

# Veritas Storage Foundation™ Installation Guide

Solaris

5.0 Maintenance Pack 3



# Veritas Storage Foundation™ Installation Guide

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

Product version: 5.0 MP3

Document version: 5.0MP3.0

## Legal Notice

Copyright © 2008 Symantec Corporation. All rights reserved.

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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

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.

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

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

Symantec Corporation  
20330 Stevens Creek Blvd.  
Cupertino, CA 95014

<http://www.symantec.com>

# Technical Support

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

Symantec's maintenance offerings include the following:

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

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

[www.symantec.com/techsupp/](http://www.symantec.com/techsupp/)

## Contacting Technical Support

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

[www.symantec.com/business/support/index.jsp](http://www.symantec.com/business/support/index.jsp)

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

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

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

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

## Licensing and registration

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

[www.symantec.com/techsupp/](http://www.symantec.com/techsupp/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/techsupp/](http://www.symantec.com/techsupp/)

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

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

## Maintenance agreement resources

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

Asia-Pacific and Japan	<a href="mailto:contractsadmin@symantec.com">contractsadmin@symantec.com</a>
Europe, Middle-East, and Africa	<a href="mailto:semea@symantec.com">semea@symantec.com</a>
North America and Latin America	<a href="mailto:supportsolutions@symantec.com">supportolutions@symantec.com</a>

## Additional enterprise services

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

Enterprise services that are available include the following:

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

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

[www.symantec.com](http://www.symantec.com)

Select your country or language from the site index.

# Contents

Technical Support .....	4	
Chapter 1	About Storage Foundation and High-Availability Solutions .....	15
	Veritas Storage Foundation product suites .....	15
	About Veritas Enterprise Administrator (VEA) .....	19
Chapter 2	Before you install .....	21
	About planning for a Storage Foundation installation .....	21
	Veritas Installation Assessment Service .....	22
	Release notes .....	22
	Symantec product licensing .....	22
	Setting environment variables .....	23
	Creating the /opt directory .....	23
	Configuring secure shell (ssh) or remote shell before installing products .....	24
	Configuring and enabling ssh .....	24
	Restarting ssh .....	29
	Enabling and disabling rsh for Solaris .....	29
	Cluster environment requirements for Sun Clusters .....	30
	Prerequisites for Storage Foundation Cluster File System .....	31
	Hardware overview and requirements for Storage Foundation Cluster File System .....	32
	Shared storage .....	33
	Fibre Channel switch .....	33
	Cluster platforms .....	34
	Preinstallation or upgrade planning for Veritas Volume Replicator .....	34
	VEA installation planning .....	34
	Planning an upgrade from the previous VVR version .....	35
	Additional settings for using VVR in a localized environment .....	36
	Database requirements .....	36
	About centralized management .....	38
	Downloading the Storage Foundation and High Availability software .....	38

	Downloading Storage Foundation Manager .....	39
Chapter 3	System requirements .....	41
	Software and hardware requirements .....	41
	Supported Solaris operating systems .....	41
	Storage Foundation Cluster File System node requirements .....	42
	Disk space requirements .....	42
	Disk space .....	42
	Disk space requirements for Veritas Volume Replicator .....	43
Chapter 4	Installing Storage Foundation using the common product installer .....	45
	Installation quick reference .....	46
	Mounting a software disc .....	46
	About the common product installer .....	47
	Installing and configuring Storage Foundation using the common product installer .....	48
	Installing and configuring Storage Foundation and High Availability Solutions using the common product installer .....	52
	Installing and configuring Storage Foundation Cluster File System using the common product installer .....	59
	Installing and configuring Veritas Volume Replicator using the common product installer .....	65
	Installing VVR when VxVM is already installed .....	71
	Installing language packages .....	71
	Installing the simplified Chinese and French language packages .....	71
	Installing the Japanese language packages .....	73
	Installing the VEA language packages .....	74
	Installing Veritas Enterprise Administrator .....	75
	Installing the Veritas Enterprise Administrator client .....	75
	Installing the Veritas Enterprise Administrator client on Solaris .....	76
	Installing the VEA client on Microsoft Windows .....	77
Chapter 5	Installing Storage Foundation, other methods .....	79
	Installing Storage Foundation - summary .....	79
	Installing the Storage Foundation products using JumpStart .....	80
	Summary of JumpStart installation tasks .....	80
	Copying and unzipping the packages .....	82
	Determining the installation order .....	84

	Creating the response files for JumpStart .....	86
	Installing Storage Foundation using the pkgadd command .....	88
Chapter 6	Configuring Storage Foundation and High Availability products .....	95
	Configuring the Storage Foundation products .....	95
	Configuring Storage Foundation .....	96
	Configuring Storage Foundation and High Availability Solutions .....	98
	Required information for configuring Storage Foundation and High Availability Solutions .....	98
	Configuring Veritas Storage Foundation and High Availability Solutions .....	99
	About adding and removing nodes in a cluster .....	106
	Configuring Storage Foundation for Databases .....	106
	Database configuration requirements .....	111
	Creating and configuring the repository database for DB2 and Oracle .....	111
	Setting administrative permissions for databases .....	114
	Configuring Veritas Volume Manager .....	115
	About hot-relocation .....	116
	Starting and enabling the configuration daemon .....	116
	Starting the volume I/O daemon .....	117
	Configuring new disks .....	117
	Using vxinstall to configure Veritas Volume Manager .....	118
	Preventing multipathing/suppress devices from VxVM's view .....	120
	Using Dynamic Multipathing with EMC Symmetrix Disk Arrays .....	122
	Enabling Sun Java System Cluster support in VxVM (optional) .....	122
	Configuring shared disks on Solaris .....	122
	Reserving shared disks .....	123
	Adding new array support .....	123
	About placing disks in another disk group .....	123
	Protecting your system and data .....	123
	Enabling the Intelligent Storage Provisioning (ISP) feature .....	124
	Enabling cluster support in VxVM (Optional) .....	125
	Configuring Veritas File System .....	128
	Loading and unloading the file system module .....	128
	vxtunefs command permissions and Cached Quick I/O .....	129
	Configuring Storage Foundation Cluster File System .....	129
	Configuring Veritas Volume Replicator .....	134

Veritas Volume Replicator and Volume Manager setup after installation .....	138
Configuring and starting Veritas Enterprise Administrator .....	138
Stopping and starting the VEA server .....	139
Starting the VEA client on Windows or Solaris .....	139
VMSA and VEA co-existence .....	140
Configuring Veritas Enterprise Administrator for databases .....	141
Configuring Veritas Enterprise Administrator for Oracle .....	141
Setting up Veritas Enterprise Administrator for DB2 .....	143

## Chapter 7

Upgrading Storage Foundation .....	147
Upgrading Storage Foundation or the operating system .....	147
Planning the upgrade .....	148
Saving system information before upgrade .....	149
Determining if the root disk is encapsulated .....	149
Upgrade paths for Storage Foundation 5.0MP3 .....	149
Performing pre-installation checks and configuration .....	153
Verifying that the file systems are clean .....	154
Upgrading external ASLs and APMs .....	155
Upgrading Storage Foundation from 5.0 to 5.0MP3 using the product installer or manual steps .....	156
Upgrading external 4.x ASL or APM packages from a previous 5.0 release to Storage Foundation 5.0MP3 .....	156
Upgrading Storage Foundation software from 5.0 to 5.0MP3 using the product installer .....	159
Upgrading Storage Foundation from 5.0 to 5.0MP3 with manual steps (patchadd) .....	160
Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using the product installer or manual steps .....	162
Upgrading external ASL or APM packages from Storage Foundation 4.x to 5.0MP3 .....	162
Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required .....	164
Upgrading Veritas Storage Foundation with manual steps when OS upgrade is not required .....	169
Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using upgrade scripts (OS upgrade or encapsulated root disk) .....	171
Upgrading Storage Foundation and/or Solaris using Live Upgrade .....	174
Upgrading Storage Foundation or the Solaris OS or Both Using Live Upgrade .....	175
Installing Live Upgrade on the current root disk .....	176

