table 1-2 Supported database and Solaris combinations

Database	Solaris 8 (64-bit)	Solaris 8 (32-bit)	Solaris 9 (64-bit)	Solaris 9 (32-bit)	Solaris 10 (64-bit)
DB2 UDB					
8.1 ESE with FixPak 6 or lower	Yes	Yes	Yes	Yes	Yes
8.2 (or 8.1 ESE with FixPak 7 or higher)	Yes	Yes	Yes	Yes	No
8.2.2 with FixPak 9	No	No	No	No	Yes
9.1	Yes	Yes	Yes	Yes	No
Oracle					
9.2 (64-bit)	Yes	Yes	Yes	Yes	Yes
9.2 (32-bit)	Yes	No	Yes	No	Yes
10.1	Yes	Yes	Yes	Yes	Yes
Sybase					
12.5	Yes	Yes	Yes	Yes	Yes
15	Yes	Yes	Yes	Yes	Yes

Additional DB2 information

With DB2 8.1 ESE, the Database Partitioning feature is only supported in an SMP configuration.

At this time, only Quick I/O and Cached Quick I/O are supported in a Massively Parallel Processor (MPP) configuration. No other Veritas Storage Foundation for DB2 tools are supported in an MPP environment.

Refer to the appropriate DB2 installation guide that accompanied your DB2 software for additional information. IBM maintains DB2 UDB FixPaks for download at the following location:

<http://www.ibm.com/software/data/db2/udb/support/>

Additional Oracle information

In order to use Veritas Extension for Oracle Disk Manager, you must use Oracle9i release 9.2 or later. Refer to Oracle bug number 1824061 for more details.

To use the Storage Mapping functionality, you must install Oracle 9.2.0.5 or higher.

Refer to the appropriate Oracle installation guide that accompanied your Oracle software for additional preinstallation information.

Additional Sybase information

Refer to the appropriate Sybase installation guide that accompanied your Sybase software for additional information.

Disk space

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) option of the product installer to determine whether there is sufficient space.

The following table shows the approximate disk space used by the Storage Foundation products for all (both the required and optional) packages:

Disk space requirements

Product name	Minimum space required (without optional packages)	Maximum space required (including all packages)
Storage Foundation Standard or Enterprise	802MB	1197MB
Storage Foundation Enterprise HA	1125MB	1197MB
Storage Foundation for DB2 Standard or Enterprise	1125MB	1160MB
Storage Foundation for DB2 Enterprise HA	1404MB	1520MB

Disk space requirements

Product name	Minimum space required (without optional packages)	Maximum space required (including all packages)
Storage Foundation for Oracle Standard or Enterprise	1167MB	1567MB
Storage Foundation for Oracle Enterprise HA	1495MB	1567MB
Storage Foundation for Sybase Standard or Enterprise	811MB	846MB
Storage Foundation for Sybase Enterprise HA	1090MB	1206MB

/opt directory

The directory `/opt` must exist, be writable and must not be a symbolic link. When upgrading, volumes not temporarily converted by the `upgrade_start` script will be unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

Environment variables

Most of the commands used in the installation are in the `/sbin` or `/usr/sbin` directory. However, there are additional variables needed to use the Veritas Storage Foundation product after installation. Add the following directories to your `PATH` environment variable:

- If you are using Bourne or Korn shell (`sh` or `ksh`), use the commands:

```
$ PATH=$PATH:/usr/sbin:/opt/VRTS/bin:/opt/VRTSvxfs/sbin:\
/opt/VRTSob/bin:/opt/VRTScvs/bin:/etc/vx/bin
$ MANPATH=/usr/share/man:/opt/VRTS/man:$MANPATH
$ export PATH MANPATH
```

- If you are using a C shell (`csh` or `tcsh`), use the commands:

```
% set path = ( $path /usr/sbin /opt/VRTSvxfs/sbin \
/opt/VRTScvs/bin /opt/VRTSob/bin /opt/VRTS/bin /etc/vx/bin )
% setenv MANPATH /usr/share/man:/opt/VRTS/man:$MANPATH
```

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

Prerequisites for Remote and Cluster Installation and Uninstallation

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. See “Configuring SSH or RSH Before Installing Veritas Products” in the *Getting Started Guide* for more information.

Release Notes

Read the *Release Notes* for all products included with this product. Portable Document Format (.pdf) versions of the *Release Notes* are included on the software disc in the `storage_foundation/release_notes` directory and on the documentation disc that came with your software.

Because product *Release Notes* are not installed by any packages, it is recommended that you copy them from the disc to the `/opt/VRTS/docs` directory on your system so that they are available for future reference.

Veritas Volume Manager requirements

Make sure you meet all Volume Manager-related requirements before installing or upgrading a Veritas Storage Foundation product or Veritas Volume Manager.

Taking a disk inventory

How to select disks

Decide which disks you want to place under VxVM control. The other disks in your configuration are not affected. Disks may be brought under VxVM control in two ways:

- Encapsulated disks—Data in all existing file systems and partitions on the disk are preserved.
- Initialized disks—Data on the disks is removed.

Verifying disk contents

Verify the disk contents. Answer the following questions and list the data for your convenience.

- 1 Make sure you are aware of the contents of each disk. Determine which disks can be encapsulated (data is preserved) or initialized (data is removed).

2 Do you want to place the system root disk under VxVM control?

3 Do you want to either encapsulate or initialize *all* disks on a controller together? Identify the controllers (for example c0t0d0).

4 Identify the disks to be encapsulated, initialized, or excluded, in a table similar to the following.

Disk ID	Encapsulate, Initialize, Exclude

5 Verify that the disks that are to be encapsulated by VxVM have two free partitions and a recommended amount of 32MB of free space. The free space must be at the beginning or end of the disk and must not belong to a partition. This space is used for storing disk group configurations and a disk label that ensures VxVM can identify the disk, even if it is moved to a different address or controller. It is also used to keep track of disk configuration and to ensure correct recovery. The boot disk is a special case. If no other space is available, VxVM attempts to allocate space usually reserved for swap by shrinking the swap partition. This process is known as *swap relocation* and, if necessary, happens automatically during root disk encapsulation.

Note: Although it is possible to put disks with no free space under VxVM control, this is only used as a migration strategy. Many VxVM capabilities based on disk identity are available only for disks with the required free space. See the `vxdisk(1M)` manual page for information on the `nopriv` disk type.

- 6 If you are encapsulating the boot (root) disk:
 - a Before encapsulating your boot disk, set the `EEPROM` variable `use-nvramrc?` to `true`. This will enable VxVM to take advantage of boot disk aliases to identify the mirror of the boot disk if a replacement is needed. If this variable is set to `false`, you must determine which disks are bootable yourself. Set this variable to `true` as follows:

```
# eeprom "use-nvramrc?=true"
```

If your root disk is connected over fabric, you should check the Hardware Compatibility List at <http://support.veritas.com> to see if your device type is supported for boot encapsulation
 - b Use the `prtvtoc(1M)` command to record the layout of the partitions on the unencapsulated boot disk (`/dev/rdisk/c0t0d0s2` in this example):

```
# prtvtoc /dev/rdisk/c0t0d0s2
```

Record the output from this command for future reference.
 - c To encapsulate the boot disk, tag the swap partition as `swap` so that it is possible to dump to that partition later. See `format(1M)` for information on tagging the swap partition.

Note: If the path to an aliased boot device is different from the path in the `/devices` directory, aliases may not function correctly. This might happen if the pathname is incorrectly entered at the command line, or if the device's pathname changes at a later date.

DMP and Alternate Pathing

Note: The information in this section only applies if you are running Solaris 8.

VxVM 5.0 does not allow Dynamic Multipathing (DMP) to co-exist with Sun's Alternate Pathing (AP) software. The AP driver must be disabled before installing or upgrading VxVM.

Array Support Libraries (ASL)

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software packages. You can obtain ASL packages from:

- The VxVM release package on the software disc
- The disk array provided by the vendor
- The Veritas Technical Support site, <http://support.veritas.com>

Veritas Enterprise Administrator

Veritas Storage Foundation must be installed and run on a UNIX (Solaris) machine. The Veritas Enterprise Administrator (VEA) client can be installed and run on any Solaris, Windows XP, Windows NT, Windows ME, Windows 2000, or Windows 98 machine that supports the Java Runtime Environment.

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.

One of the following needs to be installed and running on the client:

- VERITAS Enterprise Administrator (VRTSobgui, VRTSat, VRTSpbx, and VRTSicSCO)
These are the client packages for Solaris and are located in the *product_name/pkgs* directory.
- VERITAS Enterprise Administrator for Windows (*windows\VRTSobgui.msi*)
This is the client package for Windows.

See the *Veritas Storage Foundation Release Notes* for patch information before you install VEA.

Checking minimum requirements

The following are minimum system recommendations for the GUI:

Solaris	SPARCstation 5 with 64M memory
Windows XP, NT, Me, 2000, or 98	300MHz Pentium with 256M memory

For the VEA client to function properly with the Java Runtime Environment 1.5 (JRE 1.5), install the latest patches for JRE 1.5. To obtain patch information, see the Sun Microsystems Web site.

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.

Note: 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
```

Cluster environment requirements

Use these steps if the configuration contains a cluster, which is a set of hosts that share a set of disks.

To configure a cluster

- 1 Obtain a license for the optional VxVM cluster feature for a Sun Cluster from your Sun Customer Support channel.
- 2 If you plan to encapsulate the root disk group, decide where you want to place it for each node in the cluster. The root disk group, usually aliased as `bootdg`, contains the volumes that are used to boot the system. VxVM sets `bootdg` to the appropriate disk group if it takes control of the root disk. Otherwise `bootdg` is set to `nodg`. To check the name of the disk group, enter the command:

```
# vxvg bootdg
```
- 3 Decide the layout of shared disk groups. There may be one or more shared disk groups. Determine how many you wish to use.
- 4 If you plan to use Dirty Region Logging (DRL) with VxVM in a cluster, leave a small amount of space on the disk for these logs. The log size is proportional to the volume size and the number of nodes. Refer to the *Veritas Volume Manager Administrator's Guide* and the *Veritas Storage Foundation Cross-Platform Data Sharing Administrator's Guide* for more information on DRL.
- 5 Install the license on every node in the cluster.
See [“Enabling cluster support in VxVM \(optional\)”](#) on page 117.

Veritas File System requirements

Before installing Veritas File System:

- Review the *Veritas Storage Foundation Release Notes*, `sf_notes.pdf`, located under the `storage_foundation/release_notes` directory on the Veritas software disc. Because product release notes are not installed by any packages, Veritas recommends that you copy them from the software disc to the `/opt/VRTS/docs` directory so that they are available for future reference.
- In the *Veritas Storage Foundation Release Notes*, review the information on VRTSexplorer and installing the `VRTSspt` package. `VRTSspt` is a group of tools for troubleshooting a system and collecting information on its configuration. The tools can gather VxFS metadata information and establish various benchmarks to measure file system performance. The tools are not required for the 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, and should be used only under the direction of a Veritas Technical Support Engineer.
- Ensure that the `/opt` directory exists and has write permissions for `root`.
- Install all the latest required Solaris patches listed in the *Veritas Storage Foundation Release Notes*.

Converting from QuickLog to Multi-Volume Support

The 4.1 release of the Veritas File System was the last major release to support QuickLog. The Version 6 disk layout does not support QuickLog. The functionality provided by the Veritas Multi-Volume Support (MVS) feature replaces most of the functionality provided by QuickLog.

The following procedure describes how to convert from QuickLog to MVS. Unlike QuickLog, which allowed logging of up to 31 VxFS file systems to one device, MVS allows intent logging of only one file system per device. Therefore, the following procedure must be performed for each file system that is logged to a QuickLog device if Version 6 disk layout is used.

Note: The QuickLog device did not need to be related to the file system. For MVS, the log volume and the file system volume must be in the same disk group.

To convert Quicklog to MVS

- 1 Select a QuickLog-enabled file system to convert to MVS and unmount it.

```
# umount myfs
```
- 2 Detach one of the QuickLog volumes from the QuickLog device that the file system had been using. This volume will be used as the new intent log volume for the file system.

```
# qlodetach -g diskgroup log_vol
```
- 3 Create the volume set.

```
# vxvset make myvset myfs_volume
```
- 4 Mount the volume set.

```
# mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1
```
- 5 Upgrade the volume set's file system to Version 6 disk layout. See the *Veritas Storage Foundation Installation Guide*. See the vxfsconvert(1M) and vxupgrade(1M) manual pages.

```
# vxupgrade -n 6 /mnt1
```
- 6 Add the log volume from [step 2](#) to the volume set.

```
# vxvset addvol myvset log_vol
```
- 7 Add the log volume to the file system. The size of the volume must be specified.

```
# fsvoladm add /mnt1 log_vol 50m
```
- 8 Move the log to the new volume.

```
# fsadm -o logdev=log_vol,logsize=16m /mnt1
```

Locating installation and upgrade instructions

Follow the installation or upgrade instructions based on your site's needs.

Table 1-3 Installation and upgrade instructions for Veritas Storage Foundation products

To install or upgrade...	Go to...
Storage Foundation and Storage Foundation HA	“Installing a Veritas Storage Foundation product” on page 31 or “Upgrading Veritas Storage Foundation” on page 83
Storage Foundation for DB2 and Storage Foundation for DB2 HA	“Installing a Veritas Storage Foundation product” on page 31 or “Upgrading Veritas Storage Foundation” on page 83

Table 1-3 Installation and upgrade instructions for Veritas Storage Foundation products

To install or upgrade...	Go to...
Storage Foundation for Oracle and Storage Foundation for Oracle HA	“Installing a Veritas Storage Foundation product” on page 31 or “Upgrading Veritas Storage Foundation” on page 83
Storage Foundation for Sybase and Storage Foundation for Sybase HA	“Installing a Veritas Storage Foundation product” on page 31 or “Upgrading Veritas Storage Foundation” on page 83
Solaris Operating System	“Upgrading the Solaris OS only” on page 156
Volume Manager	“Installing Veritas Volume Manager” on page 140 or “Upgrading Veritas Volume Manager” on page 148
File System	“Installing Veritas File System” on page 145 or “Upgrading Veritas File System” on page 166

Installing a Veritas Storage Foundation product

Topics covered in this chapter include:

- [“Installation requirements”](#) on page 30
- [“Overview of Veritas Storage Foundation installation tasks”](#) on page 31
- [“Installing a Veritas Storage Foundation product”](#) on page 31
- [“Verifying the Veritas Storage Foundation installation”](#) on page 43
- [“Loading and unloading the file system module”](#) on page 47
- [“Installing the Veritas software with JumpStart”](#) on page 47
- [“Installing language packages”](#) on page 53

Note: Only users with superuser (`root`) privileges can install Veritas products.

If you already have an earlier version of any of the Veritas products installed, see [“Upgrading Veritas Storage Foundation”](#) on page 83.

Installation requirements

Veritas software and operating system installation requirements are described in the following sections.

About mapping services and performance statistics for databases

- You must install Veritas Array Integration Layer (VAIL), Generic Array for Mapping Services by Symantec (VRTSgapms), and Veritas Mapping Services (VxMS) if you want to use deep mapping services and performance statistics for supported storage arrays.
- Veritas Storage Foundation for Sybase supports file-level mapping.
- Install the EMC Solutions Enabler (SYMCLD) before you install VAIL. If you install Solutions Enabler after you install VAIL, rescan the EMC Symmetrix arrays so that they can be discovered.

See the *Veritas Storage Foundation DB2 Administrator's Guide*.

See the *Veritas Storage Foundation Oracle Administrator's Guide*.

About the operating system

If patches to the Solaris operating system are required, the patches should be applied just prior to installing the Veritas products. Patches may be required to resolve Solaris kernel, product performance, or other issues.

The system should be in a quiescent state before adding patches.

Note: See product *Release Notes* for last minute information on patches required for Veritas, Solaris (Sun), IBM (DB2), Oracle, or Sybase software.

Overview of Veritas Storage Foundation installation tasks

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

- Reviewing *Release Notes* for the most current product information.
- Reviewing preinstallation requirements and making any necessary modifications.
See [“Preinstallation Instructions”](#) on page 11.
- Adding Solaris operating system patches, if needed. Refer to the *Release Notes* for specific software patch information.
- Obtaining a license key.

Note: Veritas Storage Foundation Basic does not require a license key.

- Installing or upgrading the Veritas Storage Foundation software packages. For new installations, use the instructions in this chapter.
See [“Upgrading a Veritas Storage Foundation product”](#) on page 67.

After you complete the installation procedure, you need to initialize, set up, and use the Veritas software shipped with the Veritas Storage Foundation products. See [“Configuring the Veritas Storage Foundation software”](#) on page 101.

Note: Installing Veritas Storage Foundation using the product installer will automatically configure the software. If you install using an alternative method, you will have to run the product installer to configure the software.

Installing a Veritas Storage Foundation product

This section describes how to install the Veritas Storage Foundation products for the first time. These instructions apply to:

- Veritas Storage Foundation
- Veritas Storage Foundation for DB2
- Veritas Storage Foundation for Oracle
- Veritas Storage Foundation for Sybase.

The product installer simplifies the process of installing Veritas software and Veritas strongly recommends this method of installation.

After you complete the installation procedure, review important details about initializing (where required), setting up, and using the Veritas software shipped with Veritas Storage Foundation.

See [“Configuring the Veritas Storage Foundation software”](#) on page 101.

Mounting the software disc

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

To mount the software disc

- 1 Log in as superuser.
- 2 Place the Veritas software disc containing your product into a DVD drive connected to your system.
- 3 If Solaris volume management software is running on your system, the software disc automatically mounts as `/cdrom/cdrom0`.
- 4 If Solaris volume management software is not available to mount the DVD, you must mount it manually. After inserting the software disc, enter:

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
```

where `c0t6d0s2` is the default address for the disc drive.

To install, use one of the following procedures:

Table 2-4 Installation instructions for Veritas Storage Foundation products

To install	See
<ul style="list-style-type: none"> ■ Veritas Storage Foundation ■ Veritas Storage Foundation for DB2 ■ Veritas Storage Foundation for Oracle ■ Veritas Storage Foundation for Sybase 	“Installing using the product installer” on page 33
Veritas Volume Manager	“Installing Veritas Volume Manager” on page 140
Veritas File System	“Installing Veritas File System” on page 145

Installing using the product installer

The product installer simplifies the process of installing Veritas software and Veritas strongly recommends this method of installation. The installer enables you to configure the product, verify preinstallation requirements, and view the product's description.

Select a product to install from the product installer menu to invoke the product installation script.

Note: If you obtained a Veritas product from an electronic download site, the single product download files do not contain the product installer, so you must use the product installation script to install the product.

See [“Installation script options”](#) on page 173.

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. There is a short delay before the script exits.

The following sample procedure is based on the installation of a Veritas Storage Foundation Enterprise HA cluster with two nodes: “host1” and “host2.” If you are installing on stand-alone systems only, some steps are unnecessary, and these are indicated. Default responses are enclosed by parentheses. Press Return to accept defaults.

To install a Veritas Storage Foundation product

- 1 Insert the product software disc into a DVD drive attached to your system and mount it (see [“Mounting the software disc”](#) on page 32).
- 2 Move to the DVD directory:

```
# cd /cdrom/cdrom0
```
- 3 From the this directory, invoke the installer script:

```
# ./installer (to install using ssh)
```

or

```
# ./installer -rsh (to install using rsh)
```

For information on installing on multiple hosts, see [“Completing the installation procedure”](#) on page 39.
- 4 Enter **I** to install a product and press Return to begin.
- 5 When the list of available products is displayed, select the product you want to install and enter the corresponding number and press Return. The product installation begins automatically.

- 6 You are prompted to enter the systems 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 product_name: host1 host2
```

- 7 Enter the product license information.

```
Enter a product_name 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)
Enter a product_name 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)
```

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.

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

- 8 You can choose to either install only required packages or all packages. Each option displays the disk space 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 filesets to conserve
diskspace. Additional filesets are typically installed to
simplify future upgrades.

1) Required Veritas Storage Foundation filesets - 928 MB
   required
2) All Veritas Storage Foundation filesets - 1063 MB
   required

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

Note: The list of optional filesets may differ depending on the license key that you entered.

Note: If you are installing Veritas Storage Foundation for Oracle, do not select the “Storage Foundation for Oracle RAC packages” option unless you have the correct license or setup.

- 9 Reboot the system.
- 10 If you are installing on a cluster, you must configure the cluster. If you are not installing on a cluster, initialize (where required), set up, and use the Veritas software shipped with Veritas Storage Foundation.
See [“Configuring a Veritas Storage Foundation product on a cluster”](#) on page 35.
See [“Configuring the Veritas Storage Foundation software”](#) on page 101.
- 11 If you are installing Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle in a stand-alone configuration, create a new repository database.
See [“Upgrading to the new repository database for DB2 and Oracle”](#) on page 94.

Configuring a Veritas Storage Foundation product on a cluster

Note: Use the procedure in this section if you are installing an HA version of the Storage Foundation software.

As the installation and configuration procedure continues, a message displays notifying you that configuring Storage Foundation at this point in the installation procedure is optional.

To configure a Storage Foundation product on a cluster

- 1 To invoke the common installer, run the `installer` command on the disc as shown in this example:

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

- 2 At the product status page, enter **C** for the Configure an Installed Product and press Return. The product installer is displayed.

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

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

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

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

- 5 When prompted to configure VCS, enter **y** to configure VCS on these systems.

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

- 6 Enter the unique cluster name and Cluster ID number.

```
Enter the unique cluster name: [?] vcs_cluster2  
Enter the unique Cluster ID number between 0-65535: [b,?] 76
```

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

```
Discovering NICs on host1 ...discovered en0 en1 en2
```

7

- 8 Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat link on host1:  
[b,?] en1  
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,?] en2  
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
```

Note: When answering **y**, be sure that the same NICs are available on each system; the installer does not verify this.

Notice that in this example, en0 is not selected for use as a private heartbeat NIC because it already in use as the public network interface. The default responses are chosen.

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

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

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

- 10 When prompted to configure the product to use Veritas Security Services, enter **y** or **n** to configure.

Note: 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 product_name to use Veritas
Security
Services? [y,n,q] (n) n
```

- 11** A message displays notifying you of the information required to add users. When prompted, set the user name and /or password for the Administrator.

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

- 12** Enter **n** if you want to decline. If you enter **y**, you are prompted to change the password. You are prompted to add another user to the cluster.

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

- 13** Enter **n** if you want to decline, enter **y** if you want to add another user. You are prompted to verify the user.

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

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

- 14** You are prompted to configure the cluster management console. Enter **y** or **n** to configure the cluster management console.

```
Do you want to configure the Cluster Management Console
[y,n,q] (n) y
```

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

```
Enter the NIC for Cluster Manager (Web Console) to use on
host1: [b,?] (hme0)
```

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

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

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

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

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

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

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

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

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

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

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

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

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

- 22 If you are installing Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, change permissions to allow a database administrator or a group of administrators access to the database tools on each system.

Do you want to add single user access on host1? [y,n,q,?] y
Enter login account name for DBA user: *account_name*
Do you want to add group access on host1? [y,n,q,?] y
Enter groupo name for DBA users: dba
Do you want to add single user access on host2? [y,n,q,?] y
Enter login account name for DBA user: *account_name*
Do you want to add group access on host2? [y,n,q,?] y
Enter groupo name for DBA users: dba

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

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

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

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

- 25 You are prompted to enable Storage Foundation Management Server Management.

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

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

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

.
. .
.

Startup completed successfully on all systems

You declined to set up the name of the default disk group for
host1.

You declined to set up the name of the default disk group
for host2.

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

```
/opt/VRTS/install/logs/installsf-7ai12i
```

When `installsf` installs software, some software may be applied rather than committed. It is the responsibility of the system administrator to commit the software, which can be performed later with the `-c` option of the `installp` command.

- 27 Proceed to “[Completing the installation procedure](#)” on page 39.

Completing the installation procedure

At this point in the installation and configuration procedure, the utility begins to install the packages on one node and copy them to any other specified nodes. The following sample output is for a Storage Foundation Enterprise HA installation.

To complete the installation procedure

- 1 If you are completing the installation procedure in a cluster configuration, skip this step. If you are completing the installation procedure on a stand-alone configuration, invoke the common installer, run the `installer` command on the disc as shown in this example:

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

At the product status page, enter **C** for the Configure an Installed Product and press Return. The product installer is displayed.

- 2 As the utility continues the procedure, you are prompted to choose whether to install on all systems simultaneously. Enter **y** or press Enter to accept simultaneous installation.

```
Would you like to install Storage Foundation Enterprise HA on  
all systems simultaneously? [y,n,q,?] (y) y
```

- 3 A message displays notifying you of successful installation. Press Enter to continue.

- 4 A message displays describing the VxVM enclosure-based naming scheme and showing which nodes are eligible. When prompted to configure this feature, enter **n** if you want to decline. If required, you can use the `vxdiskadm` command or the Veritas Enterprise Administrator to configure this feature after installation.

See the `vxdiskadm(1M)` manual page and the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide* for more information.

- 5 If you are installing an HA version of the product, a message displays indicating the VCS configuration files are being created and copied. Press Return to continue.
- 6 A message displays indicating which systems can be started and on the systems that VxVM selects as targets.
- 7 Press Return to continue. If you are installing an HA version of the product, a message displays notifying you that Cluster Server is starting. This message also contains information about configuring a default disk group.
- 8 You are now given the option of specifying the name of a default disk group that is to be assumed by Veritas Volume Manager commands if a disk group is not otherwise specified.

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

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

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 `vxctl defaultdg diskgroup` command on a system. See the `vxctl(1M)` manual page and the “Creating and administering disk groups” chapter of the *Veritas Volume Manager Administrator’s Guide* for more information.

Note: If you specify the name of a default disk group, this step does not create the disk group. After installation, you can use menu item 1 in the `vxdiskadm` utility or the Veritas Enterprise Administrator to create the disk group.

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

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

- 10 If you have a VVR license installed, you are now asked questions regarding the frequency of VVR statistics collection.

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

Enable Centralized Management? [y,n,q]

If you select to enable centralized management, you are asked a series of questions relating to the names of the systems that are to be managed.

- 12 Finally, you are asked whether you want to start the Storage Foundation processes on the target systems.

Do you want to start Veritas Storage Foundation processes now?

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

Enter **y** to start the processes on each of the systems.

- 13 A message displays notifying you of a successful installation and the locations of the `/opt/VRTS/install` files.

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

```
/opt/VRTS/install/logs/installsf-xSY2LZ
```

Note: Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.

- 14 After you complete the installation procedure, review important details about initializing (where required), setting up, and using the Veritas software shipped with Veritas Storage Foundation.
See [“Configuring the Veritas Storage Foundation software”](#) on page 101.
- 15 If you are installing Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle in a cluster configuration, create a new repository database.
See [“Upgrading to the new repository database for DB2 and Oracle”](#) on page 94.

Installing the Veritas Enterprise Administrator client

If you plan to run the VEA client, you must install the `VRTSobgui`, `VRTSat`, `VRTSpx`, and `VRTSicisco` packages on the machine you are planning to use.

By default, the `VRTSobgui`, `VRTSat`, `VRTSpx`, and `VRTSicisco` packages are installed when you install a Veritas Storage Foundation product. You only need to install the packages manually if you are installing on a machine other than the server.

Installing on Solaris

To install the VEA client on a Solaris machine using `pkgadd`

- 1 Log in as superuser (`root`).
- 2 First, check to determine whether the VEA client package is already installed.

```
# pkginfo | grep VRTSobgui
```

This command will return `VRTSobgui` if `VRTSobgui` is already installed. It will return nothing if the package has not been installed.

- 3 To install the VEA client package for Solaris, insert the appropriate media disc into your system’s DVD-ROM or CD-ROM drive.

- 4 Copy the `VRTSobgui.tar.gz` package, and the dependent packages (`VRTSat`, `VRTSspb`, and `VRTSicsco`) to the current working directory on your system.

```
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSobgui.tar.gz .
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSat.tar.gz .
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSicsco.tar.gz .
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSspb.tar.gz .
```
- 5 Decompress the packages, and then extract the contents.

```
# /cdrom/cdrom0/gnu/gunzip *.tar.gz
# tar xvf VRTSobgui.tar
# tar xvf VRTSat.tar
# tar xvf VRTSicsco.tar
# tar xvf VRTSspb.tar
```
- 6 Use the `pkgadd` command to install the packages, in the following order, to allow for dependencies. Answer any questions, as the installation proceeds.