Storage Foundation Live Upgrade commands and usage .....	176
Beginning the Live Upgrade .....	178
Upgrading Solaris from the software discs .....	179
Upgrading Solaris from network directory path .....	179
Completing the Solaris upgrade .....	180
Upgrading Storage Foundation .....	180
Upgrading the Solaris OS only .....	184
Upgrading Veritas Enterprise Administrator clients .....	187
Upgrading the VEA client on a Microsoft Windows system .....	187
Upgrading the VEA client on a Solaris system .....	188
Upgrading the VEA Windows client language package .....	189
Upgrading Veritas Volume Replicator .....	189
Supported upgrade methods for Veritas Volume Replicator .....	189
Upgrading Veritas Volume Replicator using the Veritas product installer .....	190
Upgrading using VVR upgrade scripts .....	194
Upgrading VVR without disrupting replication .....	198
Upgrading VVR when VCS agents are configured .....	199
Upgrading language packages .....	207
Post-upgrade tasks .....	207
Optional configuration steps .....	207
Upgrading to the new repository database for DB2 and Oracle .....	208
Changing permissions for Storage Foundation for Databases .....	209
About upgrading disk layout versions .....	211
Migrating from /etc/vx/vxdba to /var/vx/vxdba for DB2 and Oracle .....	212
Upgrading VxVM disk group versions .....	213
Updating variables .....	213
Setting the default disk group .....	214
Upgrading the Array Support Library .....	214
Converting from QuickLog to Multi-Volume support .....	224
Verifying the Veritas Storage Foundation upgrade .....	226

## Chapter 8

### Upgrading Storage Foundation Cluster File System .....

About upgrading Storage Foundation Cluster File System .....	227
Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0 .....	228
Planning the upgrade .....	228
Preparing the system and backing up files before upgrading .....	228

Upgrade paths for Veritas Storage Foundation Cluster File System	
5.0MP3 .....	229
Overview of procedures .....	230
Ensuring the file systems are clean (full only) .....	231
Modifying the main.cf file (phased or full) .....	232
Performing the upgrade (phased or full) .....	234
Making the file systems clean .....	243
Upgrading language packages .....	245
Upgrading Storage Foundation Cluster File System to 5.0MP3 on a	
Storage Foundation 5.0 system .....	245
Preparing to upgrade to the Maintenance Pack .....	246
Phased upgrade for a Maintenance Pack .....	246
Full upgrade for a Maintenance Pack .....	252
Chapter 9	
Verifying the Storage Foundation installation .....	257
Verifying that the products were installed .....	257
Installation log files .....	257
Using the installation log file .....	258
Using the response file .....	258
Using the summary file .....	258
Checking Volume Manager processes .....	258
Checking Veritas File System installation .....	259
Verifying kernel installation .....	259
Verifying command installation .....	259
Verifying the configuration files for Storage Foundation Cluster File	
System .....	260
Low Latency Transport configuration files .....	260
Checking Low Latency Transport operation .....	261
Group Membership and Atomic Broadcast configuration files .....	264
Checking Group Membership and Atomic Broadcast	
operation .....	264
Checking cluster operation .....	265
Verifying agent configuration for Storage Foundation Cluster File	
System .....	268
Synchronizing time on Cluster File Systems .....	269
Configuring VCS for Storage Foundation Cluster File System .....	269
main.cf file .....	269
Storage Foundation Cluster File System HA Only .....	271
Veritas Cluster Server application failover services .....	271

Chapter 10	Uninstalling Storage Foundation .....	273
	About removing Veritas Storage Foundation .....	273
	Uninstallation requirements for Solaris .....	273
	Dropping the repository database for DB2 and Oracle .....	283
	Shutting down cluster operations .....	284
	Shutting down Veritas Volume Manager .....	284
	Uninstalling Veritas Storage Foundation .....	284
	Uninstalling Storage Foundation Cluster File System .....	286
	Uninstalling the VCS agents for VVR .....	287
	Disabling the agents on a system .....	287
	Uninstalling Veritas Volume Replicator (VVR) .....	288
	Removing the Replicated Data Set .....	289
	Removing the VVR packages .....	290
	Additional ways to remove VVR packages .....	291
	Uninstalling language packages .....	293
	Uninstalling language packages using the pkgrm command .....	293
	Uninstalling Storage Foundation using the pkgrm command .....	294
	Removing the Veritas Enterprise Administrator client .....	295
Appendix A	Installation scripts .....	297
	About installation scripts .....	297
	Installation script options .....	298
Appendix B	Storage Foundation and High Availability components .....	303
	Veritas Storage Foundation installation packages .....	303
	Obsolete packages in Storage Foundation .....	311
Appendix C	Troubleshooting information .....	313
	Troubleshooting information .....	313
	Storage Foundation Cluster File System installation issues .....	313
	Incorrect permissions for root on remote system .....	314
	Inaccessible system .....	314
	Storage Foundation Cluster File System problems .....	314
	Unmount failures .....	314
	Mount failures .....	315
	Command failures .....	316
	Performance issues .....	316
	High availability issues .....	316

Index ..... 319

# About Storage Foundation and High-Availability Solutions

This chapter includes the following topics:

- [Veritas Storage Foundation product suites](#)
- [About Veritas Enterprise Administrator \(VEA\)](#)

## Veritas Storage Foundation product suites

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

**Table 1-1** Contents of Veritas Storage Foundation products

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

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

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

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

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

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

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

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

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

## About Veritas Enterprise Administrator (VEA)

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



# Before you install

This chapter includes the following topics:

- [About planning for a Storage Foundation installation](#)
- [Release notes](#)
- [Symantec product licensing](#)
- [Setting environment variables](#)
- [Creating the /opt directory](#)
- [Configuring secure shell \(ssh\) or remote shell before installing products](#)
- [Cluster environment requirements for Sun Clusters](#)
- [Prerequisites for Storage Foundation Cluster File System](#)
- [Hardware overview and requirements for Storage Foundation Cluster File System](#)
- [Preinstallation or upgrade planning for Veritas Volume Replicator](#)
- [Database requirements](#)
- [About centralized management](#)
- [Downloading the Storage Foundation and High Availability software](#)
- [Downloading Storage Foundation Manager](#)

## About planning for a Storage Foundation installation

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

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

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

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

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

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

See “[Veritas Storage Foundation product suites](#)” on page 15.

## Veritas Installation Assessment Service

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

<https://vias.symantec.com/>

## Release notes

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

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

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

## Symantec product licensing

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

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

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

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

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

## Setting environment variables

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

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

If you install the Veritas manual pages, set the path of your `MANPATH` environment variable to include the relevant directories.

Add the following directories to your `PATH` environment variable:

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

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

## Creating the `/opt` directory

The directory `/opt` must exist, be writable and must not be a symbolic link.

If you are upgrading, you cannot have a symbolic link from `/opt` to an unconverted volume. If you do have a symbolic link to an unconverted volume, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

Ensure that the `/opt` directory exists and has write permissions for `root`.

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

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

You can install products to remote systems using either secure shell (ssh) or remote shell (rsh). ssh is the preferred method of remote communication because it provides a greater level of security than the rsh suite of protocols.

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

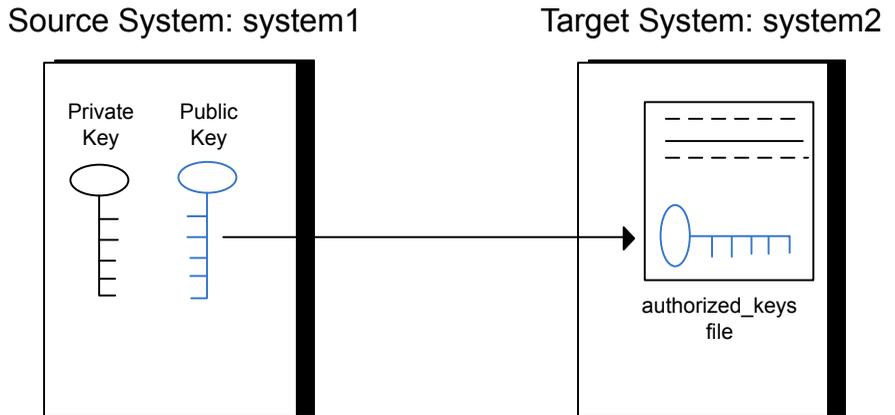
### Configuring and enabling ssh

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

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

[Figure 2-1](#) illustrates this procedure.