```
# pkgadd -d . VRTSicsco
# pkgadd -d . VRTSat
# pkgadd -d . VRTSspb
# pkgadd -d . VRTSobgui
```
- 7 The VEA client package for Solaris, and dependent packages, are installed.

Installing on Windows

The VEA client runs on Windows NT, Windows XP, Windows 2003, Windows 2000, Windows ME, Windows 98, and Windows 95 machines. If you plan to run VEA from a Windows machine, install the optional Windows package after you have installed the VEA server on a Solaris machine.

Before you install VEA on a Windows machine, you must uninstall any existing Veritas Volume Manager Storage Administrator™ (VMSA) packages and remove the old `setup.exe` from that machine. Only one VEA package can be installed on a Windows machine at any given time.

Note: 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.

Verifying the Veritas Storage Foundation installation

Verify that the Veritas Storage Foundation product was installed.

To verify that the packages were installed

- ◆ Use the `pkginfo` command as follows:

```
# pkginfo -l VRTSvlic package_name package_name ...
```

Use the following sections to further verify the product installation.

Checking Volume Manager processes

To confirm that key Volume Manager processes are running

- ◆ Type the following command:

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

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

Checking VxFS installation

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

Verifying kernel installation

To ensure that the file system driver is loaded, enter:

```
# modinfo | grep vxfs
```

The `modinfo` command displays information about all modules loaded on the system. If the `vxfs` module is loaded, you will see an entry corresponding to `vxfs`. If not, follow the instructions in [“Loading and unloading the file system module”](#) on page 47 to complete the process.

Verifying command installation

The Veritas File System commands are installed in four directories:

Table 2-5 VxFS command locations

Location	Contents
<code>/etc/fs/vxfs</code>	Contains the Veritas <code>mount</code> command and QuickLog commands required to mount file systems.
<code>/usr/lib/fs/vxfs/bin</code>	Contains the VxFS type-specific switch-out commands.
<code>/opt/VRTSvxfs/sbin</code>	Contains the Veritas-specific commands.
<code>/opt/VRTS/bin</code>	Contains symbolic links to all Veritas-specific commands installed in the directories listed above.

Determine whether these subdirectories are present:

```
# ls /etc/fs/vxfs
# ls /usr/lib/fs/vxfs/bin
# ls /opt/VRTSvxfs/sbin
# ls /opt/VRTS/bin
```

Make sure you have adjusted your environment variables accordingly.

See [“Environment variables”](#) on page 19.

Checking cluster operation

You need to check cluster operation only if you installed and configured an HA version of the Storage Foundation software.

To verify that the cluster is operating

- ◆ Type the following command on any node:

```
# hastatus -summary

-- SYSTEM STATE
-- System          State          Frozen

A host1           RUNNING       0
A host2           RUNNING       0

-- GROUP STATE
-- Group           System        Probed  AutoDisabled  State

B ClusterService host1    Y       N              ONLINE
```

```
B ClusterService host2 Y N OFFLINE
```

Identify the system state of each node in the output of this command. If the value is `RUNNING` for all the nodes, VCS is successfully installed and running.

Refer to the `hastatus(1M)` manual page and the *Veritas Cluster Server User's Guide* for more information on system states and state transitions.

About Low Latency Transport

The file `llthosts(4)` is a database containing one entry per node that links the Low Latency Transport (LLT) system ID (in the first column) with the LLT host name. This file is identical on each cluster node.

Based on the sample installation, the file `/etc/llthosts` contains the entries:

```
0 host1
1 host2
```

The file `llttab(1M)` contains information derived from the installation and used by the utility `lltconfig(1M)`. After installation, this file lists the network links that correspond to the specific node.

The first line identifies the node. The second line identifies the cluster, based on the cluster ID entered during installation. The next two lines, beginning with the `link` command, identify the two network cards used by the LLT protocol.

See the `llttab(4)` manual page for details on how to modify the LLT configuration. The manual page describes ordering the directives in the `llttab` file.

Verifying LLT

Use the `lltstat` command to verify that LLT links are active for each system in the cluster.

To verify that links are active for LLT

- ◆ Use the `lltstat -n` as follows:

```
# lltstat -n
```

With LLT configured correctly, the output of `lltstat -n` shows all of the nodes in the cluster and two links for each node. If the output shows otherwise, type `lltstat -nvv | more` on any node to view additional information about LLT.

To obtain LLT port information

- ◆ Use the `lltstat -p` command as follows:

```
# lltstat -p
```

About Group Membership and Atomic Broadcast

After installation, the file `/etc/gabtab` contains a `gabconfig(1M)` command that configures the Group Membership and Atomic Broadcast (GAB) driver.

The file `/etc/gabtab` contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

where the `-c` option configures the driver and `-nN` specifies the cluster will not be formed until at least `N` nodes are ready. The variable `N` represents the number of cluster nodes.

Verifying GAB

To verify that GAB is operating

- ◆ Type the following command on each system:

```
# /sbin/gabconfig -a
```

If GAB is operating, the following GAB port membership information is returned:

```
GAB Port Memberships
=====
Port a gen a36e0003 membership 01
Port h gen fd570002 membership 01
```

Port `a` indicates that GAB is communicating, `gen a36e0003` is a randomly generated number, and `membership 01` indicates that nodes 0 and 1 are connected.

Port `h` indicates that VCS is started, `gen fd570002` is a randomly generated number, and `membership 01` indicates that nodes 0 and 1 are both running VCS.

If GAB is not operating, no GAB port membership information is returned:

```
GAB Port Memberships
=====
```

If only one network is connected, the following GAB port membership information is returned:

```
GAB Port Memberships
=====
Port a gen a36e0003 membership 01
Port a gen a36e0003 jeopardy 1
Port h gen fd570002 membership 01
Port h gen fd570002 jeopardy 1
```

For more information on GAB, including descriptions of ports, refer to the *Veritas Cluster Server User's Guide*.

Loading and unloading the file system module

On Solaris 8, 9, and 10, 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 disk layout. In some instances, you may want to load the file system module manually. To do this, first load `vxfs`, then `vxportal`. `vxportal` is a pseudo device driver that enables VxFS commands to issue ioctls to the VxFS modules even when there are no file systems mounted on the system.

```
# modload /kernel/fs/vxfs
# modload /kernel/drv/vxportal
```

If you have a license for the Veritas Quick I/O feature, you can load its kernel modules:

```
# modload /usr/kernel/drv/sparcv9/fdd
```

To determine if the modules successfully loaded, enter:

```
# modinfo | grep vxportal
# modinfo | grep vxfs
```

The above commands provide information about the modules. The first field in the output is the module ID.

You can unload the module by entering:

```
# modunload -i portal_module_id
# modunload -i vxfs_module_id
```

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

```
# df -F vxfs
```

Installing the Veritas software with JumpStart

If you plan to install Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, or Veritas Storage Foundation for Sybase using JumpStart, unzip the packages and then follow the documentation that came with your OS.

Unzipping the packages

Before you can install the packages, you must unzip them.

To unzip the packages

- 1 Log in as superuser (`root`).
- 2 Create a directory for installation.

```
# mkdir /parent_directory/install
```

- 3 Insert the software disc with the Veritas Storage Foundation software into a drive connected to the system. The Solaris volume-management software automatically mounts the disc as `/cdrom/cdrom0`. Type the command:

```
# cd /cdrom/cdrom0
```

- 4 Copy the compressed package files from the software disc to the temporary directory.

For Veritas Storage Foundation:

```
# cp -r storage_foundation/pkgs/* /parent_directory/install
```

For Veritas Storage Foundation for DB2:

```
# cp -r storage_foundation_for_db2/pkgs/* \  
/parent_directory/install
```

For Veritas Storage Foundation for Oracle:

```
# cp -r storage_foundation_for_oracle/pkgs/* \  
/parent_directory/install
```

For Veritas Storage Foundation for Sybase:

```
# cp -r storage_foundation_for_sybase/pkgs/* \  
/parent_directory/install
```

For Veritas Volume Manager:

```
# cp -r volume_manager/pkgs/* /parent_directory/install
```

For Veritas File System:

```
# cp -r file_system/pkgs/* /parent_directory/install
```

- 5 Go to the temporary directory and unzip the compressed package files.

Note: If your system does not have the `gunzip` utility, copy it from the disc:

```
# cp /cdrom/cdrom0/gnu/gunzip /parent_directory/install
```

```
# cd /parent_directory/install  
# gunzip VRTS*.gz
```

- 6 Decompress and extract each package.

```
# tar xf package_name.tar  
# tar xf package_name.tar  
# tar xf package_name.tar  
.  
.
```

- 7 List the files in the temporary directory.

```
# ls /parent_directory/install
```

Installation order

You must install the packages in the correct order. For example, some packages must be installed before other packages because of various product dependencies. This section provides the package lists for each product in the correct order.

Note: For details on all language packages, see “[Veritas Storage Foundation install packages](#)” on page 177.

Veritas Storage Foundation

To get the package installation order for Veritas Storage Foundation, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation

- 1 Mount the software disc.
See “[Mounting the software disc](#)” on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsf -requiredpkgs
```

or

```
# ./installsf -installpkgs
```

Veritas Storage Foundation for DB2

To get the package installation order for Veritas Storage Foundation for DB2, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation for DB2

- 1 Mount the software disc.
See “[Mounting the software disc](#)” on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```

- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsfdb2 -requiredpkgs  
or  
# ./installsfdb2 -installpkgs
```

Veritas Storage Foundation for Oracle

To get the package installation order for Veritas Storage Foundation for Oracle, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation for Oracle

- 1 Mount the software disc.
See [“Mounting the software disc”](#) on page 32.

- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```

- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsfora -requiredpkgs  
or  
# ./installsfora -installpkgs
```

Veritas Storage Foundation for Sybase

To get the package installation order for Veritas Storage Foundation for Sybase, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation for Sybase

- 1 Mount the software disc.
See [“Mounting the software disc”](#) on page 32.

- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```

- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsfsyb -requiredpkgs  
or  
# ./installsfsyb -installpkgs
```

Volume Manager

To get the package installation order for Veritas Volume Manager, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Volume Manager

- 1 Mount the software disc.
See [“Mounting the software disc”](#) on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installvm -requiredpkgs
```

or

```
# ./installvm -installpkgs
```

File System

To get the package installation order for Veritas File System, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas File System

- 1 Mount the software disc.
See [“Mounting the software disc”](#) on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installfs -requiredpkgs
```

or

```
# ./installfs -installpkgs
```

Summary of JumpStart installation tasks

Caution: For detailed instructions, follow the JumpStart documentation that came with your OS. These steps are provided as a summary only.

Installing the Veritas packages on a JumpStart server consists of the following tasks:

- 1 Unzip the packages.
- 2 Copy the packages to a JumpStart server under a shared directory on the network.
- 3 Issue a command similar to the following on a JumpStart client:

```
# pkgadd -R /pathname -a admin_file -d mount_point  
package_name1 package_name2 ...
```

where *mount_point* is the mount point of the network shared directory to which you copied the packages in step 2.
For the list of packages, see “[Installation order](#)” on page 49.

Note: This command can be added as part of the `finish_install` (post-install) script of the JumpStart client.

- 4 Reboot the system using a command such as `shutdown`.
- 5 Run the appropriate command from the disc directory to configure the Veritas software.

For	Use
Veritas Storage Foundation	<code>./installsf -configure</code>
Veritas Storage Foundation for DB2	<code>./installsfdb2 -configure</code>
Veritas Storage Foundation for Oracle	<code>./installsfora -configure</code>
Veritas Storage Foundation for Sybase	<code>./installsfsyb -configure</code>
Veritas Volume Manager	<code>./installvm -configure</code>
Veritas File System	<code>./installfs -configure</code>

Installing language packages

If you are installing a Veritas Storage Foundation product in a language other than English, you must install the required language packages *after* installing the English packages. If you are planning to use the GUI, you must install the language package for the VEA client. The procedures for installing the Simplified Chinese and French language packages differ from those for the Japanese language packages.

See “[Installing the French and Simplified Chinese language packages](#)” on page 53.

See “[Installing the Japanese language packages](#)” on page 55.

Installing the French and Simplified Chinese language packages

To install the French or Simplified Chinese language packages on the server

- 1 Make sure the VEA Service is not running.

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : RUNNING
```

- 2 If the VEA Service is running, stop it by using the `vxsvcctl stop` command.

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 3 Insert the “Language” disc into the DVD-ROM or CD-ROM drive. If you are using Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

- 4 For the French language packages, copy the `fr/volume_manager/pkgs` directory to a temporary directory on your system, such as `/tmp/pkgs`.

```
# cp -r /cdrom/cdrom0/fr/volume_manager/pkgs /tmp/pkgs
```

For the Simplified Chinese language packages, copy the `zh/volume_manager/pkgs` directory to a temporary directory on your system, such as `/tmp/pkgs`.

```
# cp -r /cdrom/cdrom0/zh/volume_manager/pkgs /tmp/pkgs
```

- 5 Decompress the packages, and extract the contents.

```
# cd /tmp/pkgs
# /cdrom/cdrom0/gnu/gunzip *.gz
# tar xvf *.tar
```

- 6 Use the `pkgadd` command to install the packages.

To install the French language packages, use the following command:

```
# pkgadd -d . VRTSfrvmc VRTSfrvmd VRTSfrvmm VRTSmualc \
VRTSmuc33 VRTSmuddl VRTSmufsp VRTSmulic VRTSmuob \
VRTSmuobg VRTSmuvmp VRTSmuvmw
```

To install the Simplified Chinese language packages, use the following command:

```
# pkgadd -d . VRTSzhvmc VRTSzhvmd VRTSzhvmm VRTSmualc \  
VRTSmuc33 VRTSmuddl VRTSmufsp VRTSmulic VRTSmuob \  
VRTSmuobg VRTSmuvmp VRTSmuvmw
```

- 7 After installing the language packages, restart the VEA Service.

```
# /opt/VRTS/bin/vxsvcctl start
```

To install the French or Simplified Chinese language version of the VEA client package on a Solaris machine other than the server

- 1 Make sure the VEA Service is not running.

```
# /opt/VRTS/bin/vxsvcctl status  
Current state of server : RUNNING
```

- 2 If the VEA Service is running, stop it by using the `vxsvcctl stop` command.

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 3 Insert the “Language” disc into the DVD-ROM drive. If you are using Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

- 4 For the French language VEA client package, copy the `fr/volume_manager/pkgs/VRTSmuobg.tar.gz` package to a temporary directory on your system, such as `/tmp/pkgs`.

```
# cp /cdrom/cdrom0/fr/volume_manager/pkgs/VRTSmuobg.tar.gz \  
/tmp/pkgs
```

For the Simplified Chinese language VEA client package, copy the `zh/volume_manager/pkgs/VRTSmuobg.tar.gz` package to a temporary directory on your system, such as `/tmp/pkgs`.

```
# cp /cdrom/cdrom0/zh/volume_manager/pkgs/VRTSmuobg.tar.gz \  
/tmp/pkgs
```

- 5 Decompress the package, and extract the contents.

```
# /cdrom/cdrom0/gnu/gunzip VRTSmuobg.tar.gz  
# tar xvf VRTSmuobg.tar
```

- 6 Use the `pkgadd` command to install the package.

```
# pkgadd -d . VRTSmuobg
```

To install the French or Simplified Chinese language version of the VEA client package on a Windows machine

- 1 Insert the “Language” disc into the DVD drive.
- 2 Go to the directory containing the client language package, `D:\language\windows`, where `D` is the DVD drive.
- 3 Double-click on the `VRTSmuobg.msi` package to install it.
- 4 Follow any instructions during installation.

Installing the Japanese language packages

To install the Japanese language packages on the server

- 1 Make sure the VEA Service is not running.

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : RUNNING
```
- 2 If the VEA Service is running, stop it by using the `vxsvcctl stop` command.

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 3 Insert the “Language” disc into the DVD drive. If you are using Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.
- 4 Install the language packages using the `install_lp` command.

```
# cd /cdrom/cdrom0
# ./install_lp
```
- 5 After installing the language packages, restart the VEA Service.

```
# /opt/VRTS/bin/vxsvcctl start
```

To install the Japanese language version of the VEA client package on a Solaris machine other than the server

- 1 Make sure the VEA Service is not running.

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : RUNNING
```
- 2 If the VEA Service is running, stop it by using the `vxsvcctl stop` command.

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 3 Insert the “Language” disc into the DVD-ROM drive. If you are using Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

- 4 Copy the `ja/volume_manager/pkgs/VRTSmuobg.tar.gz` package to a temporary directory on your system, such as `/tmp/pkgs`.

```
# cp /cdrom/cdrom0/ja/volume_manager/pkgs/VRTSmuobg.tar.gz \  
/tmp/pkgs
```
- 5 Decompress the package, and extract the contents.

```
# /cdrom/cdrom0/gnu/gunzip VRTSmuobg.tar.gz  
# tar xvf VRTSmuobg.tar
```
- 6 Use the `pkgadd` command to install the package.

```
# pkgadd -d . VRTSmuobg
```

To install the Japanese language version of the VEA client package on a Windows machine

- 1 Insert the “Language” disc into the DVD drive.
- 2 Go to the directory containing the client language package, `D:\language\windows`, where `D` is the DVD drive.
- 3 Double-click on the `VRTSmuobg.msi` package to install it.
- 4 Follow any instructions during installation.

Installing Veritas Storage Foundation Basic

Topics covered in this chapter include:

- “About Veritas Storage Foundation Basic”
- “Installing Veritas Storage Foundation Basic”
- “Upgrading Veritas Storage Foundation Basic”
- “Uninstalling Veritas Storage Foundation Basic”

This chapter describes how to install, upgrade, and uninstall the Veritas Storage Foundation Basic software.

Storage Foundation Basic is a special offering that is available on a separate Storage Foundation Basic disc or downloadable from the Symantec website. Storage Foundation Basic is not part of the Storage Foundation and High Availability Solutions product suite. For complete information on ordering this product, licensing, and technical support, visit:

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

About Veritas Storage Foundation Basic

The Storage Foundation Basic software contains the same features as Storage Foundation Standard software, but you will receive messages on the system console if you exceed the following soft limitations of the license:

- Maximum of 4 VxVM volumes
- Maximum of 4 VxFS file systems
- Maximum server capacity of 2 CPU sockets

For a product overview on Storage Foundation Basic, see the *Veritas Storage Foundation and High Availability Solutions Getting Started Guide*.

Storage Foundation Basic technical support

Technical support is self-service only, available from the Veritas Support website. You can purchase additional support corresponding to the terms of the Storage Foundation Basic license. To access the self-service knowledge base, go to:

<http://support.veritas.com>

When contacting Support with questions relating to Storage Foundation Basic, be prepared to provide your product license key. You can determine your currently installed license by running the Veritas license report utility as shown in the following example:

```
# /opt/VRTS/bin/vxlicrep
....
License Key      = XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXX
Product Name    = Veritas Storage Foundation Basic
```

The Storage Foundation Basic license key is installed automatically by the Veritas product installer.

Storage Foundation Basic system requirements

Dependencies

Veritas Storage Foundation Basic can only be installed on a system running Solaris 10 (64-bit). Installing this product on any other Solaris release will fail. If necessary, upgrade Solaris before you install the Veritas products.

Disk space

The product installation script automatically checks that sufficient disk space is available. However, before installing the Veritas Storage Foundation Basic software, you may wish to confirm that there is sufficient disk space in the file systems on the target systems.

The following table shows the approximate disk space used by all the Storage Foundation Basic software packages:

/ (excluding /opt, /usr and /var)	/opt	/usr	/var	Total disk space required
75MB	642MB	125MB	3MB	845MB

The following table describes each package:

Package	Contents
Volume Manager packages	
VRTSalloc	Veritas Volume Manager: Veritas Intelligent Storage Provisioning
VRTSddlpr	Veritas Device Discovery Layer Services Provider
VRTSvdiid	Veritas Device Identification API
VRTSvmdoc	Veritas Volume Manager Documentation (optional)
VRTSvmman	Veritas Volume Manager -Manual Pages (optional)
VRTSvmpro	Veritas Volume Manager Management Services Provider
VRTSvxvm	Veritas Volume Manager, Binaries
File System packages	

Package	Contents
VRTSfsdoc	Veritas File System Documentation (optional)
VRTSfsman	Veritas File System - Manual Pages (optional)
VRTSfsmnd	Veritas File System Software Developer Kit Manual Pages (optional)
VRTSfspro	Veritas File System Management Services Provider
VRTSfssdk	Veritas File System Software Developer Kit
VRTSvxfs	Veritas File System, Binaries
Veritas Enterprise Administrator packages	
VRTSaa	Veritas Enterprise Administrator Action Agent
VRTSccg	Veritas Enterprise Administrator Central Control Grid
VRTSob	Veritas Enterprise Administrator Service
VRTSobc33	Veritas Enterprise Administrator Core Service
VRTSobgui	Veritas Enterprise Administrator
Veritas FlashSnap Agent for Symmetrix packages	
VRTSfas	Veritas FlashSnap Agent for Symmetrix
VRTSfasdc	Veritas FlashSnap Agent for Symmetrix Documentation
Miscellaneous packages	
SYMClma	Symantec License Inventory Agent
VRTSat	Symantec Product Authentication Service
VRTSdcli	Veritas Distributed Command Line Agent
VRTSdsa	Veritas Datacenter Storage Agent
VRTSgapms	Veritas Generic Array Plugin
VRTSicsco	Symantec Common Infrastructure
VRTSmapro	Veritas Storage Foundation GUI for Mapping

Package	Contents
VRTSmh	Veritas Centralized Management for Storage Foundation Managed Host
VRTSspb	Symantec Private Branch Exchange
VRTSperl	Perl 5.8.0 for Veritas
VRTSsmf	Symantec Service Management Framework
VRTSspt	Veritas Software Support Tools
VRTSvail	Veritas Array Integration Layer
VRTSvlic	Veritas License Utilities
VRTSvsvc	Veritas Volume Server and Client Provider
VRTSvxmsa	Veritas Mapping Service, Application Libraries
windows/vrtsobgui.msi	Veritas Enterprise Administrator for Windows

Installing Veritas Storage Foundation Basic

This section describes how to install Veritas Storage Foundation Basic.

Note: The Storage Foundation Basic license key is automatically installed by the Veritas product installer.

SF Basic systems can be administered with the Storage Foundation Management Server.

See [“Preinstallation Instructions”](#) on page 11.

To install Veritas Storage Foundation Basic

- 1 If you are installing from the Veritas Storage Foundation Basic disc, insert the disc into your system's DVD-ROM drive.

If Solaris volume management software is running on your system, the software disc is automatically mounted with a path such as `/cdrom/cdrom0`.

If Solaris volume management software is not available, mount the disc manually, as shown here:

```
# mount -F cdfs -o ro /dev/dsk/c0t6d0s2 /mount_point
```

where `c0t6d0s2` is the default address for the DVD drive.

- 2 Move to the top-level directory on the mounted disc, or in the extracted download, for example:

```
# cd /cdrom/cdrom0
```

- 3 From this directory, type the following command if you are installing on the local system only, or if you are using the secure shell (SSH) utilities to install on remote systems:

```
# ./installer
```

If you are using the remote shell (RSH) utilities to install on remote systems, additionally specify the `-rsh` option:

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

The sample installation assumes that SSH is being used.

- 4 At the product status page, enter **1** for the product installer and press **Return**. The product installation begins.

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

- 5 Enter the system names on which the software is to be installed. The script performs an initial check on each system, installs the licensing packages, and displays the packages that will be installed. You are prompted to press **Return** to page through the list of packages.

```
It is optional to configure SF Basic now. If you choose to  
configure SF Basic later, you can either do so manually or  
run the installsf -configure command.
```

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

- 6 Press **Return** to continue installing the packages and to configure the software.

```
Installation completed successfully on all systems
```

```
The enclosure-based naming scheme is a feature of Volume  
Manager. It allows one to reference disks using a symbolic  
name that is more meaningful than the operating system's  
normal device access name. This symbolic name is typically  
derived from the array name.
```

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

- 7 Enter **y** or enclosure-based naming scheme, or **n** to use.

```
Volume Manager default disk group configuration:
```

```
Many Volume Manager commands affect the contents or  
configuration of a disk group. Such commands require that  
the user specify a disk group. This is accomplished by using  
the -g option of a command or setting the VXVM_DEFAULTDG  
environment variable. An alternative to these two methods is  
to configure the default disk group of a system.
```

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

- 8** Press **Return** to set up the default disk group for each system. You are prompted to specify one disk group name for all eligible systems.

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

- 9** Press **Return** to specify a disk group name for all eligible systems.

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

- 10** Enter the name of the default disk group for each system.

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

- 11** Press **Return** to confirm the name of the default disk group.

- 12** If you are prompted to configure Veritas FlashSnap™ Agent for Symmetrix (VxFAS), follow the instructions to configure or unconfigure VxFAS, or skip this step.

```
Verify the install systems Fully Qualified Hostnames.
```

```
Querying fully qualified domain name of host "host1" ... ok
```

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

- 13** Press **Return** to confirm the fully qualified host name of each system.

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

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

If you select to enable centralized management, you are asked a series of questions relating to the configuration of the Storage Foundation Management Server.

- 15** Finally, you are asked whether you want to start the Storage Foundation processes on the target systems.

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

Press **Return** to start the Storage Foundation Basic processes.

- 16** A message displays notifying you of a successful installation and the locations of the /opt/VRTS/install files.

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

```
/opt/VRTS/install/logs/installsf-xSY2LZ
```

Note: Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes. See “[Storage Foundation Basic technical support](#)” on page 58.

Upgrading Veritas Storage Foundation Basic

This section describes how to upgrade Veritas Storage Foundation Basic to a full version of Veritas Storage Foundation.

Note: If you upgrade Storage Foundation Basic to any other product, contact Veritas Sales for product information.

Planning the upgrade

Complete the following tasks in advance of upgrading:

- Ensure that you have created a valid backup.
- Review the *Veritas Storage Foundation Release Notes* for any late-breaking information on upgrading your system.
- Be sure that the administrator doing the upgrade has root access and a working knowledge of Solaris operating system administration.
- Schedule sufficient outage time for the upgrade.

To determine which version of Storage Foundation is installed on your system, run the `vxlicrep` command.

Upgrade paths

The upgrade paths for Veritas Storage Foundation Basic are:

From	Upgrade to	Tasks
Storage Foundation Basic	Storage Foundation Standard	<ul style="list-style-type: none">■ Install the license key by running <code>vxlicinst</code> command.■ Run <code>vxdtl license init</code> to reread the Storage Foundation Standard license key.

From	Upgrade to	Tasks
Storage Foundation Basic	Storage Foundation Enterprise	<ul style="list-style-type: none"> ■ Install the licence key by running <code>vxlicinst</code> command. ■ Run <code>vxdtl license init</code> to reread the Storage Foundation Enterprise license key.
Storage Foundation Basic	Storage Foundation Enterprise with licensed features: Veritas Volume Replicator	<ul style="list-style-type: none"> ■ Run the Storage Foundation product <code>installer</code> to install the Storage Foundation Enterprise and Veritas Volume Replicator key. See “Installing Veritas Storage Foundation” on page 19. See “Installing a Veritas Storage Foundation product” on page 29. ■ Run <code>vxdtl license init</code> to reread the Veritas Volume Replicator license key.
Storage Foundation Basic	Any Storage Foundation product	<ul style="list-style-type: none"> ■ Uninstall Storage Foundation Basic. See “Uninstalling Veritas Storage Foundation Basic” on page 65. ■ Install any Storage Foundation product, see the <i>Veritas Storage Foundation and High Availability Getting Started Guide</i>.

Uninstalling Veritas Storage Foundation Basic

This section describes how to uninstall Veritas Storage Foundation Basic.

To uninstall Veritas Storage Foundation Basic

- 1 Move all of your data off any VxVM volumes or VxFS file systems that you have created. Unmount all VxFS file systems, and stop any applications that may be accessing VxVM volumes.
- 2 Change directory to `/opt/VRTS/install`, and run the `uninstallsf` program.

```
# cd /opt/VRTS/install
# ./uninstallsf
```

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

- 3 Enter the system names on which the software is installed. You are prompted to press **Return** to continue.