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



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

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

**To create the DSA key pair**

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

```
system1 # cd /
```

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

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

System output similar to the following is displayed:

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

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

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

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

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

- 5 Make sure the `/.ssh` directory is on all the target installation systems (system2 in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

```
system2 # cd /
system2 # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

```
system2 # chmod go-w /.ssh
```

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

- 1 Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems (system2 in this example).

To enable SFTP, the `/etc/ssh/sshd_config` file must contain the following two lines:

```
PermitRootLogin          yes
Subsystem sftp           /usr/lib/ssh/sftp-server
```

- 2 If the lines are not there, add them and restart ssh.

To restart ssh on Solaris 10, type the following command:

```
system1 # svcadm restart ssh
```

To restart on Solaris 9, type the following commands:

```
system1 # /etc/init.d/sshd stop
system1 # /etc/init.d/sshd start
```

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

Use the secure file transfer program.

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

Use the following command for secure file transfer:

```
system1 # sftp system2
```

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

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

- 4** Enter `yes`.

Output similar to the following is displayed:

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

- 5** Enter the root password of system2.

- 6** At the `sftp` prompt, type the following command:

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

The following output is displayed:

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

- 7** To quit the SFTP session, type the following command:

```
sftp> quit
```

- 8 To begin the `ssh` session on the target system (system2 in this example), type the following command on system1:

```
system1 # ssh system2
```

Enter the root password of system2 at the prompt:

```
password:
```

- 9 After you log in to system2, enter the following command to append the `id_dsa.pub` file to the authorization key file:

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

- 10 After the `id_dsa.pub` public key file is copied to the target system (system2), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, type the following command on system2:

```
system2 # rm /id_dsa.pub
```

- 11 To log out of the `ssh` session, type the following command:

```
system2 # exit
```

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

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

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

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

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

```
system1 # ssh-add
```

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

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

### To verify that you can connect to a target system

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

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

where system2 is the name of the target system.

- 2 The command should execute from the source system (system1) to the target system (system2) without the system requesting a passphrase or password.
- 3 Repeat this procedure for each target system.

## Restarting ssh

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

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

### To restart ssh

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

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

- 2 Make the key globally available for the user root

```
system1 # ssh-add
```

## Enabling and disabling rsh for Solaris

The following section describes how to enable remote shell on Solaris system.

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

See [“Configuring and enabling ssh”](#) on page 24.

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

### To enable rsh

- 1 To determine the current status of `rsh` and `rlogin`, type the following command:

```
# inetadm | grep -i login
```

If the service is enabled, the following line is displayed:

```
enabled online svc:/network/login:rlogin
```

If the service is not enabled, the following line is displayed:

```
disabled disabled svc:/network/login:rlogin
```

- 2 To enable a disabled `rsh/rlogin` service, type the following command:

```
# inetadm -e rlogin
```

- 3 To disable an enabled `rsh/rlogin` service, type the following command:

```
# inetadm -d rlogin
```

- 4 Modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. This file must be modified for each user who remotely accesses the system using `rsh`. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system having access to the local system. For example, if the root user must remotely access `system1` from `system2`, you must add an entry for `system2.companyname.com` in the `.rhosts` file on `system1`.

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

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

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

## Cluster environment requirements for Sun Clusters

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:

```
# vxvgs 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 Sun Java System Cluster support in VxVM \(optional\)”](#) on page 122.

## Prerequisites for Storage Foundation Cluster File System

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

Provide the following information when installing the SFCFS:

- The cluster name, beginning with a letter (a-z, A-Z).
- A unique ID from 0-65535 for the cluster. Within the public subnet, a new cluster using a duplicate cluster ID can cause existing clusters to fail.
- The host names of the cluster nodes.
- The device names of the network interface cards (NICs) used for the private networks among nodes.
- Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node

from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities as root on all cluster nodes or remote systems.

- Symantec recommends configuring the cluster with I/O fencing enabled. I/O fencing requires shared devices to support SCSI-3 Persistent Reservations (PR). Enabling I/O fencing prevents data corruption caused by a split brain scenario.

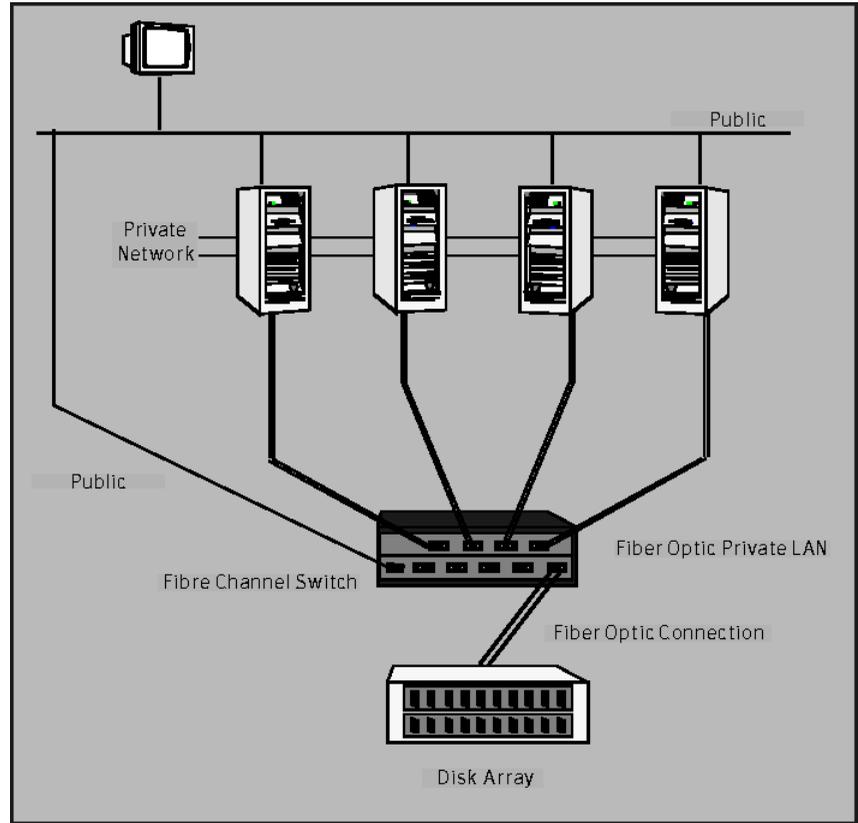
The Storage Foundation Cluster File System is supported without I/O fencing enabled. However, without I/O fencing enabled, split brain scenarios can result in data corruption.

## Hardware overview and requirements for Storage Foundation Cluster File System

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

[Figure 2-2](#) shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

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



## Shared storage

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

## Fibre Channel switch

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

## Cluster platforms

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

See the *Veritas Storage Foundation Release Notes*.

---

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

---

# Preinstallation or upgrade planning for Veritas Volume Replicator

Before installing or upgrading VVR:

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

The following related documents are available on the documentation disc:

*Veritas Volume Replicator Planning and Tuning Guide* Provides detailed explanation of VVR tunables

*Veritas Volume Replicator Administrator's Guide* Describes how to change tunable values

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

## VEA installation planning

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

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

The VEA server packages include the following:

- The Veritas Volume Replicator Management Services Provider package, `VRTSvrpro`, must be installed on all hosts in the Replicated Data Set (RDS).

- For `VRTSvrpro` to function, the Veritas Volume Manager Management Services Provider package, `VRTSvmpro`, must be installed on each system.
- To use the functionality for receiving SNMP notifications and email notifications, the Veritas Action Agent package, `VRTSaa` must be installed.
- The VEA client can be installed on a host on which VVR is installed, or a separate host that is used to administer the VVR hosts. To use the VEA client on a system, the `VRTSobgui` package must be installed on that system.

## Planning an upgrade from the previous VVR version

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

VVR supports replicating data between VVR 5.0MP3 and VVR 4.1 MP1 or later.

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

---

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

---

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

---

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

---

See [“Upgrading VVR when VCS agents are configured”](#) on page 199.

## Additional settings for using VVR in a localized environment

If the language packages for VVR are installed, VVR displays localized messages, if the client locale is a supported non-English locale. The client locale is the locale from which you are accessing the VVR command line or GUI. For example, if the Japanese version of VVR is installed, then the messages are displayed in the Japanese locale, if the client locale is Japanese.

Make sure that the appropriate locale has been installed on all the hosts that are intended to be a part of the VVR RDS setup. Otherwise, some VVR error messages will be displayed in English, because it is the default locale. Make sure the following settings are done on all hosts that are intended to be part of the RDS:

- Install the required client locale from the Operating System disc.
- Install the required Volume Manager and VVR localized packages. To use VVR VEA, make sure to install the localized package for the VEA client.
- Set the client locale, before using any of the VVR interfaces:
  - for the VVR command line or VVR VEA, set the locale using the appropriate method for your operating system. When you start VVR VEA, the GUI detects and uses the client locale.
  - for VRW, select the locale from the VRW login page.