Press [Return] to continue:

- 4 Press **Return** to continue. You are prompted to uninstall the Storage Foundation Basic packages.

```
uninstallsf is now ready to uninstall SF.
All SF processes that are currently running must be stopped.
```

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

- 5 Enter **y** to uninstall the Storage Foundation Basic packages. You are prompted to press **Return** to continue.

```
Storage Foundation Basic package uninstall completed
successfully.
```

Press [Return] to continue:

Press **Return** to continue after the Storage Foundation Basic package uninstall completed successfully.

```
Uninstall completed successfully on all systems
Uninstallation log files, summary file, and response file
are saved at:
```

```
/opt/VRTS/install/logs/uninstallsf-7TdmtZ
```

```
NOTICE: As part of the uninstallation process on Solaris,
the current configuration of VxVM was saved in the directory
/VXVM-CFG-BAK. This configuration may be used in a future
installation of VxVM. If you do not plan to reuse it, you
may manually remove this subdirectory.
```

It is suggested that you archive the saved files for review in the event that they are needed at a future date.

Upgrading a Veritas Storage Foundation product

Topics covered in this chapter include:

- [“Preparing to upgrade Veritas Volume Manager”](#) on page 68
- [“Retaining older Veritas Volume Manager releases”](#) on page 80
- [“Upgrading Veritas Storage Foundation”](#) on page 83
- [“Upgrading to the new repository database for DB2 and Oracle”](#) on page 94
- [“Verifying the Veritas Storage Foundation upgrade”](#) on page 95
- [“Upgrading from SUNWvxcvm”](#) on page 97
- [“Upgrading language packages”](#) on page 99

If you are running an earlier release of Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, Veritas Storage Foundation for Sybase, Veritas Volume Manager, or Veritas File System, you can upgrade your product using the procedures described in this chapter.

Note: The information in [“Upgrading Veritas Storage Foundation”](#) on page 83 pertains to Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation for Sybase.

Preparing to upgrade Veritas Volume Manager

This section lists required pre-upgrade tasks for VxVM.

Upgrading VxVM on a system with Alternate Pathing

Note: The information in this section only applies if you are running Solaris 8.

VxVM 5.0 does not allow Dynamic Multipathing (DMP) to co-exist with Sun's Alternate Pathing (AP) software. The AP driver must be disabled before installing or upgrading VxVM.

Pre-upgrade requirements

You can prepare for VxVM 5.0 upgrade requirements in two phases:

- Tasks that can be done well ahead of starting the upgrade (see “[Planning the upgrade](#)” on page 68).
- Tasks that need to be done just before starting the upgrade (see “[Preparing for the upgrade](#)” on page 69).

Planning the upgrade

Complete the following tasks in advance of upgrading:

- Be sure that the administrator doing the upgrade has root access and a working knowledge of Solaris operating system administration.
- Check that all terminal emulation issues are worked out. The terminal selected should be fully functional during OpenBoot prompts and single-user and multi-user run levels.
- Check the latest *Storage Foundation Release Notes* to verify that the system has all required Solaris patches.
- Schedule sufficient outage time for the upgrade and possibly several system reboots.
- If using EMC PowerPath, ensure that you are using at least mandatory patch level 2.0.3. See <http://support.veritas.com/docs/234734> for more information. The patch level makes changes to `/etc/system` that prevent panics and failure of `vxconfigd`. Upgrading PowerPath may require a system reboot.

- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. (This may not be practical, but if done, offers a failback point.)
- Make sure that all applicable license information is available and copy the contents of `/etc/vx/licenses` to a safe directory. If it exists, you should also copy the contents of `/etc/vx/elm` to a safe directory.

Preparing for the upgrade

To prepare for the upgrade

- 1 Ensure that you have created a valid backup.
- 2 Review the *Veritas Storage Foundation Release Notes*.
- 3 Ensure that you have enough file system space to upgrade Veritas Volume Manager. Also, identify where you will be copying the distribution and patch tar files. The usual place is `/patches/Veritas` when the root file system has enough space or `/var/tmp/patches` if the `/var` file system has enough space.

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

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

- 4 If you are installing using `pkgadd` instead of the installation script, untar Veritas Volume Manager and patch files (preferably into separate sub-directories). Within the untarred Volume Manager files is a script directory. Note the location of the script directory.
- 5 For any startup scripts in `/etc/rcS.d`, you should comment out any application commands or processes that are known to hang if their file systems are not present.
- 6 Make sure that all users are logged off and that all major user applications are properly shut down.
- 7 Copy `vfstab` to `vfstab.orig`:

```
# cp /etc/vfstab /etc/vfstab.orig
# cp /etc/vfstab /VXVM5.0-UPGRADE-CHECK/vfstab.orig
```

- 8 Veritas recommends that all file systems not on the `root` disk (therefore not required for booting) be unmounted, their entries commented out in `/etc/vfstab`, the associated volumes stopped, and the associated disk groups deported. Any file systems that the Solaris operating system or Veritas assumes should be in `rootdg` but are not, *must* be unmounted and the associated entry in `/etc/vfstab` commented out.
- 9 Any swap partitions not in `rootdg` must be commented out of `/etc/vfstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/vfstab` and not mounted during the upgrade. Active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.

Determining if the root disk is encapsulated

Before you upgrade, you need to determine if the root disk is encapsulated by running the following command:

```
# mount | grep "/ on"
```

If the output from this command includes a path name that contains `vx` and `rootvol` as in `/dev/vx/dsk/rootvol`, then the root disk is encapsulated.

After you complete the upgrade procedure, read “[Configuring the Veritas Storage Foundation software](#)” on page 101 for important details about initializing (where required), setting up, and using the Veritas software shipped with Veritas Storage Foundation for Oracle.

Adding Sun T3+ and T3 arrays as JBODS of type A/P

This release does not include the `libvxpurple.so` array support library (ASL) to support Sun StorEdge T3 and T3+ arrays. Any existing version of the `libvxpurple.so` ASL is removed when VxVM is upgraded to 5.0. Any T3 and T3+ arrays must be configured in autotrespass mode, and treated as JBODs of type A/P.

If an array is of type A/A-A, A/P or A/PF, and a suitable ASL is not available, the array must be claimed as an JBOD of type A/P. This is to prevent path delays and I/O failures arising. As JBODs are assumed to be type A/A by default, and neither T3 nor T3+ arrays are of this type, you must create appropriate JBOD entries for such arrays.

To configure an A/A-A, A/P or AP/F array as a JBOD

- 1 Stop all applications, such as databases, from accessing VxVM volumes that are configured on the array, and unmount all VxFS file systems and checkpoints that are configured on the array.
- 2 Configure the T3 or T3+ array in autotrespass mode.

- 3 Add the array as a JBOD of type A/P:

```
# vxddladm addjbod vid=SUN pid=T300 policy=ap
```
- 4 If you have not already done so, upgrade the Storage Foundation or VxVM software to 5.0. Device discovery will be performed during the upgrade, and the array will be claimed as an A/P JBOD.
 If you have already upgraded your system to 5.0, run the following command to perform device discovery:

```
# vxdctl enable
```
- 5 Verify that the array has been added with the policy set to APdisk:

```
# vxddladm listjbod
```

VID	PID	Opcode	Page Code	Page Offset	SNO	length	Policy
SUN	T300	18	-1	36	12		APdisk
- 6 Check that the correct devices are listed for the array:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
APdisk_0	auto:cdsdisk	-	-	online invalid
APdisk_1	auto:cdsdisk	-	-	online invalid
APdisk_2	auto:cdsdisk	-	-	online invalid
...				

Unsuppressing DMP for EMC PowerPath disks

Note: This section is only applicable if you are upgrading a system that includes EMC PowerPath disks.

If you are upgrading a system from VxVM 4.0 to VxVM 5.0, which has PowerPath installed, and the Cx600 ASL and its associated Cx600 APM are also installed, you must uninstall both the Cx600 ASL and APM, otherwise the Cx600 will claim the disks and the PowerPath disks will not be identified.

If you are upgrading a system from VxVM 4.0 to VxVM 5.0, which does not have PowerPath installed, but the Cx600 ASL and its APM are both installed, then the Cx600 ASL and its associated APM should not be uninstalled.

In releases of VxVM before 4.1, a combination of DMP subpaths and the controllers of DMP subpaths were usually suppressed to prevent interference between DMP and the EMC PowerPath multipathing driver. Suppression has the effect of hiding these subpaths and their controllers from DMP, and as a result the disks on these subpaths and controllers cannot be seen by VxVM.

VxVM 4.1 and later releases have the ability to discover EMCpower disks, and configure them as autodiscovered disks that DMP recognizes are under the

control of a separate multipathing driver. This has the benefit of allowing such disks to be reconfigured in cluster-shareable disk groups. Before upgrading to VxVM 5.0, you must remove the suppression of the subpaths and controllers so that DMP can determine the association between EMCpower metadisks and `c#t#d#` disk devices.

There are several scenarios where you need to unsuppress DMP subpaths and controllers:

- [“Converting a foreign disk to auto:simple”](#) on page 72.
- [“Converting a defined disk to auto:simple”](#) on page 74
- [“Converting a powervxvm disk to auto:simple”](#) on page 76

Because emcpower disks are auto-discovered, the `powervxvm` script should be disabled and removed from the startup script. To remove the `powervxvm` script, use the command:

```
# powervxvm remove
```

Converting a foreign disk to auto:simple

Release 4.0 of VxVM provided the `vxdldadm addforeign` command to configure foreign disks with default disk offsets for the private and public regions, and to define them as simple disks. A foreign disk must be manually converted to `auto:simple` format before upgrading to VxVM 5.0.

Note: If the foreign disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
# vxdisk list
DEVICE          TYPE          DISK    GROUP    STATUS
c6t0d12s2      auto:sliced   -       -        online
emcpower10c    simple       fdisk   fdg      online
...
```

The `vxprint` command is used to display information about the disk group, `fdg`:

```
# vxprint
Disk group: fdg
TY NAME      ASSOC          KSTATE LENGTH      PLOFFS STATE  TUTILO
PUTILO
dg fdg      fdg            -         -         -      -      -
dm fdisk    emcpower10c   -         17673456 -         -         -      -
...
```

To convert a foreign disk to `auto:simple` format:

- 1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g fdg stopall
# vxdg deport fdg
```

- 2 Use the `vxddladm` command to remove definitions for the foreign devices:

```
# vxddladm rmforeign blockpath=/dev/dsk/emcpower10c \
    charpath=/dev/rdisk/emcpower10c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
DEVICE      TYPE          DISK          GROUP          STATUS
c6t0d12s2  auto:sliced   -             -              online
...
```

- 3 Run the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
# /etc/vx/bin/vxprtvtoc -f /tmp/vtoc /dev/rdisk/emcpower10c
```

- 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/rdisk/emcpower10c
```

```
# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE      TAG      FLAGS      START      SIZE
    0         0x0     0x201     0          0
    1         0x0     0x200     0          0
    2         0x5     0x201     0         17675520

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE      TAG      FLAGS      START      SIZE
    0         0xf     0x201     0         17675520
    1         0x0     0x200     0          0
    2         0x5     0x201     0         17675520
```

```
DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

- 5 Upgrade to VxVM 5.0 using the appropriate upgrade procedure.

- 6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to `auto:simple` format:

```
# vxdisk list
DEVICE      TYPE          DISK          GROUP          STATUS
c6t0d12s2  auto:sliced   -             -              online
emcpower10s2  auto:simple   -             -              online
...
```

To display the physical device that is associated with the metadevice, `emcpower10s2`, enter the following command:

```
# vxddmpadm getsubpaths dmpnodename=emcpower10s2
```

7 Import the disk group and start the volumes:

```
# vxvg import fdg
# vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
# vxdisk list
DEVICE                TYPE          DISK   GROUP   STATUS
c6t0d12s2             auto:sliced  -     -       online
emcpower10s2         auto:simple  fdisk  fdg     online
```

Converting a defined disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (`darec`), and identified as simple disks. If an EMCpower disk is defined with a persistent `darec`, it must be manually converted to `auto:simple` format before upgrading to VxVM 5.0.

Note: If the defined disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
# ls -l /dev/vx/dmp/emcdisk1
lrwxrwxrwx 1 root other 36 Sep 24 17:59 /dev/vx/dmp/emcdisk1->
/dev/dsk/c6t0d11s5
# ls -l /dev/vx/rdmp/emcdisk1
lrwxrwxrwx 1 root other 40 Sep 24 17:59 /dev/vx/rdmp/emcdisk1->
/dev/dsk/c6t0d11s5
```

Here the fifth partition of `c6t0d11s5` is defined as the persistent disk access record `emcdisk1`.

The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
# vxdisk list
DEVICE                TYPE          DISK   GROUP   STATUS
c6t0d12s2             auto:sliced  -     -       online
emcdisk1              simple       fdisk  fdg     online
...
```

The `vxprint` command is used to display information about the disk group, `fdg`:

```
# vxprint
Disk group: fdg
TY NAME    ASSOC    KSTATE    LENGTH    PLOFFS    STATE    TUTILO    PUTILO
dg fdg     fdg      -         -         -         -         -
dm fdisk   emcdisk1 -         17673456 -         -         -
...
```

To convert a disk with a persistent disk access record to auto:simple format

- 1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g fdg stopall
# vxdg deport fdg
```

- 2 Use the `vxdisk rm` command to remove the persistent record definitions:

```
# vxdisk rm emcdisk1
```

If you now run `vxprint`, the EMCpower disk is no longer displayed:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c6t0d12s2	auto:sliced	-	-	online
...				

- 3 Use the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
# /etc/vx/bin/vxprtvtoc -f /tmp/hdisk /dev/rdisk/c6t0d11s2
```

- 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/hdisk /dev/rdisk/c6t0d11s2
```

```
# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
```

# SLICE	TAG	FLAGS	START	SIZE
4	0x0	0x200	0	0
5	0x0	0x200	3591000	2100375
6	0x0	0x200	0	0

```
# THE NEW PARTITIONING WILL BE AS FOLLOWS:
```

# SLICE	TAG	FLAGS	START	SIZE
4	0x0	0x200	0	0
5	0xf	0x200	3591000	2100375
6	0x0	0x200	0	0

```
DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
```

```
WRITING THE NEW VTOC TO THE DISK #
```

- 5 Upgrade to VxVM 5.0 using the appropriate upgrade procedure.

- 6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to auto:simple format:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c6t0d12s2	auto:sliced	-	-	online
emcpower10s2	auto:simple	-	-	online:aliased
...				

To display the physical device that is associated with the metadvice, `emcpower10s2`, enter the following command:

```
# vxddmpadm getsubpaths dmpnodename=emcpower10s2
```

7 Import the disk group and start the volumes:

```
# vxvg import fdg
# vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
# vxdisk list
DEVICE          TYPE          DISK   GROUP  STATUS
c6t0d12s2      auto:sliced   -      -      online
emcpower10s2   auto:simple   fdisk  fdg    online:aliased
```

Note: To allow DMP to receive correct enquiry data, the common Serial Number (C-bit) Symmetrix Director parameter must be set to enabled.

Converting a powervxvm disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (darec) using `powervxvm script`, and identified as simple disks. If an EMCpower disk is used using `powervxvm`, it must be manually converted to `auto:simple` format before upgrading to VxVM 5.0.

Note: If there are any controllers or devices that are suppressed from VxVM as `powervxvm` requirement, then such controllers/disks must be unsuppressed. This is required for Veritas DMP to determine the association between PowerPath metanodes and their subpaths. After the conversion to `auto:simple` is complete, the `powervxvm script` is no longer useful, and should be disabled from startup script.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
# ls -l /dev/vx/rdmp/
crw----- 1 root  root  260, 76 Feb  7 02:36 emcpower0c
#
# vxdisk list
DEVICE          TYPE          DISK   GROUP  STATUS
c6t0d12s2      auto:sliced   -      -      online
emcpower0c     simple        ppdisk01  ppdg   online
#
# vxprint
Disk group: fdg
TY NAME      ASSOC      KSTATE  LENGTH  PLOFFS  STATE  TUTIL0  PUTIL0
dg ppdg      ppdg      -        -        -        -        -        -
dm ppdisk01 emcpower0c -        2094960 -        -        -        -
```

To convert an EMCpower disk (defined using powervxvm) to auto:simple format

- 1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g ppdg stopall
# vxdg deport ppdg
```

- 2 Use the `vxdisk rm` command to remove all emcpower disks from VxVM:

```
# vxdisk rm emcpower0c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c6t0d12s2	auto:sliced	-	-	online

- 3 Use the `vxprtvtoc` command to retrieve the partition table entry for this device:

```
# /etc/vx/bin/vxprtvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
```

- 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
```

```
# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
```

# SLICE	TAG	FLAGS	START	SIZE
0	0x0	0x201	0	0
1	0x0	0x200	0	0
2	0x5	0x201	0	17675520

```
# THE NEW PARTITIONING WILL BE AS FOLLOWS:
```

# SLICE	TAG	FLAGS	START	SIZE
0	0xf	0x201	0	17675520
1	0x0	0x200	0	0
2	0x5	0x201	0	17675520

```
DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

- 5 Upgrade to VxVM 5.0 using the appropriate upgrade procedure.

After upgrading VxVM, use the `vxdisk list` command to validate the conversion to auto:simple format:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c6t0d12s2	auto:sliced	-	-	online
emcpower0s2	auto:simple	-	-	online

- 6 Import the disk group and start the volumes.

```
# vxdg import ppdg
# vxvol -g ppdg startall
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c6t0d12s2	auto:sliced	-	-	online
emcpower0s2	auto:simple	ppdsk01	ppdg	online

VxVM and Solaris release support matrix

[Table 4-1](#) indicates which VxVM releases support which Solaris operating system releases. If you are running a release earlier than VxVM 3.5, you should first upgrade to VxVM 4.0 before upgrading to VxVM 5.0.

Table 4-1 VxVM and Solaris support matrix

	Solaris 2.6	Solaris 7	Solaris 8	Solaris 9	Solaris 10
VxVM 5.0			Supported	Supported	Supported
VxVM 4.1			Supported	Supported	Supported
VxVM 4.0		Supported	Supported	Supported	
VxVM 3.5	Supported	Supported	Supported	Supported	
VxVM 3.5 MP 4	Supported	Supported	Supported	Supported	

Note: To upgrade from VxVM 3.5 MP4 to VxVM 5.0, the Solaris operating system must first be upgraded to Solaris 8 or 9.

VxVM and Solaris operating system upgrade paths

Table 4-2 shows possible VxVM/Solaris upgrade paths.

Note: If you are upgrading from a release earlier than VxVM 3.5, you must upgrade to VxVM 4.0 before upgrading to VxVM 5.0.

Table 4-2 Upgrade paths

You can upgrade VxVM and Solaris versions:	To VxVM and Solaris versions:	Notes:
VxVM 3.5, 4.0, 4.1 & Solaris 2.6, 7, 8, 9, 10	VxVM 5.0 & Solaris 8, 9, 10	
You can upgrade Solaris OS only:	To Solaris OS versions:	
VxVM 4.0 & Solaris 7,8, 9	VxVM 5.0 & Solaris 8, 9,10	

Note: If the `upgrade_start` script fails for any reason, run the `upgrade_finish` script to undo any changes already made. Verify that the system is restored by comparing `/etc/system`, `/etc/vfstab`, and the output of the `format` command. Then determine and correct the cause of the `upgrade_start` failure. If you cannot correct the problem in a timely manner, restore the `vfstab` file to the version saved in “[Preparing for the upgrade](#)” on page 69, restore any other applications, and perform an `init 6` to completely restore the system

Note: You cannot use the `upgrade_start` `upgrade_finish` scripts to upgrade from VxVM 3.5 MP4.

Determining if the root disk is encapsulated

Before you upgrade, you also need to determine if the root disk is encapsulated. See “[Determining if the root disk is encapsulated](#)” on page 70 for more information.

After you complete the upgrade procedure, read “[Configuring the Veritas Storage Foundation software](#)” on page 101 for important details about initializing (where required), setting up and using the Veritas Storage Foundation software.

Retaining older Veritas Volume Manager releases

To retain an older version of Veritas Volume Manager and not upgrade to the current release shipped with Veritas Storage Foundation, you must set a special environment variable before starting the upgrade procedure.

- If you are using Bourne or Korn shell (`sh` or `ksh`), use the commands:

```
$ NO_VXVM_UPGRADE=yes  
$ export NO_VXVM_UPGRADE
```

- If you are using a C shell (`csh` or `tcsh`), use the command:

```
% setenv NO_VXVM_UPGRADE yes
```

The installation program checks for the `NO_VXVM_UPGRADE` environment variable during the upgrade and displays messages similar to the following if this variable is set:

```
Checking existing package installation...  
NOTICE: Package VRTSvxvm is being skipped (NO_VXVM_UPGRADE set).  
NOTICE: Package VRTSvmsa is being skipped (NO_VXVM_UPGRADE set).
```

Note: We recommend that you upgrade to the current version of Veritas Volume Manager if you want to use the Database FlashSnap or Storage Mapping options.

Upgrading the disk group version separately

If you plan to implement Veritas FastResync and Database FlashSnap, you must use Disk Group Version 90 or later, which is supported in Veritas Volume Manager 3.2 and higher. However, if you choose to retain a version of Veritas Volume Manager older than 3.2, which does not contain Disk Group Version 90 or later, upgrade your disk group version separately. The latest Disk Group Version is 140.

Note: You cannot return to an older disk group version once you have upgraded.

To determine which disk group version you are using

Use the `vx dg list` command as follows:

```
# vx dg -q list diskgroup | grep version
```

where the `-q` option means no header is printed describing output fields.

Example

To determine the disk group version on `PRODDg`:

```
# vx dg -q list PRODDg | grep version  
version: 90
```

To upgrade the disk group version

Use the `vxldg upgrade` command as follows:

```
# vxldg -T diskgroup_version upgrade diskgroup
```

where the `-T` option means upgrading the disk group to a specific version.

For example, to upgrade the disk group to version 140 on `PRODDg`:

```
# vxldg -T 140 upgrade PRODDg
```

Preparing to upgrade Veritas File System

This section lists required pre-upgrade tasks for VxFS.

Upgrading VxFS disk layout versions

VxFS 5.0 allows you to mount and create the following file system disk layouts:

- Disk layout Version 7
- Disk layout Version 6
- Disk layout Version 5
- Disk layout Version 4

Disk layout Version 1, Version 2, and Version 3 are not supported on VxFS 5.0. In the next major release of VxFS, disk layouts Version 4 and 5 will no longer be supported.

To determine the disk layout version of a VxFS file system, run the `fstyp` command on the file system physical device. For example:

```
# /opt/VRTS/bin/fstyp -v /dev/vx/dsk/rootdg/volname | \
grep version
magic a501fcf5 version 7 ctime Thu May 12 11:29:31 2006
```

Deciding when to upgrade disk layout versions

To use the extended features available in the VxFS 5.0 release, upgrade older disk layout versions to disk layout Version 7. See the *Veritas Storage Foundation Release Notes* for information on new Veritas File System 5.0 features.

Deciding when to use `vxupgrade` or `vxfsconvert`

You can use the `vxupgrade` command to upgrade an earlier VxFS disk layout to disk layout Version 6 or Version 7 while the file system remains mounted.

You can use the `vxfsconvert` command to upgrade an earlier VxFS disk layout to a higher disk layout version when the file system is unmounted.

Disk layout Version 1, Version 2, and Version 3 cannot be mounted on VxFS 5.0. You can upgrade these layout versions online before installing VxFS 5.0, or

upgrade them using `vxfsconvert` after installing VxFS 5.0, as shown in the following table:

	Disk Layout Version 1	Disk Layout Version 2	Disk Layout Version 4	Disk Layout Version 5
VxFS Release 3.5 MP2 or lower	Use <code>vxupgrade</code> to upgrade to disk layout Version 4 or Version 5.			
VxFS Release 4.0 or higher	Use <code>vxfsconvert</code> to upgrade to disk layout Version 4.		Use <code>vxupgrade</code> to upgrade to disk layout Version 5 or Version 6.	

The `vxupgrade` command does not upgrade previous disk layouts directly to Version 7. You must upgrade older disk layouts in stages. For example, a Version 4 file system disk layout must first be upgraded to Version 5, then to Version 6, then to Version 7 in three separate invocations of the command:

```
# vxupgrade -n 5 /mount_point
# vxupgrade -n 6 /mount_point
# vxupgrade -n 7 /mount_point
```

The `vxfsconvert` command converts any older disk layout versions directly to Version 5, but you must use `vxupgrade` to convert from Version 5 to Version 6 or Version 7. See the `vxfsconvert(1M)`, `vxupgrade(1M)`, and `fsadm(1M)` manual pages for more information on upgrading VxFS file systems.

Note: The contents of intent logs created on previous disk layout versions cannot be used after the disk layout version is upgraded.

Estimating space and time requirements for upgrading to disk layout Version 7

Converting a Version 5 or Version 6 disk layout to a Version 7 disk layout requires adequate free space to complete. The space and time required to complete the upgrade increases with the number of files, extended attributes, and hard links in the file system. Typical maximum space is at least two additional inodes with one block for every inode. Allow at least ten minutes to upgrade for every million inodes in the file system.

Make sure the file systems are clean

Prior to upgrading to release 5.0, verify that all file systems have been cleanly unmounted by running the `fsdb` command from the existing release of File System.

- 1 Verify that all file systems have been cleanly unmounted:

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

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

- 2 If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem  
# mount -F vxfs [Block_Device] [mountpoint]  
# umount [mountpoint]
```

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

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

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system  
file system still in progress.
```

- 3 If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
- 4 Repeat [step 1](#) to verify that the unclean file system is now clean.

Upgrading Veritas Storage Foundation

You can upgrade a Veritas Storage Foundation product with the product installer. We recommend that you perform this upgrade from single-user mode. No VxFS file systems can be in use at the time of the upgrade.

The following procedures are for Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation for Sybase.

Upgrading Veritas Storage Foundation on Solaris 8, 9, or 10

This section describes upgrading to the current Veritas Storage Foundation if you are already running Solaris 8 or 9, the root disk is unencapsulated, and do not intend to upgrade your Solaris version. If you are running `SUNWvxvm`, go to [“Upgrading from SUNWvxvm”](#) on page 97. If your root disk is encapsulated, go to [“Upgrading Veritas Storage Foundation on an encapsulated root disk”](#) on page 85.

To upgrade a Veritas Storage Foundation product

- 1 If you are upgrading Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, resynchronize all existing snapshots before upgrading.

For Veritas Storage Foundation for DB2:

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

For Veritas Storage Foundation for Oracle:

```
# /opt/VRTS/bin/dbed_vmsnap -S $ORACLE_SID -f SNAPPLAN \  
-o resync
```

- 2 Stop the VEA service:

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 3 Make sure the root disk is not encapsulated. See [“Determining if the root disk is encapsulated”](#) on page 70 for more information.
- 4 Unmount any mounted VxFS file systems.

Note: If you are upgrading multiple hosts, the installer supports the upgrade of the same version of VxVM and VxFS from 3.5, 4.0 or 4.1 to 5.0. Hosts must be upgraded separately if they are running mixed versions such as 3.5 and 4.0.

If any VxFS file systems are mounted with the QuickLog feature, QuickLog must be disabled before upgrading. See the “Veritas QuickLog” chapter of the *Veritas File System Administrator’s Guide* for more information.

- 5 If you are upgrading a high availability (HA) product, take all service groups offline.

List all service groups:

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

For each service group listed, take it offline:

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

- 6 If the VxFS NetBackup libraries package (`VRTSfsnbl`) is installed, remove it before you install the new packages.

To remove the package, use the `pkgrm` command as follows:

```
# pkgrm VRTSfsnbl
```

Respond to any system messages as needed.

The libraries contained in this package are included in the `VRTSvxfs` package in 5.0.

- 7 Reboot the machine to single-user mode (using a command such as `/usr/sbin/shutdown -y -g grace_period -i s`).
- 8 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation. When the system is in single-user mode, `/opt` and `/var` are not normally mounted.
- 9 Load and mount the disc as described in “[Mounting the software disc](#)” on page 32. Then, follow the installation instructions.
- 10 To invoke the common installer, run the `install` command on the disc as shown in this example:

```
# cd /cdrom/cdrom0  
# ./installer
```

Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.

At the product status page, enter **C** for the Configure an Installed Product and press Return. The product installer is displayed.

Upgrading Veritas Storage Foundation on an encapsulated root disk

This section explains how to upgrade Veritas Storage Foundation if the root disk is encapsulated.

The following procedures are for Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation for Sybase.