## Database requirements

The following tables identify 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 2-1** Supported database and Solaris combinations for DB2

Database version	Solaris 8 (64-bit)	Solaris 8 (32-bit)	Solaris 9 (64-bit)	Solaris 9 (32-bit)	Solaris 10 (64-bit)
8.1 ESE with FixPak 6 or lower	Yes	Yes	Yes	Yes	No
8.2 (or 8.1 ESE with FixPak 7 or higher)	Yes	Yes	Yes	Yes	No
8.2.2 with FixPak 9	Yes	Yes	Yes	Yes	Yes
9.1	No	No	No	Yes	Yes

**Table 2-1** Supported database and Solaris combinations for DB2 (*continued*)

Database version	Solaris 8 (64-bit)	Solaris 8 (32-bit)	Solaris 9 (64-bit)	Solaris 9 (32-bit)	Solaris 10 (64-bit)
9.5	No	No	No	Yes	Yes

**Table 2-2** Supported database and Solaris combinations for Oracle

Database version	Solaris 8 (64-bit)	Solaris 8 (32-bit)	Solaris 9 (64-bit)	Solaris 9 (32-bit)	Solaris 10 (64-bit)
9iR2 9.2 (32-bit)	Yes	No	Yes	Yes	Yes
9iR2 9.2 (64-bit)	Yes	Yes	Yes	No	Yes
10gR1 10.1 (64-bit)	Yes	Yes	Yes	No	Yes
10gR2 10.2 (64-bit)			Yes	No	Yes
11gR1 11.1 (64-bit)			Yes	No	Yes

**Table 2-3** Supported database and Solaris combinations for Sybase

Database version	Solaris 8 (64-bit)	Solaris 8 (32-bit)	Solaris 9 (64-bit)	Solaris 9 (32-bit)	Solaris 10 (64-bit)
12.5	Yes	Yes	Yes	Yes	Yes
15	Yes	Yes	Yes	Yes	Yes

**Table 2-4** Supported database and Solaris combinations for Opteron

Database version	Solaris 10 (64-bit)
10gR1 10.1 (64-bit)	Yes
10gR2 10.2 (64-bit)	Yes

## About centralized management

Storage Foundation Manager (SFM) is a free license add-on to Veritas Storage Foundation that provides centralized application, server and storage management capabilities across a heterogeneous infrastructure. SFM is not available on the Storage Foundation and High Availability Solutions release and must be obtained separately.

See [“Downloading Storage Foundation Manager”](#) on page 39.

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

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

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

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

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

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

## Downloading the Storage Foundation and High Availability software

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

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

See [“About installation scripts”](#) on page 297.

### To download the software

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

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

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

See “[Disk space requirements](#)” on page 42.

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

```
# df -b filesystem
```

---

**Caution:** When you select a location to download files, do not select a directory that contains Veritas products from a previous release or maintenance pack. You must download the Veritas 5.0 software and the Veritas 5.0MP3 software into separate directories.

---

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

## Downloading Storage Foundation Manager

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

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

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



# System requirements

This chapter includes the following topics:

- [Software and hardware requirements](#)
- [Supported Solaris operating systems](#)
- [Storage Foundation Cluster File System node requirements](#)
- [Disk space requirements](#)

## Software and hardware requirements

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

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

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

## Supported Solaris operating systems

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

Veritas Storage Foundation is supported on the following Solaris operating systems:

- Solaris 8 (SPARC Platform 32-bit and 64-bit)
- Solaris 9 (SPARC Platform 32-bit and 64-bit)
- Solaris 10 (SPARC or x64 Platform 64-bit)

If necessary, upgrade Solaris before you install the Veritas products.

Install all the latest required Solaris patches listed in the *Veritas Storage Foundation Release Notes*.

For information about the use of this product in a VMware Environment on Solaris x64, refer to <http://entsupport.symantec.com/docs/289033>

## Storage Foundation Cluster File System node requirements

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

## Disk space requirements

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

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

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

## Disk space

[Table 3-1](#) shows the approximate disk space used by the Storage Foundation products for all (both the required and optional) packages.

**Table 3-1** 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

## Disk space requirements for Veritas Volume Replicator

[Table 3-2](#) shows the approximate disk space used by VVR for the required and optional packages.

**Table 3-2** Approximate disk space use for VVR

English	/root	/opt	/usr	/var
Required Packages	170 MB	60 MB	85 MB	0.5 MB
Optional Packages	240 MB	256 MB	85 MB	0.5 MB
All Packages	410 MB	316 MB	170 MB	1 MB

[Table 3-3](#) shows the approximate disk space used by VVR for the localized packages.

**Table 3-3** Approximate disk space use for localized packages

Localized	/(root)	/opt	/usr	/var
All Packages	41 MB	37 MB	6 MB	1 MB



# Installing Storage Foundation using the common product installer

This chapter includes the following topics:

- [Installation quick reference](#)
- [Mounting a software disc](#)
- [About the common product installer](#)
- [Installing and configuring Storage Foundation using the common product installer](#)
- [Installing and configuring Storage Foundation and High Availability Solutions using the common product installer](#)
- [Installing and configuring Storage Foundation Cluster File System using the common product installer](#)
- [Installing and configuring Veritas Volume Replicator using the common product installer](#)
- [Installing language packages](#)
- [Installing Veritas Enterprise Administrator](#)

## Installation quick reference

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

[Table 4-1](#) provides a quick overview of a stand-alone installation using the product installer.

**Table 4-1** Installation overview

Installation task	For more information, refer to the following section:
Obtain product licenses.	See <a href="#">“Symantec product licensing”</a> on page 22.
Download the software, or insert the product DVD.	See <a href="#">“Downloading the Storage Foundation and High Availability software”</a> on page 38. See <a href="#">“Mounting a software disc”</a> on page 46.
Set environment variables.	See <a href="#">“Setting environment variables”</a> on page 23.
Create the <code>/opt</code> directory, if it does not exist.	See <a href="#">“Creating the /opt directory”</a> on page 23.
Configure the secure shell (SSH) on all nodes.	See <a href="#">“Configuring secure shell (ssh) or remote shell before installing products”</a> on page 24.
Verify that hardware, software, and operating system requirements are met.	See <a href="#">“Software and hardware requirements”</a> on page 41.
Check that sufficient disk space is available.	See <a href="#">“Disk space requirements”</a> on page 42.
Use the installer to install the products.	See <a href="#">“Installing and configuring Storage Foundation using the common product installer”</a> on page 48.

## Mounting a software disc

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

### To mount the software disc

- 1 Log in as superuser.
- 2 Place the Veritas software disc 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. Insert the disc and enter the following command:  

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom/cdrom0
```

where `c0t6d0s2` is the default address for the disc drive.
- 5 Change to the appropriate directory and product subdirectory to view the product release notes and installation guides, or install the products.

## About the common product installer

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

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

See [“About installation scripts”](#) on page 297.

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

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

Additional options are available for the common product installer.

See [“Installation script options”](#) on page 298.

Install Veritas products in one of the following ways:

- For new installations or upgrades from releases prior to 5.0:
  - Use the product installer, the `installer` script  
The product installer displays a menu that simplifies the selection of installation and upgrade options. It is the recommended installation method. Selecting a product to install or upgrade from the menu invokes that product's installation script.  
See [“Installing and configuring Storage Foundation using the common product installer”](#) on page 48.
  - Use the product's installation script, such as `installsf` or `installsfdfs`.

See [“About installation scripts”](#) on page 297.

- Use custom methods

See [“Installing Storage Foundation - summary”](#) on page 79.

- For upgrades from 5.0 releases, including maintenance packs and rolling patches:

Use the `installmp` script

See [“Upgrading Storage Foundation software from 5.0 to 5.0MP3 using the product installer”](#) on page 159.

## Installing and configuring Storage Foundation using the common product installer

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

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

### To install Storage Foundation

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

- 2 Load and mount the software disc.

See [“Mounting a software disc”](#) on page 46.

- 3 Move to the top-level directory on the disc.

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

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

```
# ./installer
```

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

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

The sample installation assumes that ssh is used.

- 5 Enter **I** to install and press Return.
- 6 When the list of available products is displayed, select Veritas Storage Foundation, enter the corresponding number, and press Return.

Veritas Storage Foundation for Oracle, Veritas Storage Foundation for DB2 (on SPARC systems only), and Veritas Storage Foundation for Sybase can also be installed using this procedure. Select the number corresponding to one of those products, if desired.

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

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

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

- 8 Enter the product license information.

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

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

- 9 You are prompted to enter additional license information, until all licenses for all systems have been entered. Then reply that you have no additional licenses to enter.

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

- 10** You can choose to install required packages or all packages. Optional packages include man pages, for example.

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

For example, you should see output similar to the following:

```
SF can be installed without optional packages to conserve disk space.  
  
1) Install required Veritas Storage Foundation packages - 883 MB required  
2) Install all Veritas Storage Foundation packages - 911 MB required  
  
Select the packages to be installed on all systems? [1-2,q,?] (2) 1
```

- 11** The AP driver provides support for alternate paths, and the installation confirms availability of this Sun AP driver. Press Return to continue.

```
Checking for any AP driver issues on host1 ..... None
```

- 12** Configure Storage Foundation when prompted.

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

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

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

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

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

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

See [“Setting the default disk group”](#) on page 214.

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

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

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

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

- 16** You are prompted to confirm the default disk group.

---

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

---

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

- 17** Verify the fully qualified hostname of the systems.

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

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

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

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

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

**19** The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

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

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

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

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

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

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

### To install Storage Foundation and High Availability products

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

**2** Load and mount the software disc.

See [“Mounting a software disc”](#) on page 46.

**3** Move to the top-level directory on the disc.

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

- 4 From this directory, type the following command to install on the systems, if you use the ssh utilities:

```
# ./installer
```

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

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

The sample installation assumes that ssh is used.

- 5 Enter `I` to install and press Return.
- 6 When the list of available products is displayed, select Veritas Storage Foundation (SF), enter the corresponding number, and press Return.

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

Veritas Storage Foundation for Oracle, Veritas Storage Foundation for DB2 (on SPARC systems only), and Veritas Storage Foundation for Sybase can also be installed using this procedure. Select the number corresponding to one of those products, if desired.

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

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

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

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

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

**9** Enter the product license information.

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

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

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

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

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

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

For example, you should see output similar to the following:

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

- 1) Required Veritas Storage Foundation packages - 982 MB required
- 2) All Veritas Storage Foundation packages - 1071 MB required

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

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

**11** The AP driver provides support for alternate paths (AP), and the installation confirms availability of this Sun AP driver.

```
Checking for any AP driver issues on host1 ..... None
```

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

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

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

All systems are configured to create one cluster.

**13 Enter the unique cluster name and Cluster ID number.**

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

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

```
Discovering NICs on host1 ... discovered bge0 bge1 bge2 bge3
```

**15 Enter private heartbeat NIC information for each host.**

```
Enter the NIC for the first private heartbeat link
on host1: [b,?] bge1
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,?] bge2

Would you like to configure a third private heartbeat
link? [y,n,q,b,?] (n) n
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) n
Are you using the same NICs for private heartbeat links
on all systems? [y,n,q,b,?] (y) y
```

---

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

---

Notice that in this example, **bge0** is not selected for use as a private heartbeat NIC because it already in use as the public network interface.

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

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

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

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

---

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

---

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

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

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

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

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

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

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

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

Enter the SNMP trap daemon port: [b,?] (162) 162
Enter the SNMP console system name: [b,?] host1
Enter the minimum severity of events for which SNMP traps should
be sent to host1 [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
Would you like to add another SNMP console? [y,n,q,b] (n) n
```

- 21** Verify and confirm that the information is correct, by entering **y**, or enter the information again.
- 22** Choose whether you want to use enclosure-based naming. The enclosure-based naming scheme (rather than disk device naming) is a feature of Veritas Volume Manager. You can reference disks using a symbolic name that is more meaningful than the operating system's normal device access name.

See the *Veritas Volume Manager Administrator's Guide*

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

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

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

See "[Setting the default disk group](#)" on page 214.

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

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

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

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

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

- 25** You are prompted to confirm the default disk group.

---

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

---

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

- 26** Validate the default disk group information, and press Return.
- 27** You may be prompted to verify the fully qualified hostnames of the systems. Verify them and press Return to continue.
- 28** This product can be configured as a Storage Foundation Manager managed host or a stand-alone host.

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

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

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

- 29** The Veritas Storage Foundation software is verified and configured.  
Start the Veritas Storage Foundation processes.

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

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

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

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

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

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

See [“Creating and configuring the repository database for DB2 and Oracle”](#) on page 111.

- 32 Reboot the system (or systems).

```
# /usr/sbin/shutdown -y -i6 -g0
```

## Installing and configuring Storage Foundation Cluster File System using the common product installer

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

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

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

### To install the Storage Foundation Cluster File System

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

- 2 Load and mount the software disc.

See [“Mounting a software disc”](#) on page 46.

- 3 Move to the top-level directory on the disc.

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

- 4 From this directory, type the following command to install if you are using the secure shell (ssh) utilities:

```
# ./installer
```

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

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

The sample installation assumes that ssh is used.

- 5 From the Installation menu, choose the `I` option for Install and enter the number for Veritas Storage Foundation Cluster File System. Press Return.  
Do not select the "Storage Foundation for Oracle RAC" option unless you have the correct license and setup.
- 6 You are prompted to enter one or more system names to install SFCFS.

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

- 7 During the initial system check, the installer verifies that communication between systems has been set up.  
If the installer hangs or asks for a login password, stop the installer and set up ssh or rsh. Then run the installer again.  
See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.
- 8 Enter the product license information.

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

---

**Note:** You must have a Storage Foundation HA license to access the Symantec Security services. You can still enable the Security services after the installation by executing `/opt/VRTS/install/installsfcfs -security host1 host2 ... hostn`

---

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

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

```
Do you want to enter another license key for system02?  
[y,n,q] (n) n
```

- 10** You can choose to install required packages or all packages.

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

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

- 12** The AP driver provides alternate pathing support, and the installation confirms availability of this Sun AP driver.

```
Checking for any AP driver issues on host1 ..... None
```

- 13** Configure Storage Foundation Cluster File System (SFCFS) when prompted.

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

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

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

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

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

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

- 15** Configure the cluster. No configuration changes are made to the systems until all configuration questions are completed and confirmed. Press Return to continue.

All systems are configured to create one cluster.

Enter the unique cluster name and Cluster ID number.

```
Enter the unique cluster name: [?] cluster2
```

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

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

```
Discovering NICs on host1 ...discovered bge0 bge1 bge2 bge3
```

- 17** Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat link
on host1: [b,?] bge1
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,?] bge2

Would you like to configure a third private heartbeat
link? [y,n,q,b,?] (n) n
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) n
Are you using the same NICs for private heartbeat links
on all systems? [y,n,q,b,?] (y) y
```

---

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

---

Notice that in this example, *bge0* is not selected for use as a private heartbeat NIC because it already in use as the public network interface.

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

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

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

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

---

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

---

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

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

See the *Veritas Volume Manager Administrator's Guide*

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

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

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

See [“Setting the default disk group”](#) on page 214.

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

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

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

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

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

- 23** You are prompted to confirm the default disk group.

---

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

---

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

- 24** Validate the default disk group information, and press Return.
- 25** You may be prompted to verify the fully qualified hostname of the systems.

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

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

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

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

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

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

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

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

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

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

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

- 29** Reboot the system (or systems).

```
# /usr/sbin/shutdown -y -i6 -g0
```

- 30 If you do not plan to use SFCFS file systems to store the Oracle database or binaries, you have completed the SFCFS installation and configuration.
- 31 Before installing Oracle binaries (`ORACLE_HOME`), consider these points:
  - Local installations provide a comfort level using traditional installation methods and the possibility of improved protection against a single point of failure.
  - CFS installations provide a single Oracle installation to manage, regardless of number of nodes. This scenario offers a necessary reduction in storage requirements and easy addition of nodes.

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

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

## Installing and configuring Veritas Volume Replicator using the common product installer

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

### To install VVR using the product installer

- 1 To install on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.  
See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.
- 2 Load and mount the software disc.  
See [“Mounting a software disc”](#) on page 46.
- 3 Move to the top-level directory on the disc.

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

- 4 From this directory, type the following command to install on the systems, if you use the ssh utilities:

```
# ./installer
```

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

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

The sample installation assumes that ssh is used.

- 5 Enter `I` to install and press Return.
- 6 When the list of available products is displayed, select Veritas Volume Replicator, enter the corresponding number, and press Return.  
To install Veritas Storage Foundation for Oracle, do not select the "Storage Foundation for Oracle RAC packages" option unless you have the correct license or setup.
- 7 You are prompted to enter the system names (in the following example, "system01" and "system02") on which the software is to be installed. Enter the system name or names and then press Return.