Note: If your root disk is not encapsulated, see “[Upgrading Veritas Storage Foundation on Solaris 8, 9, or 10](#)” on page 84.

If you are upgrading a Veritas Storage Foundation product on a system with an encapsulated root disk, you must first remove the previous release of Veritas Volume Manager.

Note: It is important that you follow these steps in the specified order.

To upgrade on a system with an encapsulated root disk

- 1 Load and mount the disc as described in “[Mounting the software disc](#)” on page 32. Then, follow the installation instructions.
- 2 Run the `upgrade_check` command to see the status of the root disk.
`# upgrade_check`
- 3 Run the `upgrade_start` script to prepare the previous release of Veritas Volume Manager for removal:

- If the disc is mounted automatically, enter:

```
# cd /cdrom/cdrom0/storage_foundation/scripts
# ./upgrade_start
```

- If the disc is mounted manually, enter:

```
# /mount_point/storage_foundation/scripts
# ./upgrade_start
```

The `upgrade_start` script looks for volumes containing file systems. If certain key file systems must be converted back to using partitions, the script indicates that a reboot is necessary. If any volumes are in use, you must either unmount those volumes or reboot to single-user mode.

If any error messages are encountered during the `upgrade_start` process, run the `upgrade_finish` script immediately to restore the system. Examine the error messages and address them accordingly. Then, run the `upgrade_start` script again. Do not proceed to Step 3 until `upgrade_start` runs without producing any errors.

- 4 Boot the machine to single-user mode (using a command such as `shutdown`).
- 5 If the VxFS NetBackup libraries package (`VRTSfnsbl`) is installed, remove it before you install the new packages.

To remove the package, use the `pkgrm` command as follows:

```
# pkgrm VRTSfnsbl
```

Respond to any system messages as needed.

The libraries contained in this package are included in the `VRTSvxfs` package in 4.0.

- 6 Remove any existing VxVM packages.

For example, if you are removing older VxVM packages, which could be stand-alone or part of Veritas Foundation Suite 3.5, follow this sequence:

```
# pkgrm VRTSvmdoc VRTSvmman VRTSvmpro VRTSvxvm
```

or

If you have multiple versions of VxVM loaded, remove the packages using this sequence:

```
# pkgrm VRTSvmdoc VRTSvmman VRTSvmpro VRTSvxvm.*
```

Note: If you have made any additional modifications or have applied any patches, remove these before removing the `VRTSVxvm` package. If you are using `SUNWvxvm` and `SUNWvxva` packages, see the Sun documentation for details on how to remove `SUNWvxvm` and `SUNWvxva` patches and packages.

- 7 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation. When the system is in single user mode, `/opt` and `/var` are not normally mounted.
- 8 Remount the disc manually.

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /mount_point
```

where `c0t6d0s2` is the default address for the DVD-ROM drive. For more information, see “[Mounting the software disc](#)” on page 32.
- 9 To invoke the common installer, run the `install` command on the disc as shown in this example:

```
# cd /cdrom/cdrom0
# ./installer
```
- 10 Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
If you do not intend to view or print the online documentation, you can omit the `VRTSdbdoc`, `VRTSfsdoc`, and `VRTSvmdoc` packages. If you do not intend to use the GUI, you can omit the `VRTSobgui` package.
- 11 If VxFS file systems are commented out in the `/etc/vfstab` file, uncomment them.
- 12 To disable the hot-relocation feature, comment out the `vxrelocd` line. For Solaris 8 and 9, this line is in the startup file `/etc/init.d/vxvm-recover`.

```
# vxrelocd root & commented out to disable hot-relocation
```

For Solaris 10, this line is in the startup file `/lib/svc/method/vxvm-recover`:

```
# vxrelocd root & commented out to disable hot-relocation
```
- 13 Complete the upgrade using the `upgrade_finish` script.

```
# /mount_point/storage_foundation/scripts/upgrade_finish
```
- 14 Perform a reconfiguration reboot.

```
# reboot -- -r
```

At this point, your pre-upgrade configuration should be in effect and any file systems previously defined on volumes should be defined and mounted.
- 15 Importing a pre-5.0 Veritas Volume Manager disk group does not automatically upgrade the disk group version to the VxVM 5.0 level. You

may need to manually upgrade each of your disk groups following a VxVM upgrade. See “[Upgrading the disk group version separately](#)” on page 80.

- 16 Follow the instructions in “[Veritas Enterprise Administrator](#)” on page 23 to set up and start VEA.

Upgrading Veritas Storage Foundation and earlier Solaris releases

This section describes upgrading to the current Veritas Storage Foundation if you are not running Solaris 8, 9, or 10, have an unencapsulated root disk and need to upgrade the Solaris version in addition to the Veritas packages. If you are running SUNWvxxvm, go to “[Upgrading from SUNWvxxvm](#)” on page 97.

This upgrade procedure allows you to retain existing VxVM and VxFS configurations. After upgrading, you can resume using your file systems and volumes as before (without having to run `vxinstall` again).

Note: It is important that you follow these steps in the specified order.

To prepare to upgrade the operating system

- 1 Ensure that if any of the file systems `/`, `/usr`, `/var`, or `/opt` are defined on volumes, at least one plex for each of those volumes is formed from a single subdisk that begins on a cylinder boundary.

This is necessary because part of the upgrade process involves temporarily converting file systems on volumes back to using direct disk partitions, and Solaris requires that disk partitions start on cylinder boundaries. Upgrade scripts (located in the `/scripts` directory on the disc) automatically convert file systems on volumes back to using regular disk partitions, as necessary. If the upgrade scripts detect any problems (such as lack of cylinder alignment), they display an explanation of the problem and the upgrade does not proceed.

- 2 Check to see if any VxFS file systems are mounted using the `df` command. If so, unmount them by using the `umountall` command:

```
# df -F vxfs
# umountall -F vxfs
```

Note: You need to remove previous versions of VxFS packages prior to upgrading the operating system and the Veritas Storage Foundation product. You do not need to remove existing VxFS file systems, but all of them must remain unmounted throughout the upgrade process. If any VxFS file systems are mounted with the QuickLog option, QuickLog must be disabled.

- 3 Check to see if any Veritas packages are on the machine using the `pkginfo` command.

```
# pkginfo | grep VRTS
```

If any VxFS packages, including Veritas Quick I/O (`VRTSfdd` or `VRTSqio`) or Veritas QuickLog (`VRTSvxld` or `VRTSqlog`) are present, remove them as shown in Step 4.

- 4 Remove the VxFS related packages using the `pkgrm` command.

```
# pkgrm VRTSfsdoc VRTSfdd VRTSqio VRTSvxld VRTSqlog VRTSvxfs
\
VRTSfnsbl
```

Note: If the VxFS NetBackup libraries package (`VRTSfnsbl`) is installed, remove it before upgrading.

- 5 Load and mount the disc as described in “[Mounting the software disc](#)” on page 32.

- 6 Run the `upgrade_start` script to prepare the previous release of the Volume Manager for its removal.

- If the disc is mounted automatically, enter:

```
# /cdrom/cdrom0/storage_foundation/scripts/upgrade_start
```

- If the disc is mounted manually, enter:

```
# /mount_point/storage_foundation/scripts/upgrade_start
```

The `upgrade_start` script looks for volumes containing file systems. If certain key file systems must be converted back to using partitions, the script indicates that a reboot is necessary. If any volumes are in use, you must either unmount those volumes or reboot to single-user mode.

- 7 Reboot the machine to single-user mode (using a command such as `shutdown`).

- 8 Remove any existing VxVM packages.

For example, if you are removing older VxVM packages, which could be stand-alone or part of Veritas Foundation Suite 3.5, follow this sequence:

```
# pkgrm VRTSvmdoc VRTSvmman VRTSvmpro VRTSvxvm
```

or

If you have multiple versions of VxVM loaded, remove the packages using this sequence:

```
# pkgrm VRTSvmdoc VRTSvmman VRTSvmpro VRTSvxvm.*
```

Note: If you have made any additional modifications or have applied any patches, remove these before removing the `VRTSvxvm` package. If you are using `SUNWvxvm` and `SUNWvxva` packages, see the Sun documentation for details on how to remove `SUNWvxvm` and `SUNWvxva` patches and packages.

- 9 If you have VxFS file systems specified in the `/etc/vfstab` file, comment them out.
- 10 Shut down and halt the machine (using a command such as `shutdown`).

To upgrade the operating system

- 1 Upgrade the operating system to Solaris 8, 9, or 10, according to the Solaris installation documentation.

Note: Instructions on upgrading the operating system are beyond the scope of this document.

- 2 Reboot the machine to single-user mode (using a command such as `shutdown`).

To upgrade the Veritas Storage Foundation packages after upgrading the operating system

- 1 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation. When the system is in single user mode, `/opt` and `/var` are not normally mounted.

- 2 Remount the disc manually.

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /mount_point
```

where `c0t6d0s2` is the default address for the DVD-ROM drive.

- 3 To invoke the common installer, run the `install` command on the disc as shown in this example:

```
# cd /cdrom/cdrom0
# ./installer
```

- 4 Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.

If you do not intend to view or print the online documentation, you can omit the `VRTSdbdoc`, `VRTSfsdoc`, and `VRTSvmdoc` packages. If you do not intend to use the GUI, you can omit the `VRTSobgui` package.

- 5 If you commented out VxFS File System entries in the `/etc/vfstab` file in [step 9](#) of the pre-upgrade procedures, uncomment them.

- 6 To disable the hot-relocation feature, comment out the `vxrelocd` line.

For Solaris 8 and 9, this line is in the startup file
`/etc/init.d/vxvm-recover:`

```
# vxrelocd root & # commented out to disable hot-relocation
```

For Solaris 10, this line is in the startup file

```
/lib/svc/method/vxvm-recover:
```

```
# vxrelocd root & # commented out to disable hot-relocation
```

- 7 Complete the upgrade using the `upgrade_finish` script.

```
# /mount_point/storage_foundation/scripts/upgrade_finish
```
- 8 Reboot the machine to multiuser mode (using a command such as `shutdown`).
 At this point, your pre-upgrade configuration should be in effect and any file systems previously defined on volumes should be defined and mounted.
- 9 Importing a pre-5.0 Veritas Volume Manager disk group does not automatically upgrade the disk group version to the VxVM 5.0 level. You may need to manually upgrade each of your disk groups following a VxVM upgrade. See [“Upgrading the disk group version separately”](#) on page 80.
- 10 Follow the instructions in [“Veritas Enterprise Administrator”](#) on page 23 to set up and start VEA.

Upgrading from Database Edition 3.5 for Oracle with Database FlashSnap

In this release of Veritas Storage Foundation for Oracle, Database FlashSnap does not support snapshots of `vxdbavol` and `ORACLE_HOME`. *After* you have upgraded from Veritas Database Edition 3.5 for Oracle with Database FlashSnap to this release, you must remove the volume plexes for `vxdbavol` and `ORACLE_HOME`, and revalidate the snapplan before using Database FlashSnap with this release of Veritas Storage Foundation for Oracle.

To remove the volume plexes for `vxdbavol` and `ORACLE_HOME`

- 1 As root, snapback the snapshot plexes.

```
# vxassist [-g diskgroup] snapback snapvol
```
- 2 Turn off FastResync.

```
# vxvol [-g diskgroup] set fastresync=off volume
```
- 3 Remove the DCO object.

```
# vxassist [-g diskgroup] remove log volume logtype=dc
```
- 4 Remove the volume plexes for `vxdbavol` and `ORACLE_HOME`.

```
# vxplex -g diskgroup -o rm dis plex_name
```
- 5 Log in as the DBA user and revalidate your snapplan.

```
$ /opt/VRTS/bin/dbed_vmchecksnap -S ORACLE_SID -H  
ORACLE_HOME \  
-f SNAPPLAN -o validate -F
```
- 6 See the [“Using Database FlashSnap for Backup and Off-Host Processing”](#) chapter of the *Veritas Storage Foundation for Oracle Administrator’s Guide* for information on how to validate the snapplan.

Migrating from /etc/vx/vxdba to /var/vx/vxdba for DB2 and Oracle

If you are upgrading Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, you can migrate to /var/vx/vxdba to save space under the root partition. Migrating to /var/vx/vxdba is optional. However, if you do not perform this migration, you cannot remove any file or directory from /etc/vx/vxdba to ensure proper operation. This procedure can be done at any time.

Note: This does not apply to Veritas Storage Foundation for Sybase.

To migrate from /etc/vx/vxdba to /var/vx/vxdba

- 1 Copy the /etc/vx/vxdba directory and contents to /var/vx/vxdba

```
# cp -rp /etc/vx/vxdba /var/vx/vxdba
```
- 2 Remove /etc/vx/vxdba.

```
# rm -rf /etc/vx/vxdba
```
- 3 Link the two directories.

```
# ln -s /var/vx/vxdba /etc/vx/vxdba
```

Upgrading the VEA client on a Solaris system

For the current version of VERITAS Enterprise Administrator, four packages are required (VRTSobgui, VRTSat, VRTSspb, and VRTSicsco).

Previous versions may have required only the VRTSobgui package.

To upgrade the VEA client on a Solaris machine (other than the host) using pkgadd

- 1 Log in as superuser (root).
- 2 First, check to determine whether the VEA client package is already installed.

```
# pkginfo | grep VRTSobgui
```

This command will return VRTSobgui if VRTSobgui is already installed. It will return nothing if the package has not been installed.
- 3 If the package is installed, remove it using the pkgrm command.

```
# pkgrm VRTSobgui
```
- 4 To install the new VEA client package for Solaris, insert the appropriate media disc into your system's DVD-ROM drive.

- 5 Copy the `VRTSobgui.tar.gz` package, and the dependent packages, to the current working directory on your system.

```
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSobgui.tar.gz .
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSat.tar.gz .
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSicsco.tar.gz .
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSspb.tar.gz .
```
- 6 Decompress the packages, and then extract the contents.

```
# /cdrom/cdrom0/gnu/gunzip *.tar.gz
# tar xvf VRTSobgui.tar
# tar xvf VRTSat.tar
# tar xvf VRTSicsco.tar
# tar xvf VRTSspb.tar
```
- 7 Use the `pkgadd` command to install the packages, in the following order, to allow for dependencies. Answer any questions, as the installation proceeds.

```
# pkgadd -d . VRTSicsco
# pkgadd -d . VRTSat
# pkgadd -d . VRTSspb
# pkgadd -d . VRTSobgui
```
- 8 The VEA client package for Solaris, and dependent packages, are installed.

Upgrading the VEA client on a Windows system

To upgrade the VEA client on 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.
After removing the existing package, install the new VEA client package by following the procedure in [“Installing the Veritas Enterprise Administrator client”](#) on page 41.

Upgrading to the new repository database for DB2 and Oracle

If you are installing or upgrading Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, you need to either create a new repository database or migrate your old repository database to the new one. To use the `db2ed_update` or `dbed_update` command, you must be the instance owner or database administrator.

Note: This does not apply to Veritas Storage Foundation for Sybase.

To upgrade your repository database

- 1 Create and configure the new repository database with the `sfua_db_config` command.

```
# /opt/VRTS/dbcom/bin/sfua_db_config
```

- 2 Migrate your old repository information into the new repository database. If you are upgrading Veritas Storage Foundation for DB2 in a single-host environment, run the `db2ed_update` command.

```
# /opt/VRTS/bin/db2ed_update -D DB2DATABASE
```

If you are upgrading Veritas Storage Foundation for DB2 in a high availability (HA) environment, run the `db2ed_update` command with the `-G` option.

```
# /opt/VRTS/bin/db2ed_update -D DB2DATABASE -G service_group
```

If you are upgrading Veritas Storage Foundation for Oracle in a single-host environment, run the `dbed_update` command.

```
# /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

If you are upgrading Veritas Storage Foundation for Oracle in a high availability (HA) environment, run the `dbed_update` command with the `-G` option.

```
# /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME \  
-G service_group
```

After the upgrade, the old repository database will be marked with a hidden file name, such as `/etc/vx/vxdba/.upgrade_to_5.0`, to prevent further updates. If you need to perform an additional upgrade, the file must be removed.

Verifying the Veritas Storage Foundation upgrade

Use the following sections to verify the product upgrade.

Checking Volume Manager processes

To confirm that key Volume Manager processes are running

- ◆ Type the following command:

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

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

Checking cluster operation

You need to check cluster operation only if you installed and configured an HA version of the Storage Foundation software.

To verify that the cluster is operating

- ◆ Type the following command on any node:

```
# hastatus -summary
```

```
-- SYSTEM STATE
-- System          State          Frozen

A host1            RUNNING       0
A host2            RUNNING       0

-- GROUP STATE
-- Group           System        Probed  AutoDisabled  State

B ClusterService host1        Y      N              ONLINE
B ClusterService host2        Y      N              OFFLINE
```

Identify the system state of each node in the output of this command. If the value is `RUNNING` for all the nodes, VCS is successfully installed and running.

Refer to the `hastatus(1M)` manual page and the *Veritas Cluster Server User's Guide* for more information on system states and state transitions.

About Low Latency Transport

The file `llthosts(4)` is a database containing one entry per node that links the Low Latency Transport (LLT) system ID (in the first column) with the LLT host name. This file is identical on each cluster node.

Based on the sample installation, the file `/etc/llthosts` contains the entries:

```
0 host1
1 host2
```

The file `llttab(1M)` contains information derived from the installation and used by the utility `lltconfig(1M)`. After installation, this file lists the network links that correspond to the specific node.

The first line identifies the node. The second line identifies the cluster, based on the cluster ID entered during installation. The next two lines, beginning with the `link` command, identify the two network cards used by the LLT protocol.

See the `llttab(4)` manual page for details on how to modify the LLT configuration. The manual page describes ordering the directives in the `llttab` file.

Verifying LLT

Use the `lltstat` command to verify that LLT links are active for each system in the cluster.

To verify that links are active for LLT

- ◆ Use the `lltstat -n` as follows:

```
# lltstat -n
```

With LLT configured correctly, the output of `lltstat -n` shows all of the nodes in the cluster and two links for each node. If the output shows otherwise, type `lltstat -nvv | more` on any node to view additional information about LLT.

To obtain LLT port information

- ◆ Use the `lltstat -p` command as follows:

```
# lltstat -p
```

About Group Membership and Atomic Broadcast

After installation, the file `/etc/gabtab` contains a `gabconfig(1M)` command that configures the Group Membership and Atomic Broadcast (GAB) driver.

The file `/etc/gabtab` contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

where the `-c` option configures the driver and `-nN` specifies the cluster will not be formed until at least N nodes are ready. The variable N represents the number of cluster nodes.

Verifying GAB

To verify that GAB is operating

- ◆ Type the following command on each system:

```
# /sbin/gabconfig -a
```

If GAB is operating, the following GAB port membership information is returned:

```
GAB Port Memberships
=====
Port a gen a36e0003 membership 01
Port h gen fd570002 membership 01
```

Port a indicates that GAB is communicating, gen a36e0003 is a randomly generated number, and membership 01 indicates that nodes 0 and 1 are connected.

Port h indicates that VCS is started, gen fd570002 is a randomly generated number, and membership 01 indicates that nodes 0 and 1 are both running VCS.

If GAB is not operating, no GAB port membership information is returned:

```
GAB Port Memberships
=====
```

If only one network is connected, the following GAB port membership information is returned:

```
GAB Port Memberships
=====
Port a gen a36e0003 membership 01
Port a gen a36e0003 jeopardy 1
Port h gen fd570002 membership 01
Port h gen fd570002 jeopardy 1
```

For more information on GAB, including descriptions of ports, refer to the *Veritas Cluster Server User's Guide*.

Upgrading from SUNWvxcvm

This section provides instructions for upgrading SUNWvxcvm.

To upgrade from SUNWvxcvm if the root disk is unencapsulated

- 1 Run the `upgrade_start` script to prepare the previous release of Veritas Volume Manager for its removal:

- If the disc is mounted automatically, enter:

```
# /cdrom/cdrom0/storage_foundation/scripts/upgrade_start
```

The `upgrade_start` script looks for volumes containing file systems. If certain key file systems must be converted back to using partitions, the script will indicate that a reboot is necessary. If so, reboot the machine to single-user mode before proceeding to the next step. If any volumes are in use, you must either unmount those volumes or reboot to single-user mode.

- 2 Reboot the machine to single-user mode (using a command such as `shutdown`).
- 3 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation. When the system is in single user mode, `/opt` and `/var` are not normally mounted.
- 4 Remove all `SUNWvxxvm` and `SUNWvxxva` patches and packages. Refer to Sun documentation for this procedure.
- 5 Load and mount the disc as described in “[Mounting the software disc](#)” on page 32.
- 6 Run the `installsf` program.
 - If the disc is mounted automatically, enter:

```
# cd /cdrom/cdrom0/storage_foundation
# ./installsf
```
 - If the disc is mounted manually, enter:

```
# cd /mount_point/storage_foundation
# ./installsf
```
- 7 When you see the prompt for a license key, enter the appropriate license key. The `installsf` program then installs the Veritas Storage Foundation to include the necessary packages on your system.
- 8 Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
- 9 Complete the upgrade by using the `upgrade_finish` command:

```
# /cdrom/cdrom0/storage_foundation/scripts/upgrade_finish
```

or

```
# /mount_point/storage_foundation/scripts/upgrade_finish
```
- 10 Reboot the machine to multiuser mode (using a command such as `shutdown`).

At this point, your pre-upgrade configuration should be in effect and any file systems previously defined on volumes should be defined and mounted.

Upgrading language packages

If you are upgrading Veritas Storage Foundation in a language other than English, you must install the required language packages *after* installing the English packages. If you are planning to use the GUI, you must install the language package for the VEA client.

See “[Installing language packages](#)” on page 53.

Configuring the Veritas Storage Foundation software

This chapter covers important details about initializing (where required), setting up, and configuring the VERITAS software.

Topics covered in this chapter include:

- [“Database configuration requirements”](#) on page 101
- [“Creating and configuring the repository database for DB2 and Oracle”](#) on page 102
- [“Configuring Veritas Storage Foundation”](#) on page 105
- [“Configuring Veritas Volume Manager”](#) on page 110
- [“About hot-relocation”](#) on page 114
- [“Configuring Veritas File System”](#) on page 123

Database configuration requirements

Most relational database management system (RDBMS) software requires operating system parameters to be set prior to operation. The DB2, Oracle, and Sybase databases require modifications to kernel settings in the `/etc/system` file in Solaris 9 before the databases will run correctly. In Solaris 10, system parameters are managed through the Resource Controls facility. The most critical settings are normally located in the Shared Memory and Semaphore settings on Solaris. For precise settings, consult your current database installation and configuration documentation.

Creating and configuring the repository database for DB2 and Oracle

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

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

Before running the script, review the following requirements:

- | | |
|--------------------------------|---|
| In a stand-alone configuration | <ul style="list-style-type: none">■ 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. |
| In an HA configuration | <ul style="list-style-type: none">■ 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 example: <code>hme0 / bge0 /</code>) for all systems in the cluster.■ 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:
<pre># haconf -dump -makero</pre> |

Table 5-3 indicates the options available for the `sfua_db_config` script.

Table 5-3 `sfua_db_config` options

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

The following is an example of configuring Veritas Storage Foundation for Oracle:

```
Welcome to the SFORA configuration script.
This script creates repository for standalone and HA
configuration.
Please create a Veritas File System on a Veritas Volume
and mount it, before starting configuration using this script.
This mount point will be used to store repository.
```

The following is required to configure SFORA repository for HA solution:

- * A mount point of already mounted Veritas Volume on a shared

storage, with Veritas File system.
* A public NIC used by each system in the cluster.
* A Virtual IP address and netmask.

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

filesystem mount point for SFORA repository: /sfua_rep

Discovering public NIC on host1 ...bge0

Enter the NIC for system host1 for HA Repository configuration:
[bge0]

Discovering public NIC on host2 ...bge0

Enter the NIC for system host2 for HA Repository configuration:
[bge0]

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

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

Following information will be used for SFORA HA configuration:

Public IP address:	xx.xxx.xxx.xxx
Subnet mask:	xxx.xxx.xxx.x
Public interface:	host1 -> bge0, host2 -> bge0
Mount point:	/sfua_rep
Volume Name for mount point:	dbed_rep
Diskgroup for mount point:	sfua_rep

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

Adding repository information in VCS (HA) configuration...

Added repository information successfully in VCS (HA) configuration.

Repository configuration completed successfully for HA environment.

2 Verify that the repository was configured.

If you are installing in a high availability configuration:

```
# /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.

If you are installing in a stand-alone configuration:

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

Configuring Veritas Storage Foundation

Once you install and initialize all of the VERITAS software, you can take advantage of the various storage management features to ease the task of system and database administration.

Setting administrative permissions

To allow database administrators to administer a database using Veritas Storage Foundation, you are required to change some permission settings. You are asked during the installation process if you want to allow database administrators access to various functionality. If you did not make the permission changes during installation, you can do so at a later time.

Setting permissions for DB2

The default settings at installation time for the `/opt/VRTSdb2ed` directory allow only the `root` login to access the directory.

To allow the user “db2inst1” access to the `/opt/VRTSdb2ed` directory

Use the `chown` and `chmod` commands as follows:

```
# chown db2inst1 /opt/VRTSdb2ed
# chmod 750 /opt/VRTSdb2ed
```

To allow users in the group “db2iadm1” access to the `/opt/VRTSdb2ed` directory

Use the `chgrp` and `chmod` commands as follows:

```
# chgrp db2iadm1 /opt/VRTSdb2ed
# chmod 750 /opt/VRTSdb2ed
```

Setting permissions for Oracle

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

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

Use the `chown` and `chmod` commands, as follows:

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

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

Use the `chgrp` and `chmod` commands, as follows:

```
# chgrp dba /opt/VRTSdbed
# chmod 750 /opt/VRTSdbed
```

Setting permissions for Sybase

No changes are required.

Veritas Enterprise Administrator

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

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.

```
# pkginfo -l VRTSd2gui | egrep STATUS
STATUS: completely installed
```

- 2 Stop the VEA server.

```
# /opt/VRTS/bin/vxsvcctl 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`, `ldap`, `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 -n group_name
```

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

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 start
```

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.

Note: You cannot remove `root` from the AZDB.

To remove users other than `root` from the Veritas Enterprise Administrator AZDB

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

```
# pkginfo -l VRTSd2gui | egrep STATUS
STATUS: completely installed
```

- 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 start
```

Adding Users to Veritas Enterprise Administrator for Oracle

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.

```
# pkginfo -l VRTSorgui | egrep STATUS
STATUS: completely installed
```

- 2 Stop the VEA server.

```
# /opt/VRTS/bin/vxsvcctl stop
```

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

```
# /opt/VRTS/bin/vxdbedusr -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/VRTSdbed` owner.

-n indicates the name of the user.

-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 “oracle” as a user with `root` privileges, enter the following:

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

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

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

where -a adds a user to the registry.

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

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

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

```
# /opt/VRTS/bin/vxdbedusr -a group -n group_name
```

where -a adds the user group to the registry.

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

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

6 Restart the VEA Server.

```
# /opt/VRTS/bin/vxsvcctl start
```

Removing users from Veritas Enterprise Administrator for Oracle

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

Note: You cannot remove `root` from the AZDB.

To remove users other than root from the Veritas Enterprise Administrator AZDB

1 Make sure that the optional GUI package was installed.

```
# pkginfo -l VRTSorgui | egrep STATUS
STATUS: completely installed
```

2 Stop the VEA server.

```
# /opt/VRTS/bin/vxsvcctl stop
```

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

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

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

For example, to remove the user “oracle,” enter the following:

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

4 Restart the VEA Server.

```
# /opt/VRTS/bin/vxsvcctl start
```

vxtunefs command permissions and Cached Quick I/O

By default, you must have `superuser (root)` privileges to use the `/opt/VRTS/bin/vxtunefs` command. The `vxtunefs` command is a tool that lets you change caching policies to enable Cached Quick I/O and change other file system options. Database administrators can be granted permission to change default file system behavior in order to enable and disable Cached Quick I/O. The system administrator must change the `vxtunefs` executable permissions as follows:

```
# chown root /opt/VRTS/bin/vxtunefs
# chgrp dba /opt/VRTS/bin/vxtunefs
# chmod 4550 /opt/VRTS/bin/vxtunefs
```

Note: Setting the permissions for `/opt/VRTS/bin/vxtunefs` to 4550 allows all users in the `dba` group to use the `vxtunefs` command to modify caching behavior for Quick I/O files.

For more information, see the *Veritas File System Administrator's Guide*.

Configuring Veritas Volume Manager

This section explains how to set up VxVM enclosure-based naming. To carry out further tasks such as disk encapsulation or initialization, please see the *Veritas Volume Manager System Administrator's Guide*.

Note: In releases of VxVM (Volume Manager) prior to 4.0, a system installed with Volume Manager was configured with a default disk group, `rootdg`, that had to contain at least one disk. By default, operations were directed to the `rootdg` disk group. From release 4.0 onward, Volume Manager can function without any disk group having been configured. Only when the first disk is placed under Volume Manager control must a disk group be configured. There is no longer a requirement that you name any disk group `rootdg`, and any disk group that is named `rootdg` has no special properties by having this name. During the setup procedures, you will be asked if you want to create a default disk group, and asked to specify its name.

Using `vxinstall` to configure Veritas Volume Manager

Note: If you used the Veritas Installation Menu or the `installvm` script, you do not need to carry out the instructions in this section. Licensing, configuration of enclosure based naming and creation of a default disk group are managed by the menu installer and the `installvm` script.

Because you are no longer required to configure VxVM disks immediately, the `vxinstall` command no longer invokes the `vxdiskadm` program, making it much simpler than in previous releases. The utility provides the following three functions:

- Licensing VxVM
- Enabling Enclosure-based naming
- Setting up a system-wide default disk group

To run the command, enter:

```
# vxinstall
```

which will prompt you to enter a license key:

```
Are you prepared to enter a license key [y,n,q,?] (default: y) y
```

- If you don't have a license key, see [“Symantec product licensing”](#) on page 15.

Note: The presence of certain hardware arrays (for example, A5000) automatically generates a key.

The `vxinstall` program then asks if you want to use enclosure-based naming:

```
Do you want to use enclosure based names for all disks ?  
[y,n,q,?] (default: n)
```

After installation, disks use the traditional naming format, usually `c#t#d#s#`. Enclosure based naming provides an alternative that allows disk devices to be named for enclosures rather than for the controllers through which they are accessed. In a Storage Area Network (SAN) that uses Fibre Channel hubs or fabric switches, information about disk location provided by the operating system may not correctly indicate the physical location of the disks. Enclosure-based naming allows Volume Manager to access enclosures as separate physical entities. By configuring redundant copies of your data on separate enclosures, you can safeguard against failure of one or more enclosures. If you want to use enclosure-based naming, enter 'y' and `vxinstall` asks you whether you want to set up a systemwide default disk group:

```
Do you want to setup a system wide default disk group ?  
[y,n,q,?] (default: y)
```

VxVM will continue with the question:

```
Which disk group [<group>,list,q,?] ?
```

If you know the name of the disk group that you want to use as the default disk group, enter it at the prompt, or use the `list` option and make a selection.

In releases prior to Volume Manager 4.0, the default disk group was `rootdg` (the *root disk group*). For Volume Manager to function, the `rootdg` disk group had to exist and it had to contain at least one disk. This requirement no longer exists, however you may find it convenient to create a system-wide default disk group. For operations that require a disk group, the system-wide default disk group will be used if the VxVM command is not specified with the `-g` option. The main benefit of creating a default disk group is that VxVM commands default to the default disk group and you will not need to use the `-g` option. To verify the default disk group after it has been created, enter the command:

```
# vxdg defaultdg
```

Note: VxVM does not allow you use the following names for the default disk group because they are reserved words: `bootdg`, `defaultdg` and `nodg`.

At this stage, the installation of VxVM is complete. To carry out further tasks such as disk encapsulation or initialization, please see the *Veritas Volume Manager System Administrator's Guide*.

Using Dynamic Multipathing with EMC Symmetrix Disk Arrays

To allow DMP to receive correct enquiry data from EMC Symmetrix disk arrays, the common Serial Number (C-bit) Symmetrix Director parameter must be set to enabled.

Preventing multipathing/suppress devices from VxVM's view

This section describes how to exclude a device that is under VxVM or Dynamic Multipathing control.

To prevent multipathing or suppress devices from being seen by VxVM

- 1 Enter the command
`# vxdiskadm`
- 2 Select menu item 17 (Prevent Multipathing/Suppress devices from VxVM's view) from the `vxdiskadm` main menu.

The following message displays:

```
VxVM INFO V-5-2-1239 This operation might lead to some devices
being suppressed from VxVM's view or prevent them from being
multipathed by vxdkmp. (This operation can be reversed using the
vxdiskadm command).
```

```
Do you want to continue? [y,n,q,?] (default: n) y
```

- 3 Enter **y**.
- 4 Select the appropriate operation:
 - Suppress all paths through a controller from VxVM's view:
Select Option 1.
Enter a controller name when prompted:

```
Enter a controller name: [ctrl_name,all,list,list-exclude,q,?]
```

- Suppress a path from VxVM's view:
Select Option 2.
Enter a path when prompted.

```
Enter a pathname or pattern: [<Pattern>,all,list,list-exclude,q?]
```

- Suppress disks from VxVM's view by specifying a VID:PID combination:

Select Option 3 and read the messages displayed on the screen.

Enter a VID:PID combination when prompted.

Enter a VID:PID combination: [<Pattern>,all,list,exclude,q,?]

The disks that match the VID:PID combination are excluded from VxVM. Obtain the Vendor ID and Product ID from the Standard SCSI inquiry data returned by the disk.

For example, the VID:PID combination for Sun's T3 disk array can be specified as SUN:T3. Obtain the Vendor ID and Product ID of the disk by the command `/usr/lib/vxvm/diag.d/vxdmping`.

- Suppress all but one path to a disk:
 Select Option 4 and read the messages displayed on the screen before specifying a path.
 Enter a path when prompted:

Enter pathgroup: [<pattern>,list,list-exclude,q,?]

The next three options allow you to exclude devices from `vxdmp`.

- Prevent multipathing of all disks on a controller by VxVM.
 Select Option 5 and read the messages displayed on the screen before specifying a controller.

Enter a controller name when prompted. The controller entered is excluded from DMP control.

Enter a controller name: [<ctrl-name>,all,list,list-exclude,q,?]

- Prevent multipathing of a disk by VxVM.
 Select Option 6 to exclude the specified path from multipathing. The corresponding disks are claimed in the OTHER_DISKS category and are not multipathed. Read the messages displayed on the screen before specifying a path.

Enter a path at the prompt:

Enter a pathname or pattern: [<pattern>,all,list,list-exclude,q,?]

- Prevent multipathing of disks by specifying a VID:PID combination.
 Select Option 7 to exclude disks by a VIP:PID combination. All disks returning a VID:PID combination are claimed in the OTHER_DISKS category and are not multipathed. Read the messages displayed on the screen before specifying a VIP:PID.

Enter the VID:PID combination at the prompt.

Enter a VID:PID combination: [<pattern>,all,list,list-exclude,q,?]

Note: If you selected any of the options, reboot the system for device exclusion to take effect.

About hot-relocation

Hot-relocation automatically restores redundancy and access to mirrored and RAID-5 volumes when a disk fails. This is done by relocating the affected subdisks to disks designated as spares and/or free space in the same disk group.

The hot-relocation feature is enabled by default. The associated daemon, `vxrelocd`, is automatically started during system startup.

Follow these recommendations:

- Leave the VxVM hot-relocation feature enabled to detect disk failures automatically. It will notify you of the nature of the failure, attempt to relocate any affected subdisks that are redundant, and initiate recovery procedures.
- Configure at least one hot-relocation spare disk in each disk group. This will allow sufficient space for relocation in the event of a failure.

If you decide to disable hot-relocation, prevent `vxrelocd` from running after you load the VxVM software. See the section “Modifying the behavior of Hot-Relocation” in Chapter 9 of the *VERITAS Volume Manager Administrator’s Guide* for details.

Starting and enabling the configuration daemon

The VxVM configuration daemon (`vxconfigd`) maintains VxVM disk and disk group configurations. The `vxconfigd` communicates configuration changes to the kernel and modifies configuration information stored on disk.

Startup scripts usually invoke `vxconfigd` at system boot time. The `vxconfigd` daemon must be running for VxVM to operate properly.

The following procedures describe how to check that `vxconfigd` is started, whether it is enabled or disabled, how to start it manually, or how to enable it as required.

To determine whether `vxconfigd` is enabled, use the following command:

```
# vxctl mode
```

The following message indicates that the `vxconfigd` daemon is running and enabled:

```
mode: enabled
```

This message indicates that `vxconfigd` is not running:

```
mode: not-running
```

To start the `vxconfigd` daemon, enter the following command:

```
# vxconfigd
```

This message indicates that `vxconfigd` is running, but not enabled:

```
mode: disabled
```

To enable the volume daemon, enter the following command:

```
# vxdctl enable
```

Once started, `vxconfigd` automatically becomes a background process.

By default, `vxconfigd` writes error messages to the console. However, you can configure it to write errors to a log file. For more information, see the `vxconfigd(1M)` and `vxdctl(1M)` manual pages.

Starting the volume I/O daemon

The volume I/O daemon (`vxiod`) provides extended I/O operations without blocking calling processes. Several `vxiod` daemons are usually started at system boot time after initial installation, and they should be running at all times. The procedure below describes how to verify that the `vxiod` daemons are running, and how to start them if necessary.

To verify that `vxiod` daemons are running, enter the following command:

```
# vxiod
```

Note: The `vxiod` daemon is a kernel thread and is not visible using the `ps` command.

If, for example, 16 `vxiod` daemons are running, the following message displays:

```
16 volume I/O daemons running
```

where 16 is the number of `vxiod` daemons currently running. If no `vxiod` daemons are currently running, start some by entering this command:

```
# vxiod set 16
```

where 16 is the desired number of `vxiod` daemons. It is recommended that at least one `vxiod` daemon should be run for each CPU in the system.

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

Starting the VEA server

After installing the VEA packages, the VEA server needs to be stopped and restarted. To check the state of the VEA server, enter:

```
# /opt/VRTS/bin/vxsvcctl status
```

To stop the VEA server, enter:

```
# /opt/VRTS/bin/vxsvcctl stop
```

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

Note: The VEA server is automatically started on a reboot.

To start the VEA server, enter:

```
# /opt/VRTS/bin/vxsvcctl start
```

Starting the VEA client

Only users with appropriate privileges can run VEA. VEA can administer the local machine or a remote machine. However, VxVM and the VEA server must be installed on the machine to be administered. The VxVM `vxconfigd` daemon and the VEA server must be running on the machine to be administered.

After installing VxVM and VEA and starting the server, start the VEA client in one of the following ways.

Solaris operating system

To administer the Solaris machine, use the following command:

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

Windows operating system

To administer a *remote* Solaris machine from a Windows machine, select **Start > Programs > Veritas > Veritas Enterprise Administrator**.

Modifying connection access (optional)

To allow users other than *root* to access VEA, set up a group called *vrtsadm* in `/etc/group`, and add the users to this group. For example, adding the following entry:

```
vrtsadm::600:root,ed
```

will allow the two users, *root* and *ed*, to access VEA.

To specify a group other than *vrtsadm*, you should add the group to */etc/group*, modify the Security key and restart the ISIS server daemon, as in the following example.

- 1 Add a new group:

```
# groupadd -g gid veagr
```
- 2 Edit */etc/group* to add users to the group.
- 3 Modify the Security key in the registry:

```
# /opt/VRTSob/bin/vxregctl /etc/vx/isis/Registry setvalue \  
Software/Veritas/VxSvc/Current/Version/Security AccessGroups \  
REG_SZ veagr
```
- 4 Restart the VEA server.

```
# /opt/VRTS/bin/vxsvcctl restart
```

Enabling cluster support in VxVM (optional)

Note: This section assumes that you are using Sun Java System Cluster as the cluster monitor on your system.

This release includes an *optional* cluster feature that enables VxVM to be used in a Sun Cluster environment. The cluster functionality in VxVM allows multiple hosts to simultaneously access and manage a set of disks under VxVM control. A *cluster* is a set of hosts sharing a set of disks; each host is referred to as a *node* in the cluster.

Note: The VxVM cluster feature requires a license, which can be obtained from your Sun Customer Support channel. (The presence of a SPARCstorage™ Array may serve as a license, but it limits what can be done to private disk groups.)

To enable the cluster functionality in VxVM

- 1 Obtain a license for the VxVM cluster feature.
- 2 Install the software packages onto each system (node) to be included in the cluster.
See [“Installing VxVM using the pkgadd command”](#) on page 142.
See [“Installing the Veritas Enterprise Administrator client”](#) on page 41.
- 3 Initialize VxVM.
See [“Configuring Veritas Volume Manager”](#) on page 110.

- 4 Start VEA.
See “Starting the VEA server” on page 116.
See “Starting the VEA client” on page 116.
- 5 Configure shared disks.
See “Configuring shared disks” on page 118.

Configuring shared disks

This section describes how to configure shared disks. If you are installing VxVM for the first time or adding disks to an existing cluster, you need to configure new shared disks. If you are upgrading VxVM, verify that your shared disks still exist.

The shared disks should be configured from one node only. Since the VxVM software cannot tell whether a disk is shared or not, you must specify which are the shared disks.

Make sure that the shared disks are not being accessed from another node while you are performing the configuration. If you start the cluster on the node where you perform the configuration only, you can prevent disk accesses from other nodes because the quorum control reserves the disks for the single node.

See “Reserving shared disks” on page 121.

Configuring new disks

If you are installing and setting up VxVM for the first time, you must configure the shared disks.

To configure shared disks

- 1 Start the cluster on at least one node.
- 2 On one node, run the `vxdiskadm` program and choose option 1 to initialize new disks. When asked to add these disks to a disk group, choose `none` to leave the disks for future use.
- 3 On other nodes in the cluster, run `vxctl enable` to see the newly initialized disks.
- 4 From the master node, create disk groups on the shared disks. To determine if a node is a master or slave, run `vxctl -c mode`.
Use the `vxdg` program or VEA to create disk groups. In the `vxdg` program, use the `-s` option to create shared disk groups.
- 5 From the master node only, use `vxassist` or VEA to create volumes in the disk groups.

The volumes must be of type `gen`. Do not create RAID-5 volumes. Before creating any log subdisks, read the section on DRL in the *VERITAS Volume Manager Administrator's Guide*.

- 6 If the cluster is only running with one node, bring up the other cluster nodes. Enter the `vxldg list` command on each node to display the shared disk groups.

Verifying existing shared disks

If you are upgrading from a previous release of VxVM, verify that your shared disk groups still exist.

To verify that your shared disk groups exist

- 1 Start the cluster on all nodes.
- 2 Enter the following command on all nodes:

```
# vxldg list
```

This displays the existing shared disk groups.

Converting existing VxVM disk groups to shared disk groups

If you are upgrading from VxVM 3.x to VxVM 4.0 and you want to convert existing disk groups to shared disk groups, configure the shared disks as follows:

- 1 Ensure that all systems that are running are part of the same cluster.
- 2 Configure the disk groups using the following procedure.
To list all disk groups, use the following command:

```
# vxldg list
```
- 3 Determine which node is the master. To determine if a node is a master or a slave, run the following command:

```
# vxldctl -c mode
```


To deport disk groups to be shared, use the following command:

```
# vxldg deport disk-group-name
```


To import disk groups to be shared, use the following command on the master node:

```
# vxldg -s import disk-group-name
```


This procedure marks the disks in the shared disk groups as shared and stamps them with the ID of the cluster, enabling other nodes to recognize the shared disks.
If dirty region logs exist, ensure they are active. If not, replace them with larger ones.

To display the shared flag for all the shared disk groups, use the following command:

```
# vxpdg list
```

The disk groups are now ready to be shared.

- 4 If the cluster is only running with one node, bring up the other cluster nodes. Enter the `vxpdg list` command on each node to display the shared disk groups. This command displays the same list of shared disk groups displayed earlier.
- 5 For information on upgrading in a Cluster Volume Manager (CVM) environment, see the *Veritas Cluster File System Installation and Configuration Guide*.

Upgrading in a clustered environment and with FastResync set

This procedure applies to two upgrade scenarios:

- Upgrading from VxVM 3.5 to VxVM 5.0
- Upgrading from VxVM 3.5 Maintenance Pack 4 or from VxVM 3.2 Maintenance Pack 5 to VxVM 5.0

If there are volumes in the shared disk groups with FastResync set (`fastresync=on`), before beginning the upgrade procedure, reattach each snapshot to its data volume, using this procedure:

To upgrade in a clustered environment when FastResync is set

- 1 You should run this procedure from the master node; to find out if you are on the master node, enter the command:

```
# vxpdctl -c mode
```
- 2 On the master node, list which disk groups are shared by entering:

```
# vxpdg -s list
```
- 3 Using the diskgroup names displayed by the previous command, list the disk groups that have volumes on which FastResync is set:

```
# vxprint -g diskgroup -F "%name" -e "v_fastresync"
```
- 4 Reattach each snapshot:

```
# vxassist -g diskgroup -o nofmr snapback snapshot_volume
```
- 5 If you are upgrading from VxVM 3.5 Maintenance Pack 4 or from VxVM 3.2 Maintenance Pack 5, set FastResync to off for each volume:

```
# vxvol -g diskgroup set fastresync=off volume
```

Reserving shared disks

As part of its quorum control, the Sun Java System Cluster cluster manager reserves shared disk controllers when only one node is active. This prevents “rogue” hosts from accessing the shared disks. When this happens, the `vxdisk list` command used on a node that has left the cluster may show all disks on such a controller as having an `error` status. The more detailed options of the `vxdisk` command show the flag `unavailable`. When a new node joins the cluster, the Sun Java System Cluster software releases the controllers. VxVM attempts to access these disks, and if that is successful, the disks return to an `online` status. (See the Sun Java System Cluster documentation for further details.) If one system boots while the other system has the disks reserved, the disks can be invisible to the booting system, and the `vxdisk` command may not display any of the shared disks. When the system joins the cluster, the shared disks become visible.

Adding new array support

After installation, add any disk arrays that are unsupported by Veritas to the JBOD category as described in the *Veritas Volume Manager Administrator's Guide*.

About hot-relocation

Hot-relocation automatically restores redundancy and access to mirrored and RAID-5 volumes when a disk fails. This is done by relocating the affected subdisks to disks designated as spares and/or free space in the same disk group.

The hot-relocation feature is enabled by default. The associated daemon, `vxrelocd`, is automatically started during system startup.

Follow these recommendations:

- Leave the VxVM hot-relocation feature enabled to detect disk failures automatically. It will notify you of the nature of the failure, attempt to relocate any affected subdisks that are redundant, and initiate recovery procedures.
- Configure at least one hot-relocation spare disk in each disk group. This will allow sufficient space for relocation in the event of a failure.

If you decide to disable hot-relocation, prevent `vxrelocd` from running after you load the VxVM software. See the section “Modifying the behavior of Hot-Relocation” in Chapter 9 of the *VERITAS Volume Manager Administrator's Guide* for details.

About placing disks in another disk group

To place disks in another disk group, use VEA or the `vxdiskadm` program after completing the `vxinstall` program. See the *VERITAS Volume Manager Administrator's Guide* for information on how to create other disk groups for your disks.

Protecting your system and data

A disk failure can cause loss of data on the failed disk and loss of access to your system. Loss of access is due to the failure of a key disk used for system operations. VxVM can protect your system from these problems.

To maintain system availability, data important to running and booting your system must be mirrored. The data must be preserved so it can be used in case of failure.

The following are suggestions for protecting your system and data:

- Place the disk containing the root file system (the root or boot disk) under VxVM control through encapsulation. Encapsulation converts the `root` and `swap` devices to volumes (`rootvol` and `swapvol`).
- Mirror the root disk so that an alternate root disk exists for booting purposes. By mirroring disks critical to booting, you ensure that no single disk failure leaves your system unbootable and unusable.
For maximum availability of the system, create mirrors for the `rootvol`, `swapvol`, `usr`, and `var` volumes. For more information, see the *Veritas Volume Manager Troubleshooting Guide*.
- Use mirroring to protect data against loss from a disk failure. To preserve data, create and use mirrored volumes that have at least two data plexes. The plexes must be on different disks. If a disk failure causes a plex to fail, the data in the mirrored volume still exists on the other disk.
- Leave the VxVM hot-relocation feature enabled to detect disk failures automatically. It will notify you of the nature of the failure, attempt to relocate any affected subdisks that are redundant, and initiate recovery procedures. Configure at least one hot-relocation spare disk in each disk group. This will allow sufficient space for relocation in the event of a failure. If the `root` disk is mirrored, hot-relocation can automatically create another mirror of the `root` disk if the original `root` disk fails. The `rootdg` must contain enough contiguous spare or free space for the volumes on the root disk (`rootvol` and `swapvol` volumes require contiguous disk space).
- Use the DRL feature to speed up recovery of mirrored volumes after a system crash. Make sure that each mirrored volume has at least one log subdisk.

Note: `rootvol`, `swapvol`, and `usr` volumes cannot be DRL volumes.

- Use logging to prevent corruption of recovery data in RAID-5 volumes. Make sure that each RAID-5 volume has at least one log plex.
- Perform regular backups to protect your data. Backups are necessary if all copies of a volume are lost or corrupted. Power surges can damage several (or all) disks on your system. Also, typing a command in error can remove critical files or damage a file system directly. Performing regular backups ensures that lost or corrupted data is available to be retrieved.

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 `/etc/vfstab` file.

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

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

Uninstalling the Veritas Storage Foundation software

This chapter covers uninstallation requirements and steps to uninstall the VERITAS software.

Note: The information in “[Uninstalling Veritas Storage Foundation](#)” on page 134 pertains to Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation for Sybase.

Topics covered in this chapter include:

- “[Uninstallation requirements](#)” on page 126
- “[Dropping the repository database for DB2 and Oracle](#)” on page 133
- “[Shutting down cluster operations](#)” on page 134
- “[Uninstalling Veritas Storage Foundation](#)” on page 134
- “[Uninstalling the Veritas Enterprise Administrator client](#)” on page 136
- “[Uninstalling language packages](#)” on page 136

Uninstallation requirements

Review the uninstallation requirements before removing the Veritas software.

Remote uninstallation

For information on remote uninstalls, refer to the *Veritas Storage Foundation and High Availability Solutions Getting Started Guide* that came with your software discs.

Veritas Volume Manager

This section describes the steps you need to take before removing VERITAS Volume Manager to preserve the contents of the volumes.

Caution: Failure to follow the preparations in this section might result in unexpected behavior.

Moving volumes from an encapsulated root disk

To uninstall VxVM if `rootvol`, `swapvol`, `usr`, or `var` is a volume under Volume Manager control

- 1 Ensure that the `rootvol`, `swapvol`, `usr`, and `var` volumes have only one associated plex each.

The plex must be contiguous, non-striped, non-spanned, and non-sparse. To obtain this information, enter the following:

```
# vxprint -ht rootvol swapvol usr var
```

If any of these volumes have more than one associated plex, remove the unnecessary plexes using the following command:

```
# vxplex -o rm dis plex_name
```

Note: It is important that you have created the plex designated for `rootvol` using `vxrootmir`, which guarantees that the underlying subdisks start on cylinder boundaries and that partitions are created for them.

- 2 Run the `vxunroot` command:

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

The `vxunroot` command changes the volume entries in `/etc/vfstab` to the underlying disk partitions for `rootvol`, `swapvol`, `usr`, and `var`. It also modifies `/etc/system` and prompts for a reboot so that disk partitions are mounted instead of volumes for `root`, `swap`, `usr`, and `var`.

- 3 Once you have changed the `root`, `swap`, `usr`, and `var` volumes, move all remaining volumes to disk partitions. You can do this using one of these procedures:
 - Back up the entire system to tape and then recover from tape.
 - Back up each file system individually and then recover them all after creating new file systems on disk partitions.
 - Move volumes incrementally to disk partitions.Continue to the next section if you choose this last option. Otherwise, proceed to “[Shutting down VERITAS Volume Manager](#)” on page 132.

Moving volumes to disk partitions

To move volumes incrementally to disk partitions

- 1 Evacuate disks using `vxdiskadm`, the VEA GUI, or the `vxevac` utility. 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 VERITAS Volume Manager control by entering:

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

```
# dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/dsk/c2t2d2s7
```

where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly created partition.
- 7 Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.

- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Stop and remove the volume from VERITAS Volume Manager using the commands.


```
# vxvol -g diskgroup stop volume_name
# vxedit -rf rm volume_name
```
- 10 Remove any free disks (those having no subdisks defined on them) by removing the volumes from VERITAS Volume Manager control. To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.


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

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


```
# vxdg rmdisk diskname
# vxdisk rm devname
```

 Use the free space created for adding the data from the next volume you want to remove.
- 11 After you successfully convert all volumes into disk partitions, reboot the system.
- 12 After the reboot, make sure none of the volumes are open by using the `vxprint` command.

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

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

Example

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

These are the contents of the disk group `voldg` before the data on `vol01` is copied to `disk3`.

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

dg voldg     default     default    115000
```

1017856044.1141.hostname.veritas.com

```
dm disk1      c1t12d0s2    sliced  2591    17900352 -
dm disk2      c1t14d0s2    sliced  2591    17899056 -
dm disk3      c1t3d0s2     sliced  2591    17899056 -

v  vol1       -             ENABLED ACTIVE  4196448 ROUND  -         fsgen
pl  pl1       vol1         ENABLED ACTIVE  4196448 CONCAT -         RW
sd  sd1       pl1          disk1     0        2098224 0        c1t12d0  ENA
sd  sd2       pl1          disk2     0        2098224 2098224  c1t14d0  ENA
```

Evacuate disk1 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk1 disk3
# vxprint -g voldg -ht
```

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

```
dg voldg     default      default    115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1      c1t12d0s2    sliced  2591    17900352 -
dm disk2      c1t14d0s2    sliced  2591    17899056 -
dm disk3      c1t3d0s2     sliced  2591    17899056 -

v  vol1       -             ENABLED ACTIVE  4196448 ROUND  -         fsgen
pl  pl1       vol1         ENABLED ACTIVE  4196448 CONCAT -         RW
sd  disk3-01  pl1          disk3     0        2098224 0        c1t3d0   ENA
sd  sd2       pl1          disk2     0        2098224 2098224  c1t14d0  ENA
```

Evacuate disk2 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk2 disk3
# vxprint -g voldg -ht
```

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

```
dg voldg     default      default    115000
```

1017856044.1141.hostname.veritas.com

```
dm disk1      c1t12d0s2    sliced  2591    17900352 -
dm disk2      c1t14d0s2    sliced  2591    17899056 -
dm disk3      c1t3d0s2     sliced  2591    17899056 -

v  vol1       -             ENABLED ACTIVE  4196448 ROUND  -         fsgen
pl  pl1       vol1         ENABLED ACTIVE  4196448 CONCAT -        RW
sd  disk3-01  pl1          disk3      0        2098224 0        c1t3d0   ENA
sd  disk3-02  pl1          disk3      2098224 2098224 2098224  c1t3d0   ENA
```

Remove the evacuated disks from VERITAS Volume Manager control.

```
# vxdisk -g voldg list
DEVICE      TYPE      DISK      GROUP     STATUS
c1t3d0s2    sliced   disk3     voldg     online
c1t12d0s2   sliced   disk1     voldg     online
c1t14d0s2   sliced   disk2     voldg     online

# vxvg rmdisk disk1
# vxvg rmdisk disk2
# vxdisk rm c1t12d0
# vxdisk rm c1t14d0
```

Verify that the evacuated disks have been removed from VERITAS Volume Manager control.

```
# vxdisk -g voldg list
DEVICE      TYPE      DISK      GROUP     STATUS
c1t3d0s2    sliced   disk3     voldg     online
```

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

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

Create a partition on free disk space of the same size as the volume. In this example, a 2G partition is created on disk1 (c1t12d0s1).

format

Searching for disks...done

AVAILABLE DISK SELECTIONS:

0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
 /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
1. c1t3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
2. c1t9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
3. c1t10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
4. c1t11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
5. c1t12d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@c,0

6. clt14d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@e,0
7. clt15d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@f,0

Specify disk (enter its number): **5**
 selecting clt12d0
 [disk formatted]

FORMAT MENU:

- disk - select a disk
- type - select (define) a disk type
- partition - select (define) a partition table
- current - describe the current disk
- format - format and analyze the disk
- repair - repair a defective sector
- label - write label to the disk
- analyze - surface analysis
- defect - defect list management
- backup - search for backup labels
- verify - read and display labels
- save - save new disk/partition definitions
- inquiry - show vendor, product and revision
- volname - set 8-character volume name
- !<cmd> - execute <cmd>, then return
- quit

format> **p**

PARTITION MENU:

- 0 - change `0' partition
- 1 - change `1' partition
- 2 - change `2' partition
- 3 - change `3' partition
- 4 - change `4' partition
- 5 - change `5' partition
- 6 - change `6' partition
- 7 - change `7' partition
- select - select a predefined table
- modify - modify a predefined partition table
- name - name the current table
- print - display the current table
- label - write partition map and label to the disk
- !<cmd> - execute <cmd>, then return
- quit

partition> **1**

Part	Tag	Flag	Cylinders	Size	Blocks
1	unassigned	wm	0	0	(0/0/0) 0

Enter partition id tag[unassigned]:

Enter partition permission flags[w]:

Enter new starting cyl[0]:

Enter partition size[0b, 0c, 0.00mb, 0.00gb]: **2.00gb**

partition> **1**

Ready to label disk, continue? **y**

```
partition> p
Current partition table (unnamed):
Total disk cylinders available: 13814 + 2 (reserved cylinders)
Part      Tag      Flag      Cylinders      Size      Blocks
  0 unassigned  wm        0              0          (0/0/0)      0
  1 unassigned  wm        0 - 3236      2.00GB     (3237/0/0)   4195152
partition> q
```

Copy the data on vol101 to the newly created disk partition.

```
# dd if=/dev/vx/dsk/voldg/vol101 of=/dev/dsk/c1t12d0s1
```

In the /etc/vfstab file, remove the following entry.