```
Enter the system names separated by spaces on which to  
install SF: system01 system02
```

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

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

**9 Enter the product license information.**

```

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

```

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

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

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

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

For example, you should see output similar to the following:

Additional packages are typically installed to simplify future upgrades

- 1) Required Veritas Volume Replicator packages - 855 MB required
- 2) All Veritas Volume Replicator packages - 861 MB required
- 3) Storage Foundation Enterprise packages - 911 MB required

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

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

**11 The AP driver provides support for alternate paths, and the installation confirms availability of this Sun AP driver. Press Return to continue.**

```
Checking for any AP driver issues on system01 ..... None
```

- 12** Configure Storage Foundation and VVR when prompted. Configuring VVR also involves the configuration of Storage Foundation components.

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

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

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

See the *Veritas Volume Manager Administrator's Guide*

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

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

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

See “[Setting the default disk group](#)” on page 214.

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

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

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

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

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

- 16** You are prompted to confirm the default disk group.

---

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

---

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

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

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

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

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

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

Change the frequency of online statistics collection, if needed.

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

- 19** Change the maximum number of days that online statistics are retained, if needed.

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

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

- 20** Configure the VVR statistics options (tunables), if needed.

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

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

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

- 22** Verify the fully qualified hostnames of the systems. Press Return to continue.

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

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

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

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

- 24** To start the VVR processes, press Return, or type *y*.

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

- 25** The installation and configuration complete automatically, and the processes are started.

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

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

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

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

## Installing VVR when VxVM is already installed

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

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

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

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

See [“Installing and configuring Veritas Volume Replicator using the common product installer”](#) on page 65.

To use the new features of VVR 5.0MP3, upgrade the version of each disk group.

See [“Upgrading VxVM disk group versions”](#) on page 213.

## Installing language packages

To install a Veritas Storage Foundation product in a language other than English, install the required language packages after installing the English packages. Veritas Storage Foundation 5.0MP3 supports simplified Chinese, French, and Japanese languages. Follow the procedure for the appropriate language packages.

See [“Installing the simplified Chinese and French language packages”](#) on page 71.

See [“Installing the Japanese language packages”](#) on page 73.

## Installing the simplified Chinese and French language packages

Use the following procedure to install the simplified Chinese or French language packages on the server.

**To install the simplified Chinese or French 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. With 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` directory to a temporary directory on your system, such as `/tmp/pkgs`:

```
# cp /cdrom/cdrom0/fr/volume_manager/pkgs \  
    /tmp/pkgs
```

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

```
# cp /cdrom/cdrom0/zh/volume_manager/pkgs \  
    /tmp/pkgs
```

- 5 Decompress the package, 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 Use the `patchadd` command to install the 5.0MP3 patches for Chinese or French.

For Chinese, the required patches are located in the following directory:

```
/cdrom/cdrom0/zh
```

For French, the required patches are located in the following directory:

```
/cdrom/cdrom0/fr
```

- 8 Restart the VEA Service.

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

## Installing the Japanese language packages

Use the following procedure to install the Japanese language packages on the server.

### To install the Japanese language packages on the server

- 1 Make sure the VEA Service is not running.

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

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

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

- 3 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With 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 Install the language patches for 5.0MP3 with the `install_mlp` command.

```
# ./install_mlp
```

- 6 Restart the VEA Service.

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

## Installing the VEA language packages

If you are using the VEA GUI, the language packages must also be installed.

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

- 1 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.
```

- 2 For the French language VEA client package, copy the following package to a temporary directory on your system, such as `/tmp/pkgs`:

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

For the Simplified Chinese language VEA client package, copy the following package to a temporary directory on your system, such as `/tmp/pkgs`:

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

For the Japanese language VEA client package, copy the following package to a temporary directory on your system, such as `/tmp/pkgs`:

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

- 3 Decompress the package, and extract the contents.

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

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

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

- 5 Add the 5.0MP3 patch for the `VRTSmuobg` package:

```
# patchadd -d . V123073-xx
```

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

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

## Installing Veritas Enterprise Administrator

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

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

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

### Installing the Veritas Enterprise Administrator client

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

The 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.

The VEA client requires one of the following packages:

- Veritas Enterprise Administrator client package (`VRTSobgui`)  
This is the client package for UNIX.
- 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.

## Minimum system requirements for VEA clients

Table 4-2 shows the system minimum requirements for the GUI.

Table 4-2 VEA system minimum requirements

Operating System	System minimum requirements
Solaris	SPARCstation 5 with 64M memory
Windows XP, NT, Me, 2000, or 98	300MHz Pentium with at least 256MB of 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.

## Installing the Veritas Enterprise Administrator client on Solaris

If you plan to run the VEA client, you must install the `VRTSobgui` package on the machine you are planning to use.

By default, the VEA client is 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.

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

- 1 Log in as superuser (`root`).
- 2 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 to the current working directory on your system.

```
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSobgui.tar.gz .
```

- 5 Decompress the package, and then extract the contents.

```
# /cdrom/cdrom0/gnu/gunzip *.tar.gz  
# tar xvf VRTSobgui.tar
```

- 6 Use the `pkgadd` command to install the package. Answer any questions, as the installation proceeds.

```
# pkgadd -d . VRTSobgui
```

The VEA client package for Solaris is installed.

## Installing the VEA client on Microsoft Windows

The VEA client package for Microsoft Windows systems (`win32\VRTSobgui.msi`) is used to access servers running VxVM and the VEA service on Solaris. 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.

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

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

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

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

<http://www.microsoft.com>

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

**To install the VEA client on a Windows machine**

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

# Installing Storage Foundation, other methods

This chapter includes the following topics:

- [Installing Storage Foundation - summary](#)
- [Installing the Storage Foundation products using JumpStart](#)
- [Installing Storage Foundation using the pkgadd command](#)

## Installing Storage Foundation - summary

This section describes procedures to install Storage Foundation or Storage Foundation Cluster File System, as alternative to using the common product installer.

The recommended method to install is to use the common product installer.

See [“About the common product installer”](#) on page 47.

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

See the *Veritas Storage Foundation Management Server Installation Guide*.

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.

[Table 5-1](#) lists the alternative methods that are possible for installing Storage Foundation.

**Table 5-1** Installation methods

Method	See
Installation using pkgadd	See <a href="#">“Installing Storage Foundation using the pkgadd command”</a> on page 88.
Installation using JumpStart	See <a href="#">“Installing the Storage Foundation products using JumpStart”</a> on page 80.

## Installing the Storage Foundation products using JumpStart

These installation instructions using JumpStart assume a working knowledge of the JumpStart procedure. See the JumpStart documentation that came with your operating system for details.

Only fresh installs of Storage Foundation products are supported using JumpStart. Upgrading is not supported.

The following procedure assumes a stand-alone configuration.

### Summary of JumpStart installation tasks

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

#### To install the Veritas packages on a JumpStart server

- 1 Add a client (register to the JumpStart server).  
See the JumpStart documentation that came with your operating system for details.
- 2 Copy the compressed Veritas package files to a temporary directory, uncompress the files, and extract the packages from the tar files.  
See [“Copying and unzipping the packages”](#) on page 82.
- 3 Read the instructions for the installation procedure using JumpStart.
- 4 Copy the packages to a JumpStart server under a shared directory on the network. Note the subdirectory with the relevant packages for your installation.
- 5 Determine the installation order.  
See [“Determining the installation order”](#) on page 84.

**6** Modify the rules file for JumpStart.

See the JumpStart documentation that came with your operating system for details.

**7** Write the JumpStart start and finish scripts.

The `pkgadd` operations to install the packages must be coded in a script that can be used with the JumpStart server.

You can run the installer with the `jumpstart` option to create a sample finish file.

To create a sample finish file for Storage Foundation High Availability:

```
# cd storage_foundation
# ./installsf -jumpstart dir_path
```

Use the `-osversion` option to create a finish file that includes only packages for a specific Solaris version. For example, to create a finish file for Storage Foundation High Availability on Solaris 9:

```
# cd storage_foundation
# ./installsf -jumpstart dir_path -osversion sol9
```

To create a sample finish file for Storage Foundation only:

```
# cd storage_foundation
# ./installsf -jumpstart -nohapkgs dir_path
```

The `dir_path` indicates the path to the directory in which to create the finish file.