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

Replace it with an entry for the newly created partition.

```
/dev/dsk/c1t12d0s1 /dev/rdisk/c1t12d0s1 /vol101 vxfs 4 yes rw
```

Mount the disk partition.

```
# mount -F vxfs /dev/dsk/c1t12d0s1 /vol101
```

Remove vol101 from VERITAS Volume Manager.

```
# vxedit -rf rm /dev/vx/dsk/voldg/vol101
```

To complete the procedure, follow Steps 10 through 12 in the previous section, “[Moving volumes to disk partitions](#)” on page 127.

Shutting down VERITAS Volume Manager

To shut down VERITAS Volume Manager

- ◆ Enter the vxctl and vxiod commands as follows:

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

Veritas File System

The VRTSvxfs package cannot be removed if there are any mounted VxFS file systems or Storage Checkpoints. Unmount the VxFS file systems and Storage Checkpoints before uninstalling Veritas Storage Foundation. After you remove the VRTSvxfs package, VxFS file systems are not mountable or accessible until another VRTSvxfs package is installed.

To unmount a file system

- 1 Check if any VxFS file systems are mounted.

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

- 2 Unmount any file systems.

```
# umount special | mount_point
```

Specify the file system to be unmounted as a *mount_point* or *special* (the device on which the file system resides). See the `umount_vxfs(1M)` manual page for more information about this command and its available options.

Note: You can use the `-a` option to unmount all file systems except `/`, `/usr`, `/usr/kvm`, `/var`, `/proc`, `/dev/fd`, and `/tmp`.

To unmount a Storage Checkpoint

- 1 Check if any Storage Checkpoints are mounted.

```
# cat /etc/mnttab | grep vxfs
```
- 2 Unmount any Storage Checkpoints.

```
# umount /checkpoint_name
```

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
```

Note: 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.

Uninstalling Veritas Storage Foundation

If you need to uninstall the Veritas Storage Foundation software packages, use the uninstallation script.

The following procedures pertain to Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation for Sybase.

Note: If you installed the French or Chinese versions of Veritas Volume Manager, the uninstallation script does not remove the French or Chinese language packages.

To remove a Veritas Storage Foundation product

- 1 Comment out or remove any VERITAS File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.
- 2 Unmount all mount points for VxFS file systems.

```
# umount /mount_point
```

- 3 If the VERITAS Volume Manager package (VRTSvxxvm) is installed, read and follow the uninstallation procedures in the section titled “[Veritas Volume Manager](#)” on page 126.
- 4 Stop the VEA Service.

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 5 If you are uninstalling Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, stop the repository database and unmount the repository volume.

In a stand-alone configuration:

Stop the database repository:

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

Unmount the database repository:

```
# /opt/VRTSdbcom/config/sfua_rep_mount stop
```

In an HA configuration:

Stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 6 Move to the /opt/VRTS/install directory on your system and use the uninstallation script to remove the Veritas Storage Foundation product installed on your system. For example, to remove Veritas Storage Foundation:

```
# cd /opt/VRTS/install  
# ./uninstallsf
```

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

- 7 To verify the removal of the packages, use the `pkginfo` command.

```
# pkginfo | grep VRTS
```

Uninstalling 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 from a Solaris system other than the server

- 1 Stop the VEA Service.

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 2 Use the `pkgrm` command to remove the `VRTSobgui` software package.

```
# pkgrm VRTSobgui
```

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.

Uninstalling language packages

Language packages are uninstalled when you remove the English packages using the product installer menu or the uninstallation scripts provided by the software.

See [“Uninstalling Veritas Storage Foundation”](#) on page 134.

Note: The product installer and uninstallation scripts do not have an option to remove *ONLY* language packages.

Uninstalling language packages using the `pkgadd` command

If you would like to remove only the language packages, you can do so with the `pkgrm` command.

Note: If you use the product installer menu or the uninstallation script, you can remove the language packages along with the English packages.

To remove the language packages

- 1 Stop the VEA service on each system using the `vxsvcctl stop` command.

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 2 Use the `pkgrm` command to remove the appropriate packages.

```
# pkgrm package_name package_name ...
```

Table 6-4 Locations of language package information

For information on...	Go to...
Japanese language packages	Table B-2 on page 182
Chinese language packages	Table B-3 on page 185
French language packages	Table B-4 on page 185

Note: Some of the packages listed in the tables referenced above may not be installed on your system, depending on the actual installation.

Because the packages do not contain any dependencies, you can remove them in any order.

- 3 After removing the appropriate packages, restart the VEA service on each system using the `vxsvcctl start` command.

```
# /opt/VRTS/bin/vxsvcctl start
```


Installing, upgrading and uninstalling VxVM and VxFS

Topics covered in this chapter include:

- [“Installing Veritas Volume Manager”](#) on page 140
- [“Installing Veritas File System”](#) on page 145
- [“Upgrading Veritas Volume Manager”](#) on page 148
- [“Upgrading Veritas File System”](#) on page 166
- [“Uninstalling Veritas Volume Manager”](#) on page 168
- [“Uninstalling Veritas File System”](#) on page 169

Note: Hosts may be configured as managed hosts or as standalone hosts. A Management Server and Authentication Broker must have previously been set up if a managed host is required.

Any new licenses that are required should be obtained before attempting to install or upgrade the software.

Installing Veritas Volume Manager

If you purchased Veritas Volume Manager only, use the steps in this section to install the product. If you purchased one of the Storage Foundation products, see “[Installing a Veritas Storage Foundation product](#)” on page 31.

[Table 7-5](#) lists the various methods that are possible for installing and upgrading VxVM.

Table 7-5 Installation and upgrade methods

Method	Description
Installation using the product installer	See “ Installing VxVM using the product installer ” on page 140.
Installation using pkgadd and conversion to a managed host	See “ Installing VxVM using the pkgadd command ” on page 142.
Installation using JumpStart	See “ Installing VxVM using JumpStart ” on page 144.
Upgrade using the product installer	See “ Upgrading VxVM using the product installer ” on page 149. See “ Upgrading VxFS using the product installer ” on page 166.
Upgrade using pkgadd	See “ Upgrading VxVM using pkgadd ” on page 151.
Upgrade using scripts	See “ Upgrading VxVM using upgrade scripts ” on page 154. See “ Upgrading VxVM and the Solaris OS ” on page 154.
Upgrade using Live Upgrade	“ Upgrading VxVM and/or Solaris using Live Upgrade ” on page 158.

Installing VxVM using the product installer

To install VxVM

- 1 Mount the disc.
See “[Mounting the software disc](#)” on page 32..
- 2 Go to the top level directory of the disc:
`cd /cdrom/cdrom0`
- 3 Run the product installer:
`./installer`

The Selection Menu appears and prompts you to enter a letter indicating your selection. In the following example, the selection is “I” to select a product. The product being installed is Veritas Volume Manager.

- 4 You are then prompted to enter the system names. The installation script performs an initial system check, and confirms success by displaying information regarding the OS version, whether the VRTSvxvm package is installed, and whether the necessary patches are installed.
- 5 You will now be asked which set of packages should be installed.
 - 1) Required Veritas Volume Manager packages - 698MB required
 - 2) All Veritas Volume Manager packages - 838 MB required
 - 3) Storage Foundation Enterprise HA packages - 1192 MB required.
- 6 The packages will be installed .
- 7 The system will then begin license verification.
- 8 Enter the license key. For example:
`xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxx`
The output resembles the following:
`xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxx successfully registered on test1.`
`Do you want to enter another license key for test1? [y,n,q,?]`
`(n)`
- 9 Repeat [step 8](#) as prompted for any other systems. When this is complete, the following is displayed:
`VxVM licensing completed successfully.`
- 10 Follow instructions until the installation is complete. You are then asked if you would like to configure Veritas Volume Manager:
`Are you ready to configure VM [y, n, q]`
- 11 If you answer ‘y’, you will then be asked if you want to set up an enclosure-based naming scheme. You will then be asked a series of questions regarding the default disk group. This establishes the name of the default disk group. It does not create the disk group.
- 12 If you have a VVR license installed, the next phase concerns configuration of VVR:
`Do you want to change any of the VVR ports ... [y, n, q]`
- 13 You are now asked questions regarding the frequency of VVR statistics collection.
- 14 The next phase of the configuration procedure consists of setting up a centrally managed host:
`Enable Centralized Management? [y,n,q]`
- 15 Finally, if you selected centralized management, you will be asked a series of questions relating to host names and the VEA agent password.

- 16 When the procedure is complete, the system will ask you:
Do you want to start Veritas Volume Manager processes now?
[y, n, q]
- 17 You should see the message:
Startup completed successfully on all systems
- 18 You are now prompted to enter an encryption key.
- 19 The installation is now complete. Check the log file for information.

Installing VxVM using the pkgadd command

Note: The Veritas packages and patches are not compressed when you purchase Veritas Volume Manager through Sun Microsystems.

On Solaris 10, the packages must be installed while in the global zone.

This procedure describes how to install the software on a standalone host. The system can be converted to a managed host at a later stage.

To install VxVM

- 1 Mount the software disc. See “[Mounting the software disc](#)” on page 32.
- 2 Copy the packages to a location to which you can write and then uncompress and untar the packages:

```
# cp -R /cdrom/cdrom0/volume_manager/pkgs/* /tmp/pkgs
```
- 3 Uncompress and extract the packages by using the `gzcat` command:

```
# cd /tmp/pkgs  
# gzcat *.gz | tar xvf -
```
- 4 The `-a adminfile` option should be specified to `pkgadd`. This *adminfile* must be created in the current directory, and contain the following entries:

```
mail=  
instance=overwrite  
partial=nocheck  
runlevel=quit  
idepend=quit  
rdepend=nocheck  
space=quit  
setuid=nocheck  
conflict=nocheck  
action=nocheck  
basedir=default
```

5 Install the packages.

While installing the `VRTSobc33` package, enter `n` when prompted if the host will be centrally managed.

On Solaris 10, these packages must be installed while in the global zone. If a package's `pkginfo` file contains the variable `SUNW_PKG_ALLZONES` set not equal to `true`, the `-G` option should additionally be specified to the `pkgadd` command.

```
# pkgadd -a adminfile -d . VRTSvlic VRTSperl VRTSicsco \
VRTSspbx VRTSsmf
# pkgadd -a adminfile -d . VRTSat VRTSobc33 VRTSob \
VRTSobgui VRTSccg VRTSmh
# pkgadd -a adminfile -d . VRTSaa VRTSspt SYMClma \
VRTSvxvm VRTSdsa VRTSfspro
# pkgadd -a adminfile -d . VRTSvm doc VRTSvmman VRTSdcli \
VRTSalloc VRTSvmpro VRTSvsvc
# pkgadd -a adminfile -d . VRTSv did VRTSddlpr VRTSvrpro \
VRTSvcsvr VRTSspt VRTSvsvc
# pkgadd -a adminfile -d . VRTSjre VRTSjre15 VRTSweb \
VRTSvrw VRTSvrdoc
```

6 Start the VEA server:

```
# /opt/VRTSob/bin/vxsvcctl start
```

7 Register the agents:

```
# /opt/VRTSobc/pal33/install/cfgsecurity.sh -a StorageAgent
# /opt/VRTSobc/pal33/install/cfgsecurity.sh -a actionagent
# /opt/VRTSobc/pal33/install/cfgsecurity.sh -a gridnode
```

8 Run the `vxinstall` command to configure VxVM.

9 Start the agents:

```
# /opt/VRTSobc/pal33/bin/vxpalctl -a StorageAgent -c start
# /opt/VRTSobc/pal33/bin/vxpalctl -a actionagent -c start
# /opt/VRTSobc/pal33/bin/vxpalctl -a gridnode -c start
```

10 Configure the license manager agent:

```
# /opt/SYMClma/bin/lmautil --Config --SecurityEnabled 1 \
--RootBrokerHostname "1" --CollectorNodeUsername "2" \
--CollectorNodeUserDomainType "3" \
--CollectorNodeUserDomain "4"
```

To migrate a standalone host to a managed host, see [“Migrating a standalone host to a managed host”](#) on page 154.

Installing VxVM using JumpStart

These installation instructions using JumpStart assume a knowledge of the JumpStart procedure.

Note: Hosts may be configured as managed hosts or as standalone hosts. A Management Server and Authentication Broker must have previously been set up if a managed host is required.

Only fresh installs of VxVM are supported using JumpStart. Upgrading is not supported.

To install VxVM

- 1 Mount the software disc. See “[Mounting the software disc](#)” on page 32.
- 2 Copy the packages to a location to which you can write and then uncompress and untar the packages:

```
# cp -R /cdrom/cdrom0/volume_manager/pkgs/* \  
/$jumpstart/vm_pkg_dir
```

- 3 Uncompress and extract the packages by using the `gzcat` command:

```
# cd /$jumpstart/vm_pkg_dir  
# gzcat *.gz | tar xvf -
```

- 4 The `pkgadd` operations to install the packages must be coded in a script that can be used with the JumpStart server. The packages that are required for VxVM must be installed in the following order:

```
VRTSvlic VRTSperl VRTSicsco VRTSspb VRTSsmf VRTSat VRTSobc33  
VRTSob VRTSobgui
```

```
VRTSccg VRTSmh VRTSaa VRTSspt SYMClma VRTSvxvm VRTSdsa  
VRTSfspro VRTSvman
```

```
VRTSvmdoc VRTSdcli VRTSalloc VRTSvmpro VRTSsvvc VRTSvdid  
VRTSddlpr VRTSvrpro
```

```
VRTSvcsvr VRTSjre VRTSjre15 VRTSweb VRTSvrw VRTSvrdoc
```

The `VRTSdcli`, `VRTSjre`, `VRTSsvvc`, `VRTSvxvm` and `VRTSweb` packages require an empty response file, for example:

```
touch responsefile  
pkgadd -r responsefile $pkg_name
```

The `VRTSobc33` response file must contain the following settings:

```
VXPALSECURITY=n  
VXPALLOCALHOSTNAME=  
VXPALDOMAINCONTROLLER=  
VXPALAUTHENTICATIONBROKER=
```

The VRTSob response file must contain the following settings:

```
SECURITYADMINPASSWORD=  
ROOTAGENTPASSWORD=  
CONFIGURESECURITY=n
```

The VRTSvdid response file must contain the following setting:

```
OPT=/opt
```

The `-a adminfile` option should be specified to `pkgadd`. This *adminfile* must be created in the current directory, and contain the following entries:

```
mail=  
instance=overwrite  
partial=nocheck  
runlevel=quit  
idepend=quit  
rdepend=nocheck  
space=quit  
setuid=nocheck  
conflict=nocheck  
action=nocheck  
basedir=default
```

On Solaris 10, the packages must be installed while in the global zone. If a package's `pkginfo` file contains the variable `SUNW_PKG_ALLZONES` set not equal to `true`, the `-G` option should additionally be specified to the `pkgadd` command.

- 5 JumpStart reboots the system after the packages have been installed. If it does not, then after the packages have been installed by JumpStart, manually reboot the system:
- 6 To configure Volume Manager, run the following script from the mounted software disc:

```
# shutdown -i6 -g0 -y  
  
# cd /cdrom/cdrom0/volume_manager  
# ./installvm `hostname` -configure
```

Installing Veritas File System

If you purchased Veritas File System only, use the steps in this section to install the product. If you purchased one of the Storage Foundation products, see [“Installing a Veritas Storage Foundation product”](#) on page 31.

Installing VxFS using the product installer

To install VxFS

- 1 Mount the disc.
See [“Mounting the software disc”](#) on page 32..
- 2 Go to the top level directory of the disc:

```
# cd /cdrom/cdrom0
```
- 3 Run the product installer:

```
# ./installer
```

The Selection Menu appears and prompts you to enter a letter indicating your selection. In the following example, the selection is “**I**” to select a product. The product being installed is Veritas Volume Manager.
- 4 You are then prompted to enter the system names. The installation script performs an initial system check, and confirms success by displaying information regarding the OS version, whether the `VRTSvxfs` package is installed, and whether the necessary patches are installed.
- 5 You will now be asked which set of packages should be installed.
1) Required Veritas File System packages - 437 MB required
2) All Veritas File System packages - 441 MB required
- 6 The packages will be installed .
- 7 The system will then begin license verification.
- 8 Enter the license key. For example:

```
xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxx
```

The output resembles the following:

```
xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxx successfully registerd on test1.  
Do you want to enter another license key for test1? [y,n,q,?]  
(n)
```
- 9 Repeat [step 8](#) as prompted for any other systems. When this is complete, the following is displayed:

```
VxVM licensing completed successfully.
```
- 10 Follow instructions until the installation is complete. You are then asked if you would like to configure Veritas Volume Manager:

```
Are you ready to configure VM [y, n, q]
```
- 11 The next phase of the configuration procedure consists of setting up a centrally managed host:

```
Enable Centralized Management? [y,n,q]
```
- 12 Finally, if you selected centralized management, you will be asked a series of questions relating to host names and the VEA agent password.

13 When the procedure is complete, the system will ask you:

```
Do you want to start Veritas File System processes now? [y,
n, q]
```

14 You should see the message:

```
Startup completed successfully on all systems
Installation log files, summary file, and response file are
saved at:
```

```
/opt/VRTS/install/logs/installfs-J8aIMn
```

```
Reboot all systems on which VxFS was installed or upgraded.
shutdown -y -i6 -g0
```

```
See the Veritas File System Administrators Guide for
information on using VxFS.
```

Installing VxFS using the pkgadd command

To install VxFS

- 1 The VxFS packages are compressed using GNU compression, so you will need to uncompress them using the `gunzip` command. First, copy the packages from the directory containing the packages to a location to which you can write and then uncompress and untar the packages. If the current directory is your selected location, create a directory `pkgs`, which will be the target path, and enter the commands:

```
# cp -R /cdrom/cdrom0/file_system/pkgs/* pkgs
```

- 2 Uncompress the packages using the `gunzip` command, and place the uncompressed packages in the `pkgs` directory that you created:

```
# /cdrom/cdrom0/gnu/gunzip pkgs/*.gz
```

- 3 Change to the `pkgs` directory that now contains the VxFS packages:

```
# cd pkgs
```

Use `tar` to extract the packages. You should extract each package individually using a command such as:

```
# tar xvf VRTSvxfs.tar
```

Repeat the command for each package.

- 4 Install the packages:

```
# pkgadd -d . VRTSvlic VRTSob VRTSobgui VRTSvxfs VRTStep \
VRTSap VRTSfspro VRTSfsman VRTSfsdoc VRTSfssdk VRTSfsmnd
```

- 5 If you are installing VxFS on a file system with disk layout Version 3 or prior, continue by upgrading the disc layouts.

See [“Upgrading VxFS disk layout versions”](#) on page 81.

If you performed a fresh installation, verify your installation.

See [“Checking VxFS installation”](#) on page 43.

Upgrading Veritas Volume Manager

If you purchased Veritas Volume Manager only, use the steps in the following section to upgrade VxVM. If you purchased Storage Foundation, use the Storage Foundation upgrade procedure.

This chapter describes how to upgrade Veritas Volume Manager (VxVM). There are several ways to upgrade:

- Upgrading VxVM Only; see [“Upgrading VxVM only”](#) on page 148
- Upgrading VxVM and the Solaris operating system; see [“Upgrading VxVM and the Solaris OS”](#) on page 154.
- Upgrading the Solaris operating system only; see [“Upgrading the Solaris OS only”](#) on page 156.

Within these scenarios, there are other different upgrade paths depending on whether you have an encapsulated disk.

Note: If you are planning to use VxVM in a cluster environment, follow the instructions in [“Enabling cluster support in VxVM \(optional\)”](#) on page 117.

Upgrading VxVM only

This section explains how to upgrade VxVM. Refer to the [“VxVM and Solaris release support matrix”](#) on page 78 to determine if upgrading the operating system is required.

If you are planning to use the `installvm` script to upgrade VxVM, go to [“Upgrading VxVM using the product installer”](#) on page 149.

There are several options to choose from if you want to use the `pkgadd` method to upgrade VxVM:

- An upgrade with one reboot.
See [“Upgrading VxVM using pkgadd”](#) on page 151.
- An upgrade method which may require more than one reboot.
See [“Upgrading VxVM using upgrade scripts”](#) on page 154.
- An upgrade with one reboot, which requires an alternate boot disk.
See [“VxVM Live Upgrade commands and usage”](#) on page 159.

Upgrading VxVM using the product installer

You can use the product installer (`installvm` script) to upgrade VxVM with an encapsulated or unencapsulated root disk.

To upgrade VxVM on a remote host, `rsh` or `ssh` must be set up. See “[Completing the installation procedure](#)” on page 39 for more information.

Upgrading VxVM using the product installer includes the following general steps:

- Run `installvm` to load the packages
- Do a reboot
- Run `installvm` to configure the 5.0 items

Note: If you are performing a multihost installation, note that the `installvm` script does not support a mixture of new installations and upgrades. You should either install on all hosts or upgrade all hosts. In addition, every host to be upgraded in a multihost installation must be running the same older version of VxVM.

If you use this `installvm` script to upgrade, do not use the `upgrade_start` and `upgrade_finish` scripts.

To upgrade using the `installvm` script

- 1 Mount the software disc; see “[Mounting the software disc](#)” on page 32 for instructions.
- 2 Remove the older VRTS packages using the `pkgrm` command:

```
# pkgrm VRTSvras VRTSlic VRTSvmsa
```

Note: Some of the above packages may not exist on your system.

- 3 Insert the DVD in the DVD-ROM drive.
Once the DVD is inserted, the Solaris volume management software automatically mounts it. Change directory:

```
# cd /cdrom/cdrom0/volume_manager
```
- 4 To install the software on a single host, run the following command:

```
# ./installvm
```

Alternatively, copy the VxVM packages and patches to a network-shared file system. You can then install the software on any systems that have mounted this file system.

```
# cd /cdrom/cdrom0  
# cp -r * $NFS_FS
```

```
# cd volume_manager
# ./installvm -pkgpath $NFS_FS/volume_manager/pkgs
```

where the NFS-mountable file system is *\$NFS_FS*.

If a patch is available, the `installvm` command must also include the patch path as follows:

```
# ./installvm -pkgpath $NFS_FS/volume_manager/pkgs \
-patchpath $NFS_FS/volume_manager/patches
```

- 5 At the prompt, enter the name of the system or systems on which you want to install VxVM:

```
Enter the system names separated by spaces on which to
install VxVM:
```

The `installvm` script will carry out an initial system check, and will confirm success by displaying information regarding the OS version, whether the `VRTSVxvm` package is installed and whether the necessary patches are installed. The initial system check might tell you that you need to obtain Solaris OS patches. If this happens, you need to see “Solaris Patch Requirements” in the *Veritas Storage Foundation Release Notes*.

When the system check is complete, the `installvm` script displays this message:

```
Initial system check completed successfully.
Press Return to continue.
```

- 6 You will now be asked which set of packages should be installed.
 - 1) Required Veritas Volume Manager packages - 698MB required
 - 2) All Veritas Volume Manager packages - 838 MB required
 - 3) Storage Foundation Enterprise HA packages - 1192 MB required.

- 7 Select the packages to be installed.

- 8 After displaying a list of packages as they are copied to the system, confirm that you want to proceed with the upgrade:

```
Press [return] to continue:
installvm is now ready to upgrade VxVM
Are you sure you want to upgrade VxVM? [y,n,q]
```

- 9 Existing VxVM packages will now be removed prior to replacement by the VxVM 5.0 packages.
- 10 The VxVM packages are then installed. You will see a listing of the packages displayed.
- 11 VxVM is now upgraded, and you should see a message stating that VxVM 5.0 has been successfully installed. You will see a message advising you to reboot.

Reboot your system by entering:

```
# /usr/sbin/shutdown -y -i6 -g0
```

- 12 To take advantage of the new features in this release, upgrade to the latest disk group version. See [“Upgrading CVM protocol and disk group version”](#)

on page 165. To confirm that the VxVM upgrade has been successful, see “[Verifying the Veritas Storage Foundation upgrade](#)” on page 95.

- 13 Now configure VxVM as described in “[Configuring VxVM using the installvm script.](#)”

Configuring VxVM using the installvm script

To configure VxVM after installing the packages

- 1 Enter the following command

```
# ./installvm -configure
```

The script runs an initial system check, and will tell you that you cannot configure already configured features such as enclosure-based naming and default disk groups. However, it will lead you through the process of configuring 5.0 features, such as centralized management. For information about centralized management, refer to the *Veritas Storage Foundation Management Installation Guide*.
- 2 The next phase of the configuration procedure consists of setting up a centrally managed host:

```
Enable Centralized Management? [y,n,q]
```
- 3 Finally, if you selected centralized management, you will be asked a series of questions relating to host names. You will also be asked the VEA agent password.
- 4 You should see the following message:

```
Startup completed successfully on all systems
```
- 5 You are now prompted to enter an encryption key.

Upgrading VxVM using pkgadd

This section describes the procedure for upgrading using the `pkgadd` command. This works with both an encapsulated and unencapsulated root disk. The procedure requires only one reboot, but it overwrites the new VxVM package on the existing VxVM package.

Note: This procedure describes how to upgrade the software on a standalone host. The system can be converted to a managed host at a later stage.

To upgrade using pkgadd

- 1 If you have not already obtained and installed a VxVM 5.0 license key, do so now. See “[Symantec product licensing](#)” on page 15 for details.

- 2 Bring the system down to single-user mode using the following command:

```
# init s
```
- 3 If /opt is configured as a separate file system, ensure that it is mounted.
- 4 Stop all applications that are using VM.
- 5 Stop the VEA server if it is running.
 For release 3.5 MP4, use the following command:

```
# /opt/VRTSob/bin/vxsvc -k
```

 For release 4.0 and 4.1, use the following command:

```
# /opt/VRTSob/bin/vxsvcctl stop
```
- 6 Remove all VxVM packages except for VRTSVXVM. The packages to be removed depend upon the product release.
 For release 3.5 MP4, use the following commands:

```
# pkgrm VRTSvlic VRTSvman VRTSvmdoc VRTSvmpo
# pkgrm VRTSfspro VRTSobgui VRTSob
```

 For releases 4.0 and 4.1, use the following commands:

```
# pkgrm VRTSvlic VRTSvmdoc VRTSvman VRTScpi
# pkgrm VRTStep VRTSap VRTSvrdoc VRTSvrw
# pkgrm VRTSweb VRTSjre VRTSvcsvr VRTSvrpro
# pkgrm VRTSfspro VRTSalloc
# pkgrm VRTSvmpo VRTSddlpr VRTSjre VRTSperl
# pkgrm VRTSobgui VRTSob
```
- 7 Mount the software disc. See “[Mounting the software disc](#)” on page 32.
- 8 Copy the packages to a location `t` which you can write and then `uncompress` and `untar` the packages:

```
# cp -R /cdrom/cdrom0/volume_manager/pkgs/* /mytmp/pkgs
```
- 9 Uncompress and extract the packages using the `gzcat` command:

```
# cd /mytmp/pkgs
# gzcat *.gz | tar xvf -
```
- 10 The `-a adminfile` option should be specified to `pkgadd`. This `adminfile` must be created in the current directory, and contain the following entries:

```
mail=
instance=overwrite
partial=nocheck
runlevel=quit
idepend=quit
rdepend=nocheck
space=quit
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=default
```

11 Install the packages.

While installing the `VRTSobc33` package, enter `n` when prompted if the host will be centrally managed.

On Solaris 10, these packages must be installed while in the global zone. If a package's `pkginfo` file contains the variable `SUNW_PKG_ALLZONES` set not equal to `true`, the `-G` option should additionally be specified to the `pkgadd` command.

```
# pkgadd -a adminfile -d . VRTSvlic VRTSperl VRTSicsco \
VRTSspbx VRTSsmf
# pkgadd -a adminfile -d . VRTSat VRTSobc33 VRTSob \
VRTSobgui VRTSccg VRTSmh
# pkgadd -a adminfile -d . VRTSaa VRTSspt SYMClma \
VRTSvxvm VRTSdsa VRTSfspro
# pkgadd -a adminfile -d . VRTSvmdoc VRTSvmman VRTSdcli \
VRTSalloc VRTSvmpro VRTSvsvc
# pkgadd -a adminfile -d . VRTSvdid VRTSddlpr VRTSvrpro \
VRTSvcsvr VRTSspt VRTSvsvc
# pkgadd -a adminfile -d . VRTSjre VRTSjre15 VRTSweb \
VRTSvrw VRTSvrdoc
```

12 Reboot the host:

```
# shutdown -i6 -go -y
```

13 Start the VEA server:

```
# /opt/VRTSob/bin/vxsvctrl start
```

14 Register the agents:

```
# /opt/VRTSobc/pal33/install/cfgsecurity.sh -a StorageAgent
# /opt/VRTSobc/pal33/install/cfgsecurity.sh -a actionagent
# /opt/VRTSobc/pal33/install/cfgsecurity.sh -a gridnode
```

15 Start the agents:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c start
# /opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c start
# /opt/VRTSobc/pal33/bin/vxpalctrl -a gridnode -c start
```

16 Configure the license manager agent:

```
# /opt/SYMClma/bin/lmautil --Config --SecurityEnabled 1 \
--RootBrokerHostname "1" --CollectorNodeUsername "2" \
--CollectorNodeUserDomainType "3" \
--CollectorNodeUserDomain "4"
```

After you have upgraded VxVM using `pkgadd`, you should configure VxVM. See [“Configuring VxVM using the installvm script”](#) on page 151.

To migrate a standalone host to a managed host, see [“Migrating a standalone host to a managed host”](#) on page 154.

Migrating a standalone host to a managed host

To migrate a standalone host to a managed host

- ◆ Enter the following command:

```
# /opt/VRTSmh/bin/migrate central_mode \  
-d central_server_name -p vea_agent_password
```

Upgrading VxVM using upgrade scripts

The upgrade scripts consist of `upgrade_start` and `upgrade_finish`, which are provided on the installation media. Use these scripts when you upgrade either of the following:

- Solaris and VxVM
- Solaris only on the boot disk

The `upgrade_start` script prepares the system for a Solaris upgrade, including saving the VxVM configuration data, and converting an encapsulated boot disk to a partition format.

The `upgrade_finish` script restores the VxVM configuration, including re-encapsulating the boot disk, if needed.

For a general VxVM upgrade, use the `installvm` script.

See “[Upgrading VxVM using the product installer](#)” on page 149.

Upgrading VxVM and the Solaris OS

If you are upgrading from VxVM 3.5 MP4 and the Solaris operating system go to “[Upgrading VxVM and the Solaris OS](#)” on page 154.

If you are upgrading only the Solaris operating system, go to “[Upgrading the Solaris OS only](#)” on page 156.

This section explains how to upgrade both the Solaris operating system and VxVM. (Refer to the “[VxVM and Solaris operating system upgrade paths](#)” on page 79 to determine if an operating system upgrade is required.) If VxVM 3.5 MP4 is running on Solaris 7, you cannot upgrade directly to Solaris 10.

This upgrade procedure covers both the possible scenarios, which are:

- Upgrading from VxVM 3.5 MP4, and the root disk is encapsulated.
- Upgrading from VxVM 3.5 MP4, and the root disk is not encapsulated.

To upgrade both VxVM and the Solaris OS

- 1 If you have not already obtained and installed a VxVM 5.0 license key, do so now. See “[Symantec product licensing](#)” on page 15 for details.

Note: The directory `/opt` must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by `upgrade_start` are unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in `/opt` are not installed.

- 2 Load and mount the DVD by starting the `volmgt` daemon.

```
# /etc/init.d/volmgt start
```
- 3 Run the `upgrade_start` with the `-check` argument to detect any problems that exist which could prevent a successful upgrade. (Be sure to run the `upgrade_start` and `upgrade_finish` scripts that are supplied with the VxVM 5.0 release.) If this command reports success, you can proceed with running the `upgrade_start` script, but if it reports errors, correct the problem(s) and rerun `upgrade_start -check`.

```
# cd /cdrom/cdrom0/volume_manager/scripts
# ./upgrade_start -check
```
- 4 Run the `upgrade_start` script so that the system can come up with partitions. The `upgrade_start` script searches for volumes containing file systems, and if any are found, converts them to partitions:

```
# ./upgrade_start
```
- 5 Reboot to single-user mode.

```
# reboot -- -s
```
- 6 If needed, mount `/opt` manually before you run `pkgrm` to remove VxVM packages.
- 7 If VxVM 3.5 MP4 was previously installed, stop any VxVM processes. Then remove the old packages (including `VRTSVxvm`), as follows:

```
# pkgrm VRTSvmmman VRTSvmdoc VRTSvmpo VRTSfspro VRTSobgui
# pkgrm VRTSvxvm VRTSob VRTSvlic
```
- 8 If VxVM 4.0 or 4.1 was previously installed, remove the old packages, as follows:

```
# cd /opt/VRTS/install
# /opt/VRTSob/bin/vxsvctrl stop
# ./uninstallvm `hostname`
# ./uninstallinfr
```
- 9 Upgrade the operating system to Solaris 8, 9 or 10.

Note: Refer to the Solaris installation documentation for instructions on how to upgrade the Solaris operating system. After installing the Solaris operating system, install the required patches that are listed in “Solaris Patch Requirements” in the *Veritas Storage Foundation Release Notes*.

- 10 Once the system is up after upgrading the Solaris operating system, ensure that `/opt` is mounted. Then load and mount the software disc by starting the `volmgt` daemon.

```
# /etc/init.d/volmgt start
```

- 11 Install the VxVM packages from the VM 5.0 distribution media:

```
# cd /cdrom/cdrom0/volume_manager
# ./installvm `hostname` -installonly
```

- 12 Complete the upgrade:

```
# cd /cdrom/cdrom0/volume_manager/scripts
# ./upgrade_finish
```

- 13 Perform a reconfiguration reboot:

```
# reboot -- -r
```

- 14 Configure the VxVM packages from the VM 5.0 distribution media:

```
# cd /cdrom/cdrom0/volume_manager
# ./installvm `hostname` -configure
```

- 15 To take advantage of the new features in this release, you should upgrade to the latest disk group version. See [“Upgrading CVM protocol and disk group version”](#) on page 165. To confirm that the VxVM upgrade has been successful, see [“Verifying the Veritas Storage Foundation upgrade”](#) on page 95.

Upgrading the Solaris OS only

If you are running VxVM 5.0 with an earlier release of the Solaris operating system, you can upgrade the Solaris operating system using the following procedure.

Caution: You should only use this procedure to upgrade the Solaris operating system if you are running VxVM 5.0.

Note: The directory `/opt` must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by the `upgrade_start` are unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

To upgrade the Solaris operating system only

- 1 Bring the system down to single-user mode using the following command:

```
# init s
```

Note: You must mount `/opt` manually if `/opt` is on its own partition.

- 2 Load and mount the software disc from the currently installed version of VxVM by starting the `volmgt` daemon.

```
# /etc/init.d/volmgt start
```
- 3 Change directory:

```
# cd /cdrom/cdrom0/volume_manager/scripts
```
- 4 Run the `upgrade_start` with the `-check` argument to detect any problems that exist which could prevent a successful upgrade. Run the `upgrade_start` and `upgrade_finish` scripts that were supplied with the currently installed VxVM release. If this command reports success, you can proceed with running the `upgrade_start` script, but if it reports errors, correct the problem(s) and rerun `upgrade_start -check`.

```
# ./upgrade_start -check
```
- 5 Run the `upgrade_start` script so that the system can come up with partitions. Run the `upgrade_start` and `upgrade_finish` scripts that were supplied with the currently installed VxVM release. The `upgrade_start` script searches for volumes containing file systems, and if any are found, converts them to partitions:

```
# ./upgrade_start
```
- 6 Bring the system down to run level 0.

```
# init 0
```
- 7 Upgrade the operating system to Solaris 8, 9 or 10.

Note: You should boot up the system from run level 0 depending on the Solaris upgrade procedure that you want to follow. Refer to the Solaris installation documentation for instructions on how to upgrade the Solaris operating system. After installing the Solaris operating system, install the required patches that are listed in “Solaris Patch Requirements” in the *Veritas Storage Foundation Release Notes*.

- 8 Once the system is up with the upgraded Solaris operating system, bring the system down to single-user mode by entering:

```
# init s
```

- 9 Ensure that /opt is mounted.

- 10 Load and mount the software disc from the currently installed version of VxVM by starting the volmgt daemon.

```
# /etc/init.d/volmgt start
```

- 11 If upgrading to Solaris 10, reinstall the existing VRTSvxvm package from the currently installed version of VxVM:

```
# pkgrm VRTSvxvm
# pkgadd VRTSvxvm
```

If required, you can also apply these commands to the VRTSvxfs package. This reinstallation is necessary to install Solaris 10 SMF support for VxVM.

- 12 Change directory:

```
# cd /cdrom/cdrom0/volume_manager/scripts
```

- 13 Complete the upgrade from the software disc from the currently installed version of VxVM by entering:

```
# ./upgrade_finish
```

- 14 Perform the reconfiguration reboot:

```
# reboot -- -r
```

Upgrading VxVM and/or Solaris using Live Upgrade

Solaris Live Upgrade is the feature that performs an operating system upgrade with no downtime. The upgrade is done on an alternate disk using the current boot environment (BE). After the Live Upgrade, the alternate disk has all the information from the current BE and just one reboot will bring the system up on the alternate disk with the upgraded Solaris Version.

Volume Manager uses Solaris Live Upgrade to upgrade VxVM software in a live environment.

The advantages of using Live Upgrade are:

- Less system downtime - just one reboot will bring the system up and running on the new Solaris version

- You can revert to the previous OS version at any time
- Alternate root disk partitions can be resized to grow or shrink
- If you do not upgrade the disk group version on the alternate boot environment, you can create more than one boot environment with different versions of VxVM/Solaris and easily switch between them with just one reboot

VxVM Live Upgrade requires an alternate disk to upgrade. This disk can be a mirrored root disk or an independent unused disk which supports booting. VxVM Live Upgrade is mainly done using the Solaris Live Upgrade commands.

Upgrading VxVM or Solaris OS or Both Using Live Upgrade

VxVM Live Upgrade supports the following:

- Upgrading VxVM only (both encapsulated and unencapsulated root disk)
- Upgrading Solaris OS only (both encapsulated and unencapsulated root disk)
- Upgrading both Solaris OS and VxVM (both encapsulated and unencapsulated root disk)
- Installing any VxVM or Solaris OS patches

Live Upgrade requires an alternate disk to perform the upgrade. An alternate root disk can be a mirrored root disk (chosen by default), or you can specify an alternate disk. The size of the alternate disk should be greater than or equal to the size of the root disk. This procedure will upgrade all the mounted partitions on the root disk with entry in `/etc/vfstab`. Unmounted partitions and raw volumes will not be migrated to alternate disk.

Installing Live Upgrade on the current root disk

To begin upgrading the system you must first install the Solaris Live Upgrade software on the system. This Solaris Live Upgrade package is available on the latest Solaris software disc. Load the latest Solaris Software 2 of 2 disc. Run the command `liveupgrade20` in the directory `/cdrom/cdrom0/Solaris10/Tools/Installer`. This installs the Solaris Live Upgrade packages on the system. Read the *Solaris Installation Guide* for the procedure to install the latest Solaris Live Upgrade package.

VxVM Live Upgrade commands and usage

Volume Manager Live Upgrade uses the two commands `vxlustart` and `vxlufinish`. These commands are on the Volume Manager 5.0 software disc. The command `vxlustart` configures the machine (like virtual unencapsulation,

setting up alternate disk, and so on) and performs the Solaris Live Upgrade. The command `vxlufinish` completes the upgrade process by encapsulating the alternate root disk, if required. The `vxlustart` command can upgrade the Solaris operating system either from CD-ROM or from a network path. Copy the `vxlustart` and `vxlufinish` commands to the local directory so that the CD-ROM can be used to load the Solaris installation discs if you are upgrading the operating system. The usage of these commands is described below.

Using vxlustart

The `vxlustart` command and its options are:

```
vxlustart [-DfmVv] [-u {5.8|5.9|5.10}] [-d diskname] \  
[-g diskgroup] [-F filesystem] \  
[-s path_to_solaris_installation_image]  
  
vxlustart [-DfmUVv] [-u {5.8|5.9|5.10}] [-d diskname]  
  
vxlustart [-rv] [-u {5.8|5.9|5.10}]
```

where the option flags are:

- D This is for debugging. With this option `ksh` turns on the `-x` option to print every line it executes.
- d Use this option to specify the alternate disk name that is to be used as the new Solaris root disk. If the option is not specified, `vxlustart` prompts for the mirror root disk to be used.
- F Specify the file system type for the system volumes. The default file system type is `ufs`.
- f This option forces VTOC creation, if the partitions are not cylinder aligned.
- g This specifies the disk group where the root disk resides. This option is useful only if `vxldg bootdg` fails.
- m If this option is specified, the command assumes the VTOC is created manually. This is helpful if you want to increase the size of alternate root disk partition size. If this is not specified, the VTOC of the alternate root disk is created very similar to the current root disk.
- r Use this option to remount the alternate root disk in case the system was rebooted, or crashed after running `vxlustart` command and before completing `vxlufinish`.
- s Use this option to specify the path to the new Solaris image. This path must be network/directory path which has the complete Solaris image (one like in JumpStart image directory). If this option is not specified, the script assumes that the upgrade is from discs and the script will prompt for loading a disc. After loading the disc, the path to the image must be specified (for example, `/cdrom/cdrom0/s0`).

- U Use this option to upgrade VxVM only. The command does not prompt for Solaris discs.
- u Option to specify the Solaris version to be upgraded. In case of upgrading only VxVM, the version should be the current Solaris version.
- V Verbose option, which will print the list of commands executed by `vxlustart` without executing them. This option can be used as a preliminary check for the `vxlustart` command. This may not find all possible errors.
- v Verbose option to print the list of commands executed by `vxlustart`.

Using `vxlufinish`

The `vxlufinish` command and its options are:

```
vxlufinish [-fDv] [-u {5.8|5.9|5.10}]
```

where the option flags are:

- D This is for debugging. With this option `ksh` turns on the `-x` option to print every line it executes.
- f This command option forces the upgrade to complete in case Volume Manager is not upgraded. Before using this option make sure the Volume Manager drivers in the alternate disks are compatible to the upgraded Solaris operating system. Otherwise Volume Manager may not perform optimally.
- u Specify the upgraded Solaris version. This Solaris version must be the same as specified in the `vxlustart` command.
- v This is a verbose option to print the list of commands executed by `vxlustart`.

Beginning the Live Upgrade

To begin the Live Upgrade

- 1 If you have not already obtained and installed a VxVM 5.0 license key, do so now. See “[Symantec product licensing](#)” on page 15 for details.
- 2 Load and mount the installation media by starting the `volmgt` daemon.

```
# /etc/init.d/volmgt start
```
- 3 The two commands which are involved in Live Upgrade are `vxlustart` and `vxlufinish`. These commands are on the software disc. The `vxlustart` command configures the machine (like unencapsulation, setting up alternate disk, and so on) and performs the upgrade. The `vxlufinish` command completes the upgrade process by encapsulating the root disk if required. The `vxlustart` command can upgrade the Solaris operating

system either from the software discs or from a network directory path. Copy these two commands to the local directory so that the DVD-ROM drive can be used to load the Solaris installation discs.

```
# cp /cdrom/cdrom0/volume_manager/scripts/vxlustart .
# cp /cdrom/cdrom0/volume_manager/scripts/vxlufinish .
```

- 4 Run `vxlustart` with the `-v` option to detect any problems that might prevent a successful upgrade. If this command reports success, proceed with running the `vxlustart` script. If it reports errors, correct the problem, and run `vxlustart -v` again. Note that this option does not catch failures that are reported by Solaris Live Upgrade commands.

For example:

- If you want to upgrade from Solaris 9 to Solaris 10, run the following command to check whether the upgrade will run successfully.

```
# vxlustart -u 5.10 -v -d c#t#d# -s path_to_image_or_disc
```
 - If you want to upgrade just VxVM on Solaris 9, run:

```
# vxlustart -u 5.9 -v -d c#t#d# -U
```
- 5 Now upgrade Solaris from the software discs or a network path:
 - “Upgrading Solaris from the software discs”
 - “Upgrading Solaris from network directory path” on page 162.

Upgrading Solaris from the software discs

Load the Solaris software disc 1 of 1. Run the `vxlustart` command to upgrade to the version which is on the disc.

For example: to upgrade from Solaris 9 to Solaris 10, use the following command

```
# vxlustart -u 5.10 -d c#t#d# -s /cdrom/sol_10_404_sparc/s0
```

Once all the packages are installed from disc 1, the command prompts for the second disc. Load the second disc and specify the disc path at the prompt. For example: `/cdrom/cdrom0`. After you have installed the second disc, install the language disc if required. Enter `NONE` after installing all the required discs. The `-s` option is optional. If `-s` is not specified, the command will prompt to load the discs starting with the first.

Upgrading Solaris from network directory path

Upload the disc images to a network directory. Once completed, all the packages are available in one location so path to the directory will upgrade the complete Solaris operating system. To upgrade from Solaris 9 to Solaris 10, use the following command:

```
# vxlustart -u 5.10 -d c#t#d# -s \  
/network_directory_path/jumpstart/solaris2.10
```

This command upgrades the packages and activates the boot environment.

Completing the Solaris upgrade

To complete the Solaris upgrade using Live Upgrade

- 1 If upgrading to Solaris 10, reinstall the existing `VRTSvxvm` package from the currently installed version of VxVM:

```
# pkgrm VRTSvxvm
# pkgadd VRTSvxvm
```

If required, you can also apply these commands to the `VRTSvxfs` package. This reinstallation is necessary to install Solaris 10 SMF support for VxVM.

- 2 Use the `vxlufinish` command to complete the Live Upgrade process. Use the following command if the Solaris operating system is upgraded from Solaris 9 to 10 and Volume Manager is installed.

```
# vxlufinish -u 5.10
```

The upgrade process is now complete.

- 3 The alternate disk is mounted on the current root disk as `/altroot.5.OS_VERSION`. For example, if the alternate root disk is upgraded to Solaris 10 with two system partitions, `/` and `/usr`, these are mounted as `/altroot.5.10` and `/altroot.5.10/usr`. These two directories must be remounted manually in case the system crashes, or if the partition is accidentally unmounted before going to the next step.
- 4 Shut down the machine using `shutdown` or `init 6` to reboot the machine on the alternate root disk.

```
# init 6
```

Note: Don't use the `reboot` command to reboot the machine on an alternate root disk. If the alternate root disk has a problem rebooting or any upgrade issue, reboot the machine from the older disk.

Upgrading VxVM

To upgrade VxVM using Live Upgrade

- 1 If only Volume Manager needs to be upgraded, specify the `-U` option and the current Solaris version to the command. For example, if you want to upgrade only Volume Manager on Solaris 8, use the following command,

```
# vxlustart -u 5.8 -s path_to_image_or_disc -d c#t#d# -U
```

- 2 Load and mount the software disc by starting the `volmgt` daemon.

```
# /etc/init.d/volmgt start
```

- 3 If `/opt` is configured a separate file system, ensure that it is mounted.

- 4 If VxVM 3.5 MP4 was previously installed, remove the old VxVM packages on the alternate root by entering:

```
# rm /altroot.5.X/var/vx/isis/vxisis.lock
# /opt/VRTSob/bin/vxsvc -k
# pkgrm -R /altroot.5.X VRTSvmman VRTSvmdoc VRTSvmpro
# pkgrm -R /altroot.5.X VRTSfspro VRTSvxvm VRTSobgui \
  VRTSob VRTSvlic
```

where the *X* in `/altroot.5.X` corresponds to the Solaris version (8, 9 or 10).

If VxVM 4.0 or 4.1 was previously installed, remove the old VxVM packages on the alternate root by entering:

```
# rm /altroot.5.X/var/vx/isis/isis.lock
# /opt/VRTSob/bin/vxsvc -k
# pkgrm -R /altroot.5.X VRTSvmdoc VRTSvmman VRTScpi
# pkgrm -R /altroot.5.X VRTSstep VRTSap VRTSvrdoc \
  VRTSvrw
# pkgrm -R /altroot.5.X VRTSweb VRTSjre VRTSvcsvr \
  VRTSvrpro
# pkgrm -R /altroot.5.X VRTSfspro VRTSalloc
# pkgrm -R /altroot.5.X VRTSvmpro VRTSddlpr VRTSjre \
  VRTSperl
# pkgrm -R /altroot.5.X VRTSvxvm VRTSobgui VRTSob \
  VRTSmulic VRTSvlic
```

where the *X* in `/altroot.5.X` corresponds to the Solaris version (8, 9 or 10).

Note: Your system may not have all listed packages installed.

If you still have clients running previous versions of VxVM, refer to “[VMSA and VEA co-existence](#)” on page 24.

Licensing in VxVM 5.0 requires the new `VRTSvlic` package. You do not need to remove the existing `VRTSlic` package.

- 5 Install the VxVM 5.0 packages from the distribution media by entering the command:

```
# cd /cdrom/cdrom0/volume_manager
# ./installvm `hostname` -rootpath /altroot.5.X
```

- 6 After installing the Volume Manager software and patches, use the `vxlufinish` command to complete the Live Upgrade process.

```
# vxlufinish -u 5.8
```

The upgrade process is now complete.

- 7 Shut down the machine using `shutdown` or “`init 6`” to reboot the machine on the alternate root disk.

```
# init 6
```

Note: Do not use the `reboot` command to reboot the machine on an alternate root disk. If the alternate root disk has a problem rebooting or any upgrade issue, reboot the machine from the older disk.

- 8 Configure the VxVM 5.0 packages from the distribution media by entering the following commands:

```
# cd /cdrom/cdrom0/volume_manager
# ./installvm `hostname` -configure
```

- 9 To verify the upgrade, check the Volume Manager version:

```
# pkginfo -l VRTSvxvm
```

- 10 Verify that all the file systems that were under VxVM control prior to the upgrade are now under VxVM control.

```
# df -k
```

Upgrading VEA Windows Client Package

To upgrade the Windows Client Package, you need to uninstall the existing packages, and then install the new versions.

To upgrade the VEA Windows client package

- 1 Click **Start > Settings > Control Panel > Add or Remove Programs** and select “Veritas Enterprise Administrator language pack” for removal.
- 2 Uninstall the base VEA package the same way.
- 3 Install the base and language packages as described in the section entitled “[Installing on Windows](#)” on page 42.

Upgrading CVM protocol and disk group version

If you want to take advantage of the new features in this release, you must upgrade the Veritas Cluster Volume Manager (CVM) protocol version (70), and upgrade to the latest disk group version (140).

- 1 To upgrade the CVM protocol version to version 70, enter the following command on the master node:

```
# vxdctl upgrade
```

- 2 To upgrade the disk group version to 140, enter the command:

```
# vxldg -T 140 upgrade dgname
```

- 3 After upgrading CVM in a VCS environment, you should run the command:

```
# vxcvmconfig upgrade
```

If this command is not run, you will see a warning in the engine log file, `/opt/VRTSvcs/log/engine_A.log`.

Note: Downgrading disk group versions is not supported. If a Veritas cluster is used, the disk group version should match the lowest Volume Manager version installed in the cluster.

Updating variables

In `/etc/profile`, update the `PATH` and `MANPATH` variables as needed.

Default disk group

In releases prior to Volume Manager 4.0, the default disk group was `rootdg` (the root disk group). For Volume Manager to function, the `rootdg` disk group had to exist and it had to contain at least one disk.

This requirement no longer exists, however you may find it convenient to create a system wide default disk group. For instructions on how to change the default disk group, see the *Veritas Volume Manager System Administrator's Guide*. The main benefit of creating a default disk group is that VxVM commands default to the default disk group and you will not need to use the `-g` option. If you want to confirm that the root disk is encapsulated, enter the command:

```
# vxpdg bootdg
```

Upgrading Veritas File System

If you purchased Veritas File System only, use the steps in this section to upgrade the product.

Upgrading VxFS using the product installer

- 1 Unmount all VxFS file systems and Storage Checkpoints as described in “[Uninstalling Veritas File System](#)” on page 169.
- 2 Remove VxFS packages using either the `uninstallfs` script as described in “[Uninstalling VxFS using the uninstallation script](#)” on page 169, or using the `pkgrm` command as described in the section “[Uninstalling VxFS using the pkgrm command](#)” on page 170.
- 3 If you have VxFS file systems specified in the `/etc/vfstab` file, comment them out before rebooting, but do not remove the entries.
- 4 Add the VxFS packages as described in “[Installing VxFS using the product installer](#)” on page 146 or “[Installing VxFS using the pkgadd command](#)” on page 147.
- 5 Undo the changes to `/etc/vfstab` done in [step 3](#).

- 4 Shut down the machine using `shutdown` or `init 6` to reboot the machine on the alternate root disk.

```
# init 6
```

Note: Don't use the `reboot` command to reboot the machine on an alternate root disk. If the alternate root disk has a problem rebooting or any upgrade issue, reboot the machine from the older disk.

Verifying the upgrade

To verify the upgrade

- Verify that VxFS has been upgraded and are the new versions. Check a few of the packages. For example,

```
# pkginfo -l VRTSvxfs
# pkginfo -l VRTSfsmnd
# pkginfo -l VRTSfssdk
```

Uninstalling Veritas Volume Manager

If you are using Veritas Volume Manager only, use the steps in this section to remove the product.

Uninstalling VxVM using the uninstallation script

Note: The `uninstallvm` script also uninstalls any Veritas File System packages that are currently installed.

To uninstall VxVM

- 1 Read and follow the uninstallation procedures.
See [“Uninstalling Veritas Volume Manager”](#) on page 168.
- 2 Stop all applications from accessing VxVM volumes.
- 3 Move to the `/opt/VRTS/install` directory on your system and use the uninstallation script to remove Veritas Volume Manager.

```
# cd /opt/VRTS/install
# ./uninstallvm
```

Note: To ensure complete removal, a system reboot is recommended after all packages have been removed.

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


```
Enter the system names separated by spaces from which to
uninstall VxVM: host1
```
- 5 To verify the removal of the packages, use the `pkginfo` command.


```
# pkginfo | grep VRTS
```

Uninstalling VxVM using the `pkgrm` command

To uninstall VxVM

- ◆ Enter the following commands:

```
# pkgrm VRTSvrdoc VRTSvrw VRTSweb VRTSjre15 VRTSjre
# pkgrm VRTSvcsvr VRTSvrpro VRTSddlpr VRTSvdid
# pkgrm VRTSsvvc VRTSvmpro VRTSalloc VRTSdcli VRTSvman \
VRTSvmdoc
# pkgrm VRTSfspro VRTSdsa VRTSvxvm SYMClma VRTSspt VRTSaa
# pkgrm VRTSmh VRTSccg VRTSobgui VRTSob VRTSobc33 VRTSat
# pkgrm VRTSsmf VRTSspbx VRTSicsco VRTSperl
```

You can also include `VRTSvlic` in the removal line if you have not installed any other packages that use this package.

Uninstalling Veritas File System

If you installed Veritas File System only, use the steps in this section to remove the product.