**8** Add packages to the appropriate location and set up the JumpStart environment.

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.

**9** Create the response files for the VRTS packages.

See [“Creating the response files for JumpStart”](#) on page 86.

**10** Run JumpStart to install the packages.

- 11 JumpStart may automatically reboot the system after the packages have been installed.
- 12 Run the installer command from the disc directory to configure the Veritas software.

```
# /cdrom/cdrom0/installer -configure
```

## Copying and unzipping the packages

Before you can install the packages, you must unzip them, and extract them from the tar file. This procedure is only required when you install with a manual procedure, such as using JumpStart. This procedure is not required if you install with the common product installer or the product installation scripts.

### To unzip the packages

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

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

This directory must be clean, with no files present from previous installations or versions of any products.

- 3 Insert the disc with the Veritas Storage Foundation software into a drive that is connected to the system.

The Solaris volume management software automatically mounts the disc as `/cdrom/cdrom0`. Type the command:

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

Alternatively, you can download the software from the Symantec Web site.

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

If you download the software, you need to unzip and untar the downloaded file to the temporary directory.

For Veritas Storage Foundation, Veritas Storage Foundation for DB2, Veritas Storage Foundation for Oracle, or Veritas Storage Foundation for Sybase:

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

For Veritas Volume Manager:

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

For Veritas File System:

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

For Veritas Cluster File System:

```
# cp -r storage_foundation_cluster_file_system/pkgs/* \
  /parent_directory/install/pkgs
# cp -r storage_foundation_cluster_file_system/patches/* \
  /parent_directory/install/patches
```

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

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

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

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

**7** Decompress and extract each package.

```
# cd /parent_directory/install/pkg  
# tar xf package_name.tar  
# tar xf package_name.tar  
# tar xf package_name.tar  
.  
.
```

**8** Decompress and extract each patch.

```
# cd /parent_directory/install/patches  
# tar xf patch_name.tar  
# tar xf patch_name.tar  
# tar xf patch_name.tar  
.  
.
```

**9** List the files in the temporary directory.

```
# ls /parent_directory/install/pkg  
# ls /parent_directory/install/patches
```

**10** Use these directories to provide the packages and patches for the manual installation procedure.

## Determining the 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.

The list of the available packages has descriptions of each package.

See [“Veritas Storage Foundation installation packages”](#) on page 303.

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

The `requiredpkgs` option displays only the required packages and the `installpkgs` displays all packages.

### To get package installation order for Veritas Storage Foundation

- 1 Move to the Storage Foundation installation directory. For the disc, use the following command:

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

- 2 Run the script with the `-requiredpkgs` option or the `-installpkgs` option.

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

or

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

### To get package installation order for Veritas Storage Foundation HA

- 1 Move to the Storage Foundation installation directory. For the disc, use the following command:

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

- 2 Run the script with the `-requiredpkgs` option or the `-installpkgs` option.

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

or

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

### To get package installation order for Veritas Volume Manager

- 1 Move to the Volume Manager installation directory. For the disc, use the following command:

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

- 2 Run the script with the `-requiredpkgs` option or the `-installpkgs` option.

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

or

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

### To get package installation order for Veritas File System

- 1 Move to the File System installation directory. For the disc, use the following command:

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

- 2 Run the script with the `-requiredpkgs` option or the `-installpkgs` option.

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

or

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

### To get package installation order for Veritas Cluster File System

- 1 Move to the Veritas Cluster File System installation directory. For the disc, use the following command:

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

- 2 Run the script with the `-requiredpkgs` option or the `-installpkgs` option.

```
# ./installsfcfs -requiredpkgs
```

or

```
# ./installsfcfs -installpkgs
```

## Creating the response files for JumpStart

Use the following instructions to create the response files for JumpStart and add the relevant parameters to your finish file.

### To create the response files for JumpStart

- 1 The VRTSdcli, VRTSjre15, VRTSvxvm, and VRTSweb packages require an empty response file.

You can add the following lines to your scripts to create the empty response files before the relevant `pkgadd` command.

```
# touch responsefile
# pkgadd -r responsefile package_name
```

- 2 Copy the supplied VRTS admin files from the installation media, and modify them if needed.

```
# cp storage_foundation/scripts/VRTS* \  
    /tmp/pkgs
```

- 3 The VRTSobc33 response file must contain the following settings:

```
VXPALLOCALHOSTNAME=  
VXPALDOMAINCONTROLLER=  
VXPALAUTHENTICATIONBROKER=  
VXPALSECURITY=y  
HOSTCONTROLLER=1  
AG_PASSWORD=
```

- 4 The VRTSob response file must contain the following settings:

```
SECURITYADMINPASSWORD=  
ROOTAGENTPASSWORD=  
CONFIGURESECURITY=n
```

- 5 The VRTScssim response files must contain the following settings:

```
BASEDIR=/opt
```

- 6 The `VRTSvdiid` response files must contain the following settings:

```
OPT=/opt
```

- 7 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
```

## Installing Storage Foundation using the pkgadd command

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 stand-alone host. The system can be converted later to a Storage Foundation Manager managed host.

For information about obtaining and installing the SF Manager, refer to the *Veritas Storage Foundation Manager Installation Guide*.

### To install Storage Foundation using the pkgadd command

- 1 Mount the software disc.

See [“Mounting a software disc”](#) on page 46.

- 2 Copy the packages to a location to which you can write and then uncompress and untar the packages:

```
# cp -R /cdrom/cdrom0/storage_foundation/pkgs/* /tmp/pkgs
```

- 3 Uncompress and extract the packages by using the `gzcat` command:

```
# cd /tmp/pkgs
# gzcat *.gz | tar xvf -
```

- 4 Copy the supplied `VRTSobcadmin` and `VRTSobadmin` files from the installation media, and modify them if needed.

```
# cp /cdrom/cdrom0/storage_foundation/scripts/VRTS* \
/tmp/pkgs
```

- 5 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
```

- 6 Use the product installation script to determine the list of packages and patches, and the order in which they should be installed.

For example, to install all of the packages for Storage Foundation, use the following command:

```
./installsf -installpkgs
```

See [“Determining the installation order”](#) on page 84.

- 7 Install the packages listed in step 6.

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.

**8** Verify that each of the packages is installed:

```
# pkginfo -l packagename
```

## 9 Install the patches for Storage Foundation 5.0MP3 with the `patchadd` command.

The following patches are available for Solaris Sparc:

PATCH	PACKAGE	OS VERSION
127333-01	VRTSvlic	ALL
137338-01	VRTSspb	ALL
123722-01	VRTSat	ALL
122632-21	VRTSobc33	ALL
122631-21	VRTSob	ALL
122633-21	VRTSobgui	ALL
123076-21	VRTSccg	ALL
123079-21	VRTSmh	ALL
123075-21	VRTSaa	ALL
123200-03	VRTSvxfs	5.8
123201-03	VRTSvxfs	5.9
123202-03	VRTSvxfs	5.10
122058-10	VRTSvxvm	ALL
121714-03	VRTSfspro	ALL
123818-02	VRTSvmman	ALL
123740-03	VRTSvmpro	ALL
123742-05	VRTSdcli	ALL
123821-03	VRTSalloc	ALL
123823-03	VRTSddlpr	ALL
123743-03	VRTSvrpro	ALL
125150-07	VRTSjre15	ALL
123744-03	VRTSvrw	ALL
128078-01	VRTSfsman	ALL
137329-01	VRTSfssdk	ALL
124004-02	VRTSvail	ALL
137385-01	VRTSvxmsa	ALL
123220-04	VRTSmapro	ALL

The following patches are available for Solaris Opteron:

PATCH	PACKAGE	OS VERSION
127333-01	VRTSvlic	ALL
137339-01	VRTSspb	5.10
125862-21	VRTSobc33	ALL
125861-21	VRTSob	ALL
125863-21	VRTSobgui	ALL
125865-21	VRTSccg	ALL

125866-21	VRTSmh	ALL
125864-21	VRTSaa	ALL
127337-01	VRTSvxfs	5.10
127336-01	VRTSvxvm	5.10
127342-01	VRTSfspro	5.10
128079-01	VRTSvmman	ALL
127363-01	VRTSvmpro	ALL
128051-01	VRTSdcli	ALL
127361-01	VRTSalloc	ALL
127362-01	VRTSddlpr	ALL
127324-01	VRTSvrpro	ALL
128091-01	VRTSvcsvr	5.10
137384-03	VRTSjre15	5.10
128080-01	VRTSfsman	5.10
137330-01	VRTSfssdk	5.10
137386-01	VRTSvxmsa	5.10