Uninstalling VxFS using the uninstallation script

To uninstall VxFS

- 1 Comment out or remove any VERITAS File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.
- 2 Unmount all mount points for VxFS file systems.


```
# umount /mount_point
```
- 3 Move to the `/opt/VRTS/install` directory on your system and use the uninstallation script to remove Veritas File System.


```
# cd /opt/VRTS/install
# ./uninstallfs
```

Note: In order to ensure complete removal, reboot the system after all packages have been removed.

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

```
Enter the system names separated by spaces from which to
uninstall VxFS: host1
```
- 5 After the uninstall completes, remove any VxFS file system entries from the /etc/vfstab file.
- 6 To verify the removal of the packages, use the `pkginfo` command.

```
# pkginfo | grep VRTS
```

Uninstalling VxFS using the `pkgrm` command

The VxFS packages may also be uninstalled using the `pkgrm` command.

To uninstall VxFS

- 1 Remove the VxFS packages using the `pkgrm` command, starting with the optional packages. Do not remove the license packages `VRTSvlic` or `VRTSlic` if there are other Veritas products installed.

```
# pkgrm VRTSfsmnd VRTSfssdk VRTSfdoc VRTSfsman \
VRTSfsprv VRTSfscat VRTSap VRTStep VRTSfspro \
VRTSob VRTSvxfs
```

Note: If the `VRTSqio` package is installed, remove it. If the `VRTSfsnbl` package is installed, remove it also.

The system responds with a message similar to the following:

The following package is currently installed:

```
VRTSvxfs          Veritas File System
(sparc) 5.0,REV=5.0A14_sol
```

```
Do you want to remove this package? [y,n,?,q]
```

- 2 Type **y** to continue the removal.

```
## Removing installed package instance <VRTSvxfs>
This package contains scripts which will be executed with
superuser permission during the process of removing this
package.
Do you want to continue with the removal of this package
[y,n,?,q]
```

3 Type **y** to continue the removal.

```
## Verifying package dependencies
## Processing package information.
## Executing preremove script.
## Removing pathnames in class <s210b64>
. . .
## Removing pathnames in class <s210>
. . .
## Removing pathnames in class <all>
. . .
## Updating system information.
Removal of <VRTSvxfs> was successful.
```

4 After the uninstall completes, remove any VxFS file system entries from the `/etc/vfstab` file.

Installation script options

If you choose to install Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, or Veritas Storage Foundation for Sybase using the product installation script instead of the product installer, you have several options you can use.

Use the table to determine which installation script to use:

If you are installing	Use
Veritas Storage Foundation	installsf
Veritas Storage Foundation for DB2	installsfdb2
Veritas Storage Foundation for Oracle	installsfora
Veritas Storage Foundation for Sybase	installsfsyb
Veritas Volume Manager	installvm
Veritas File System	installfs

The following options apply to all Veritas Storage Foundation products.

Installation Script Options

The following options are available when using the product installation script. For an initial install or upgrade, options are not usually required.

Table A-6 Available command line options

Command Line Option	Function
<i>system1 system2...</i>	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 installing using the -installonly option.
-enckeyfile <i>encryption_key_file</i>	See the -responsefile and the -encrypt options.
-encrypt <i>password</i>	Encrypts <i>password</i> using the encryption key provided with the -enckeyfile option so that the encrypted password can be stored in response files.
-installpkgs	Displays all product packages in correct installation order. Output can be used to create scripts for command line installs, or for installations over a network. See the requiredpkgs option.
-installonly	Installs packages, but does not configure the product.
-keyfile <i>ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes -i <i>ssh_key_file</i> to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
-logpath <i>log_path</i>	Specifies a directory other than /opt/VRTS/install/logs as the location where installer log files, summary files, and response files are saved.
-noextrapkgs	Additional packages can be installed so that you can upgrade to another Symantec product simply by installing a new license. The -noextrapkgs option bypasses installation of extra product packages to simplify future maintenance updates.

Table A-6 Available command line options

Command Line Option	Function
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-nooptionalpkgs	Bypasses installation of optional product packages such as user documentation and manual pages.
-nostart	Bypasses startup of the product following installation and configuration.
-patchpath <i>patch_path</i>	Designates the path of a directory that contains all patches to install. The directory is typically an NFS mounted location and must be accessible all specified installation systems.
-pkgpath <i>package_path</i>	Designates the path of a directory that contains all packages to install. The directory is typically an NFS mounted location and must be accessible all specified installation systems.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
-requiredpkgs	Displays all required product packages in correct installation order. Optional packages are not listed. Output can be used to create scripts for command line installs, or for installations over a network. See <code>installpkgs</code> option.
-responsefile <i>response_file</i> [-enckeyfile <i>encryption_key_file</i>]	<p>Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. If not specified, the response file is automatically generated as <code>installerernumber.response.number</code> is random. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.</p> <p>The <code>-enckeyfile</code> option and <i>encryption_key_file</i> name are required with the <code>-responsefile</code> option when the response file contains encrypted passwords.</p>

Table A-6 Available command line options

Command Line Option	Function
<code>-rootpath <i>root_path</i></code>	<p>Specifies an alternative root directory on which to install packages.</p> <p>On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code>.</p> <p>On HP-UX operating systems, <code>-rootpath</code> passes <code>-I path</code> to <code>swinstall</code>.</p> <p>The <code>-rootpath</code> option is not supported on AIX or Linux operating systems.</p>
<code>-rsh</code>	<p>Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. The <code>-rsh</code> option requires that systems be preconfigured so that commands between systems execute without prompting for passwords or confirmations.</p>
<code>-tmppath <i>tmp_path</i></code>	<p>Specifies a directory other than <code>/var/tmp</code> 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.</p>

Veritas Storage Foundation install packages

Topics covered in this appendix include:

- [“Veritas Storage Foundation install packages”](#) on page 177
- [“Package installation for manual installations”](#) on page 187

Veritas Storage Foundation install packages

[Table B-1](#) shows the package name and contents for each English package for:

- Veritas Storage Foundation
- Veritas Storage Foundation for DB2
- Veritas Storage Foundation for Oracle
- Veritas Storage Foundation for Sybase

Table B-1 Storage Foundation packages

Package	Contents	Required/Optional
Volume Manager packages		
VRTSvxvm	Veritas Volume Manager Binaries	Required
VRTSvmweb	Veritas Volume Manager Management Services Web Client Extensions	Required
VRTSvmman	Veritas volume Manager Manual Pages	Optional

Table B-1 Storage Foundation packages

Package	Contents	Required/Optional
VRTSvm doc	Veritas Volume Manager Documentation	Optional
VRTSdcli	Veritas Distributed Command Line Interface	Required
VRTSalloc	Veritas Volume Manager Veritas Intelligent Storage Provisioning	Required
VRTSvmpro	Veritas Volume Manager Management Services Provider	Required
VRTSvsvc	Veritas Volume Server and Client Provider	Required
VRTSvdid	Veritas Device Identification API	Required
VRTSddlpr	Veritas Device Discovery Layer services Provider	Required
File System packages		
VRTSvxfs	VERITAS File System	Required
VRTSfspro	Veritas File System Management Services Provider	Required
VRTSfssdk	Veritas File System Software Developer Kit	Required
VRTSfsweb	Veritas File System Provider Web Client Extension	Required
VRTSfspro	Veritas File System Management Services Provider	Required
VRTSfsdoc	VERITAS File System Documentation	Optional
VRTSfsman	Veritas File System Manual Pages	Optional
VRTSfsmnd	Veritas File System Software Developer Kit Manual Pages	Optional
Database packages		
VRTSdbcom	Veritas Storage Foundation Common Utilities for Databases	Required (for all database products)
VRTSdb2ed	Veritas Storage Foundation for DB2	Required (for Storage Foundation for DB2)
VRTSd2gui	Veritas Storage Foundation for DB2 Graphical User Interface	Required (for Storage Foundation for DB2)

Table B-1 Storage Foundation packages

Package	Contents	Required/Optional
VRTSdbed	Veritas Storage Foundation for Oracle	Required (for Storage Foundation for Oracle)
VRTSorgui	Veritas Storage Foundation for Oracle Graphical User Interface	Required (for Storage Foundation for Oracle)
VRTSsybed	Veritas Storage Foundation for Sybase Note: Veritas Storage Foundation for Sybase (VRTSsybed) has not been localized.	Required (for Storage Foundation for Sybase)
VRTSodm	ODM Driver for VxFS	Required (for Storage Foundation for Oracle)
VRTSvxmsa	Veritas Mapping Service, Application Libraries	Required (for DB2 and Oracle products)
VRTSdbdoc	Veritas Storage Foundation Documentation for Databases	Optional
Veritas Enterprise Administrator packages		
VRTSob	Veritas Enterprise Administrator	Required
VRTSobc33	Veritas Enterprise Administrator Core	Required
VRTSccg	Veritas Enterprise Administrator Central Control Grid	Required
VRTSaa	Veritas Enterprise Administrator Action Agent	Required
VRTSobweb	Veritas Enterprise Administrator Web Console	Required
VRTSobgui	Veritas Enterprise Administrator	Optional
Infrastructure packages		
VRTSicsco	Symantec Infrastructure Core Services	Required
VRTSddlpr	Veritas Device Discovery Layer Services Provider	Required
VRTSvail	Veritas Array Integration Layer	Required
VRTSat	Symantec Product Authentication Service	Required
VRTSgapms	Veritas Generic Array Plugin	Required

Table B-1 Storage Foundation packages

Package	Contents	Required/Optional
High Availability (Veritas Cluster Server) packages		
VRTSvc	Veritas Cluster Server	Required
VRTScsow	Veritas Cluster Server Oracle and RAC Configuration Wizards	Required
VRTSvcso	Veritas High Availability Agent for Oracle	Required
VRTSvcdb	Veritas High Availability Agent for DB2	Required
VRTScscm	Veritas Cluster Server Cluster Manager	Required
VRTScscw	Veritas Cluster Server Configuration Wizards	Required
VRTScssim	Veritas Cluster Server Simulator	Required
VRTScutil	Veritas Cluster Utilities	Required
VRTSgab	Veritas Group Membership and Atomic Broadcast	Required
VRTSvcsvr	Veritas Cluster Server Agents for Veritas Volume Replicator	Required
VRTSjre	Veritas Java Runtime Environment Redistribution	Required
VRTSjre15	Veritas Java Runtime Environment Redistribution	Required
VRTSllt	Veritas Low Latency Transport	Required
VRTSvcstag	Veritas Cluster Server Bundled Agents	Required
VRTSvcsmg	Veritas Cluster Server English Message Catalogs	Required
VRTSvcsw	Veritas Cluster Manager (Web Console)	Required
VRTSvxfen	Veritas I/O Fencing	Required
VRTSacclib	Veritas Cluster Server ACC Library 5.0	Required
VRTScmcs	Veritas Cluster Management Console for single cluster environments	Required
VRTScmccc	Veritas Cluster Management Console Cluster Connector	Required

Table B-1 Storage Foundation packages

Package	Contents	Required/Optional
VRTSvcsdc	Veritas Cluster Server Documentation	Optional
VRTSvcsdr	Veritas Cluster Server Documentation	Optional
VRTSvcsmn	Manual Pages for Veritas Cluster Server	Optional
FlashSnap Agent for Symmetrix packages		
VRTSfas	Veritas FlashSnap Agent for Symmetrix	Required
VRTSfasag	Veritas Cluster Server Agents for Veritas FlashSnap Agent for Symmetrix	Required
VRTSfasdc	Veritas FlashSnap Agent for Symmetrix Documentation	Optional
Other Packages		
VRTSvlic	Veritas Licensing	Required
SYMClma	Symantec License Inventory Agent	Required
VRTScweb	Symantec Web Server	Required
VRTSweb	Symantec Web Server	Required
VRTSdcp	Veritas Disk Correlator Provider	Required
VRTSdsa	Veritas Datacenter Storage Agent	Required
VRTSdsm	Veritas datacenter Storage Manager	Required
VRTSgcscha	Veritas GCS High Availability Agents	Required
VRTSgcspr	Veritas SAN Global Configuration Server Object Bus Provider	Required
VRTSdbms3	Symantec Shared DBMS	Required
VRTSperl	Perl 5.8.8 Redistribution	Required
VRTSjre	Veritas JRE Redistribution	Required
VRTSjre15	Symantec JRE Redistribution	Required
VRTSspb	Symantec Private Branch Exchange	Required
VRTSspt	Veritas Software Support Tools	Required
VRTSdsa	Veritas Datacenter Storage Agent	Required
VRTSvdid	Veritas Device Identification API	Required

Table B-1 Storage Foundation packages

Package	Contents	Required/Optional
VRTSdbms3	Veritas Shared DBMS	Required
VRTScs	Veritas Storage Foundation Management Server	Required
VRTScsdoc	Veritas Storage Foundation Management Server	Required
VRTSmh	Veritas Centralized Management for Storage Foundation	Required
VRTScpi	Veritas Cross Product Installation Framework	Required
VRTSspb	Symantec Private Branch Exchange	Required
VRTSsmf	Symantec Service Management Framework	Required
VRTSat	Symantec Product Authentication Service	Required
VRTSspt	Veritas Software Support Tools	Required
windows/ vrtsofgui.msi	Veritas Enterprise Administrator for Windows	Optional
Veritas Volume Replicator		
VRTSvrpro	Veritas Volume Replicator Client Extension and Provider for Veritas Enterprise Administrator	Required
VRTSvcsvr	Veritas Cluster Server Agents for VVR	Required
VRTSvrw	Veritas Volume Replicator Web Console	Required
VRTSvrdoc	Veritas Volume Replicator Documentation	Optional

Japanese language packages

The table shows the package name and contents for each Japanese language package:

Table B-2 Japanese package contents

Package	Contents
Volume Manager Japanese packages	

Table B-2 Japanese package contents

Package	Contents
VRTSjavmc	Veritas Volume Manager (VRTSVXVM) Multi-language Package
VRTSmuvmp	Multi Language Veritas Volume Manager Management Services Provider
VRTSmualc	Veritas Intelligent Storage Provisioning (VRTSalloc) Multi-language Package
VRTSjavmd	Veritas Volume Manager Japanese Document Package
VRTSjavmm	Veritas Volume Manager Japanese Online Manual Document Package
File System Japanese packages	
VRTSjafsc	Veritas File System (VRTSVXFS) Japanese Language Package
VRTSmufsp	Veritas File System Provider (VRTSfspro) Multi-language Package
VRTSjafsd	Veritas File System Japanese Document Package
VRTSjafsm	Veritas File System Japanese Online Manual Document Package
Database Japanese packages	
VRTSjadb2	Veritas Storage Foundation <i>for DB2</i> (VRTSdb2ed) Japanese Package
VRTSjad2g	Veritas Storage Foundation <i>for DB2</i> Graphical User Interface (VRTSd2gui) Japanese Package
VRTSjadbe	VERITAS Storage Foundation <i>for Oracle</i> (VRTSdbed) Japanese Package
VRTSjaorg	VERITAS Storage Foundation <i>for Oracle</i> Graphical User Interface (VRTSorgui) Japanese Package
VRTSjaodm	Veritas Oracle Disk Manager (VRTSodm), Japanese Package
VRTSjamsa	Veritas VxMS Mapping Service (VRTSVxmsa) Japanese Package
VRTSjadbd	Japanese Veritas Storage Foundation Documentation for Databases

Table B-2 Japanese package contents

Package	Contents
Veritas Enterprise Administrator Japanese packages	
VRTSmuob	Veritas Enterprise Administrator Service Localized Package
Infrastructure Japanese packages	
VRTSmuddl	Veritas Device Discovery Layer Provider (VRTSddlpr) Multi-language Package
VRTSjaap	Veritas Array Providers - Japanese Support
VRTSmuap	Veritas Action Provider Language Package
VRTSmutep	Veritas Task Exec Provider Language Package
High Availability (Veritas Cluster Server) Japanese packages	
VRTSjacs	Japanese Veritas Cluster Server Message Catalogs
VRTSjacsu	Japanese Veritas Cluster Utility Language Pack
VRTSjacsj	Japanese Veritas Cluster Server Cluster Manager
VRTSjacsm	Japanese Veritas Cluster Server Simulator
VRTSjacsw	Japanese Veritas Cluster Manager (Web Console)
VRTSjawareb	Japanese Veritas Web Server Language Pack
VRTSjacsd	Japanese Veritas Cluster Server Documentation
FlashSnap Agent for Symmetrix Japanese packages	
VRTSjafag	Japanese Veritas Cluster Server Agents for Veritas FlashSnap Agent for Symmetrix
VRTSjafad	Japanese Veritas FlashSnap Agent for Symmetrix Documentation
VRTSjafas	Japanese Veritas Flashsnap Agent for Symmetrix
Other Japanese packages	
VRTSmulic	Multi-language Veritas License Utilities
VRTSmuobg (Windows client only)	Veritas Enterprise Administrator Localized GUI Package

Chinese language packages

The table shows the package name and contents for each Chinese language package.

Table B-3 Chinese package contents

Package	Contents
Volume Manager Chinese packages	
VRTSmuvmp	Multi Language Veritas Volume Manager Management Services Provider
VRTSmualc	Veritas Intelligent Storage Provisioning (VRTSalloc) Multi-language Package
VRTSmuob	Veritas Enterprise Administrator Service Localized Package
VRTSzhvmm	Chinese Veritas Volume Manager, Manual Pages
VRTSzhvmd	Chinese Veritas Volume Manager (User Documentation)
VRTSzhvmc	Chinese Veritas Volume Manager - Message Catalogs
VRTSmufsp	Multi-language Veritas File System Management Services Provider
VRTSmuddl	Veritas Device Discovery Layer Provider (VRTSddlpr) Multi-language Package
VRTSmulic	Multi-language Veritas License Utilities
VRTSmuobg (Windows client only)	Veritas Enterprise Administrator Localized GUI Package

French language packages

The table shows the package name and contents for each French language package:

Table B-4 French package contents

Package	Contents
Volume Manager French packages	
VRTSmuvmp	Multi Language Veritas Volume Manager Management Services Provider

Table B-4 French package contents

Package	Contents
VRTSmualc	Veritas Intelligent Storage Provisioning (VRTSa11oc) Multi-language Package
VRTSmuob	Veritas Enterprise Administrator Service Localized Package
VRTSfrvmm	French Veritas Volume Manager, Manual Pages
VRTSfrvmd	French Veritas Volume Manager (User Documentation)
VRTSfrvmc	French Veritas Volume Manager - Message Catalogs
VRTSmufsp	Multi-language Veritas File System Management Services Provider
VRTSmuc33	Veritas Enterprise Administrator Core Language Package by Symantec
VRTSmuvmw	Multi Language Veritas Volume Manager Management Services Web Client Extensions by Symantec
VRTSmuddl	Veritas Device Discovery Layer Provider (VRTSddlpr) Multi-language Package
VRTSatFR	Symantec Product Authentication Service Software French language Kit
VRTSfrico	Symantec Infrastructure Core Services Common French Language
VRTSfrpbx	Symantec Private Branch Exchange French Language
VRTSfrsmf	Symantec Service Management Framework French Language
VRTSmulic	Multi-language Veritas License Utilities
VRTSmuobg (Windows client only)	Veritas Enterprise Administrator Localized GUI Package

Package installation for manual installations

If you are choose to install the Veritas packages manually, they must be unzipped and installed in a particular order.

Unzipping the packages

Before you can install the packages, you must unzip them.

To unzip the packages

- 1 Log in as superuser (`root`).
- 2 Create a directory for installation.

```
# mkdir /parent_directory/install
```
- 3 Insert the software disc with the Veritas Storage Foundation software into a drive connected to the system. The Solaris volume-management software automatically mounts the disc as `/cdrom/cdrom0`. Type the command:

```
# cd /cdrom/cdrom0
```
- 4 Copy the compressed package files from the software disc to the temporary directory.

For	Use
Veritas Storage Foundation	<pre># cp -r storage_foundation/pkgs/* /parent_directory/install</pre>
Veritas Storage Foundation for DB2	<pre># cp -r storage_foundation/pkgs/* /parent_directory/install</pre>
Veritas Storage Foundation for Oracle	<pre># cp -r storage_foundation_for_db2/ pkgs/* /parent_directory/install</pre>
Veritas Storage Foundation for Sybase	<pre># cp -r storage_foundation_for_oracle/pkgs/* /parent_directory/install</pre>
Veritas Volume Manager	<pre># cp -r volume_manager/pkgs/* /parent_directory/install</pre>
Veritas File System	<pre># cp -r file_system /pkgs/* /parent_directory/install</pre>

- 5 Go to the temporary directory and unzip the compressed package files.

Note: If your system does not have the `gunzip` utility, copy it from the disc:

```
# cp /cdrom/dcrom0/gnu/gunzip /parent_directory/install
```

```
# cd /parent_directory/install  
# gunzip VRTS*.gz
```

- 6 Decompress and extract each package.

```
# tar xf package_name.tar  
# tar xf package_name.tar  
# tar xf package_name.tar  
.  
.
```

- 7 List the files in the temporary directory.

```
# ls /parent_directory/install
```

Installing the packages using the `pkgadd` command

To manually install the packages

- 1 Change to the `pkgs` directory that now contains the VxFS packages:

```
# cd /parent_directory/install
```

- 2 Install the packages:

```
# pkgadd -d . package_name package_name ...  
#
```

To get the package order for	See
Veritas Storage Foundation	“Veritas Storage Foundation” on page 189
Veritas Storage Foundation for DB2	“Veritas Storage Foundation for DB2” on page 189
Veritas Storage Foundation for Oracle	“Veritas Storage Foundation for Oracle” on page 190
Veritas Storage Foundation for Sybase	“Veritas Storage Foundation for Sybase” on page 190
Veritas Volume Manager	“Veritas Volume Manager” on page 191
Veritas File System	“Veritas File System” on page 191

Veritas Storage Foundation

To get the package installation order for Veritas Storage Foundation, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation

- 1 Mount the software disc.
See [“Mounting the software disc”](#) on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsf -requiredpkgs
```

or

```
# ./installsf -installpkgs
```

Veritas Storage Foundation for DB2

To get the package installation order for Veritas Storage Foundation for DB2, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation for DB2

- 1 Mount the software disc.
See [“Mounting the software disc”](#) on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsfdb2 -requiredpkgs
```

or

```
# ./installsfdb2 -installpkgs
```

Veritas Storage Foundation for Oracle

To get the package installation order for Veritas Storage Foundation for Oracle, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation for Oracle

- 1 Mount the software disc.
See "[Mounting the software disc](#)" on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsfora -requiredpkgs
```

or

```
# ./installsfora -installpkgs
```

Veritas Storage Foundation for Sybase

To get the package installation order for Veritas Storage Foundation for Sybase, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Storage Foundation for Sybase

- 1 Mount the software disc.
See "[Mounting the software disc](#)" on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installsfsyb -requiredpkgs
```

or

```
# ./installsfsyb -installpkgs
```

Veritas Volume Manager

To get the package installation order for Veritas Volume Manager, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas Volume Manager

- 1 Mount the software disc.
See “[Mounting the software disc](#)” on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installvm -requiredpkgs
```

or

```
# ./installvm -installpkgs
```

Veritas File System

To get the package installation order for Veritas File System, use the `-requiredpkgs` or `-installpkgs` option with the script from the disc.

To get package installation order for Veritas File System

- 1 Mount the software disc.
See “[Mounting the software disc](#)” on page 32.
- 2 Move to the disc directory.

```
# cd /cdrom/cdrom0
```
- 3 Run the script with the `-requiredpkgs` or `-installpkgs` option. The `-requiredpkgs` option displays only the required packages and the `-installpkgs` displays all packages.

```
# ./installfs -requiredpkgs
```

or

```
# ./installfs -installpkgs
```


Configuring the Symantec License Inventory Agent

This appendix includes the following topics:

- [“About the Symantec License Inventory Manager”](#) on page 194
- [“When the Symantec License Inventory Agent is installed”](#) on page 195
- [“When the server and access points are installed”](#) on page 195
- [“What you can do with the agent after it is installed”](#) on page 195
- [“How to remove the agent”](#) on page 196
- [“How to order the Symantec License Inventory Manager license and media kit”](#) on page 197

The Symantec License Inventory Manager installation disc is available separately. For information on how to order the full product, see [“How to order the Symantec License Inventory Manager license and media kit”](#) on page 197. The installation media provides online documentation with details on all topics discussed in this appendix.

Read the following Technical Support TechNote for the latest information on updates, patches, and software issues regarding this product:

<http://support.veritas.com/docs/282183>

You can also download the *Symantec License Inventory Agent 4.1 Release Notes*, from this website.

About the Symantec License Inventory Manager

The Symantec License Inventory Manager (license inventory manager) is an enterprise asset management tracking tool that inventories Symantec Information Availability products in your network and consolidates critical information on the deployment of these products to facilitate license management and compliance tracking. Using the information provided by the license inventory manager, you can:

- Determine all the Symantec software products and licenses being used in your enterprise
- Achieve easier license self-compliance management
- Know your Enterprise License Agreement deployment status
- Reduce administrative overhead for managing license compliance
- Renew support and maintenance based on the licenses you have deployed
- Gain more control over your Symantec software usage
- Manage department chargebacks based on actual software usage
- Use more flexible licensing and pricing models
- Exploit detailed deployment data to perform return on investment analyses for purchased software

The license inventory manager is a three-tiered system that consists of a server tier, access point tier, and an agent tier. The server tier is the Symantec License Inventory Server, which consolidates and stores information that it gathers from the agents and access points.

The optional access point tier includes Symantec License Inventory Access Points and serves as a consolidation layer between the agents and server.

The agent tier includes Symantec License Inventory Agents, which are deployed on individual hosts in a network. Each agent gathers product information on the supported Symantec products that are installed on the agent's host, then sends the information to an access point or the server.

When the Symantec License Inventory Agent is installed

The Symantec product installer installs or upgrades the agent on the host with the Symantec product. The agent is installed in the following directory:

`/opt/SYMC1ma`

The agent is installed with a default configuration that minimizes its impact on a running system. The minimum configuration prevents remote communication with the agent to keep its data and interfaces secure.

When the server and access points are installed

The server and access points are not installed automatically. If you want to use the Symantec License Inventory Manager, you must manually install the server and, optionally, the access points. After you install the server and access points, the agents can gather information and you can create inventory reports.

You can install the server and access points from the Symantec License Inventory Manager installation disc.

What you can do with the agent after it is installed

If you are already participating in a Symantec sales program that requires the use of the agent, or if you want to order and deploy the Symantec License Inventory Manager, you can use the agent to track Symantec products on the systems on which it was installed. To use the agent, however, you must manually configure it to enable remote communication between the agent and its server or access point.

Complete instructions for reconfiguring the agent are provided in the *Symantec License Inventory Manager 4.1 Release Notes*. You can download this document from the following website:

<http://support.veritas.com/docs/282183>

How to remove the agent

If you do not want to use the Symantec License Inventory Manager, you can remove the agent using the operating system package removal commands to remove the agent packages, which include SYMClma and VRTSsmf.

The server and access point also use the VRTSsmf package. If the server or access point is installed on this host with the agent, you can remove the SYMClma package, but not the VRTSsmf package. If neither the server nor the access point is installed on this host, you can remove both the SYMClma and VRTSsmf packages.

If you remove both packages, remove the SYMClma package first.

[Table C-1](#) lists the commands required to remove these packages on the supported platforms.

Table C-1 Package removal commands required to remove the agent

Platform	Package removal command
AIX	<code>installp -u VRTSlma</code> <code>installp -u VRTSsmf</code>
HP-UX	<code>swremove SYMClma</code> <code>swremove VRTSsmf</code>
Linux	<code>rpm evv SYMClma</code> <code>rpm evv VRTSsmf</code>
Solaris	<code>pkgrm VRTSlma</code> <code>pkgrm VRTSsmf</code>

Later, you can reinstall the agent with the Symantec License Inventory Manager installation disc. This disc is available in the Symantec License Inventory Manager kit.

How to order the Symantec License Inventory Manager license and media kit

To order a Symantec License Inventory Manager license and media kit, contact your Symantec sales representative.

The installation media provides online documentation for the Symantec License Inventory Manager. You can contact your sales representative to order printed copies of the documentation. The documents you can order include:

- *Symantec License Inventory Manager Installation and Configuration Guide*
- *Symantec License Inventory Manager Administrator's Guide*
- *Symantec License Inventory Manager User's Guide*