**10** Start the VEA server:

```
# /opt/VRTSob/bin/vxsvcctl start
```

**11** Register the agents:

```
# /opt/VRTSobc/pal33/bin/install/cfgsecurity.sh \  
-a StorageAgent  
# /opt/VRTSaa/config/setup_vxaa.sh --nocpi  
# /opt/VRTSccg/config/setup.sh --gridnode
```

**12** Run the `sfinstall -configure` command to configure.

**13** Start the agents:

```
# /opt/VRTSobc/pal33/bin/vxpalctl -a StorageAgent -c start  
# /opt/VRTSobc/pal33/bin/vxpalctl -a gridnode -c start
```

**14** If the configuration is standalone, also start the action agent:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c start
```

**15** Configure the license manager agent:

```
# /opt/SYMCлма/bin/lmautil --Config --SecurityEnabled 1 \  
  --RootBrokerHostname "1" --CollectorNodeUsername "2" \  
  --CollectorNodeUserDomainType "3" \  
  --CollectorNodeUserDomain "4"
```



# Configuring Storage Foundation and High Availability products

This chapter includes the following topics:

- [Configuring the Storage Foundation products](#)
- [Configuring Storage Foundation](#)
- [Configuring Storage Foundation and High Availability Solutions](#)
- [Configuring Storage Foundation for Databases](#)
- [Configuring Veritas Volume Manager](#)
- [Configuring Veritas File System](#)
- [Configuring Storage Foundation Cluster File System](#)
- [Configuring Veritas Volume Replicator](#)
- [Veritas Volume Replicator and Volume Manager setup after installation](#)
- [Configuring and starting Veritas Enterprise Administrator](#)
- [Configuring Veritas Enterprise Administrator for databases](#)

## Configuring the Storage Foundation products

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

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

See “[Configuring Storage Foundation for Databases](#)” on page 106.

If the Storage Foundation products were installed with another method, they may also need to be configured. Review the configuration sections that are appropriate for the Storage Foundation products that were installed. Follow the needed procedures.

## Configuring Storage Foundation

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

### To configure Storage Foundation

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

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

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

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

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

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

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

```
Checking system licensing
```

```
SF license registered on host1
```

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

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

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

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

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

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

See “[Setting the default disk group](#)” on page 214.

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

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

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

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

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

- 8 You are prompted to confirm the default disk group.

---

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

---

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

- 9 Verify the fully qualified hostname of the systems.

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

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

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

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

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

- 11** The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

- 12** The configuration completes automatically.

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

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

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

## Configuring Storage Foundation and High Availability Solutions

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

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

### Required information for configuring Storage Foundation and High Availability Solutions

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

See also the *Veritas Cluster Server Installation Guide*.

- A unique Cluster name

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

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

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

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

See the *Veritas Cluster Server Installation Guide* for more information on configuring a secure cluster.

The following information is required to configure SMTP notification:

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

The following information is required to configure SNMP notification:

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

## Configuring Veritas Storage Foundation and High Availability Solutions

After installation, you must configure the product.

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

### To configure Storage Foundation product on a cluster

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

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

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

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

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

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

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

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

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

```
Checking system licensing
```

```
SF license registered on host1
```

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

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

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

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

All systems are configured to create one cluster.

**7** Enter the unique cluster name and Cluster ID number.

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

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

```
Discovering NICs on host1 ... discovered bge0 bge1 bge2 bge3
```

**9** Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat link
on host1: [b,?] bge1
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,?] bge2
```

```
Would you like to configure a third private heartbeat
link? [y,n,q,b,?] (n) n
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) n
Are you using the same NICs for private heartbeat links
on all systems? [y,n,q,b,?] (y) y
```

---

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

---

Notice that in this example, `bge0` is not selected for use as a private heartbeat NIC because it already in use as the public network interface.

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

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

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

- 11** When prompted to configure the product to use Veritas Security Services, enter *n*, unless a Root Broker has already been set up.

---

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

---

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

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

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

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

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

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

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

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

- 14** You are prompted to enter the user information.

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

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

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

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

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

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

Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,?] smtp.mycompany.com
Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,?] user@mycompany.com
Enter the minimum severity of events for which mail should be sent
to user@163.com [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
```

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

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

Enter the SNMP trap daemon port: [b,?] (162) 162
Enter the SNMP console system name: [b,?] host1
Enter the minimum severity of events for which SNMP traps should
be sent to host1 [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
Would you like to add another SNMP console? [y,n,q,b] (n) n
```

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

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

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

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

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

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

- 20** Verify and confirm the configuration of the global cluster.

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

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

See the *Veritas Volume Manager Administrator's Guide*

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

- 22** You are now given the option of specifying the default name of a disk group that is to be assumed by Veritas Volume Manager commands if a disk group is not otherwise specified.

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

See “[Setting the default disk group](#)” on page 214.

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

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

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

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

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

- 24** Validate the default disk group information, and press Return.
- 25** You may be prompted to verify the fully qualified hostname of the systems. Press Return to continue.
- 26** This product can be configured as a Storage Foundation Manager managed host or a stand-alone host.

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

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

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

- 27** The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

**28** The configuration and startup complete automatically.

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

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

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

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

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

See [“Creating and configuring the repository database for DB2 and Oracle”](#) on page 111.

## About adding and removing nodes in a cluster

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

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

## Configuring Storage Foundation for Databases

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

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

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

See [“Configuring Storage Foundation and High Availability Solutions”](#) on page 98.

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

See [“Creating and configuring the repository database for DB2 and Oracle”](#) on page 111.

**To configure Storage Foundation product**

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

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

- 2 When the list of available products is displayed, select the number corresponding to the product you want to configure, and press Return.

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

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

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

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

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

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

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

```
Checking system licensing
```

```
SF license registered on host1
```

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

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

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

Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,?] smtp.mycompany.com
Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,?] user@mycompany.com
Enter the minimum severity of events for which mail should be sent
to user@163.com [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
```

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

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

Enter the SNMP trap daemon port: [b,?] (162) 162
Enter the SNMP console system name: [b,?] host1
Enter the minimum severity of events for which SNMP traps should
be sent to host1 [I=Information, W=Warning, E=Error,
S=SevereError]: [b,?] E
Would you like to add another SNMP console? [y,n,q,b] (n) n
```

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

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

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

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

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

See the *Veritas Volume Manager Administrator's Guide*

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

- 10** You are now given the option of specifying the default name of a disk group that is to be assumed by Veritas Volume Manager commands if a disk group is not otherwise specified.

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

See “[Setting the default disk group](#)” on page 214.

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

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

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

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

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

- 12** Validate the default disk group information, and press Return.
- 13** You may be prompted to verify the fully qualified hostname of the systems. Press Return to continue.
- 14** This product can be configured as a Storage Foundation Manager managed host or a stand-alone host.

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

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

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

- 15** The Veritas Storage Foundation for databases software is verified and configured.

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

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

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

- 16** The configuration and startup complete automatically.

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

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

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

- 17 The installation script prompts for a reboot if there are one or more errors. Reboot the system (or systems) if the install script prompts you to do so.
- 18 If you installed Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, create a new repository database.  
See [“Creating and configuring the repository database for DB2 and Oracle”](#) on page 111.

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

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

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

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

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

- Create a separate, non-shared disk group on shared storage. Create a VxVM volume and a VxFS file system and mount the volume.
- It is recommended that you have a separate disk group for the repository volume so that any failovers are independent of other service groups.
- The mount point is used to store the repository database.
- Obtain an unique virtual IP address for public NIC interface.

- Obtain the device names for the public NIC interface for all systems in the cluster.  
 For example, use these names.  
     hme0  
     bge0
- Obtain a subnet mask for the public NIC interface.
- Make sure VCS is not in read-write (-rw) mode. To make sure VCS is in read-only mode, use the following command:

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

Table 6-1 indicates the options available for the `sfua_db_config` script.

**Table 6-1** sfua\_db\_config options

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

**To create and configure the repository database**

- 1** Run the `sfua_db_config` script as follows:

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

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

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

- 3** The mount point is displayed.

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

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

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

```
[bge0]
```

```
Enter the NIC for system host2 for HA Repository configuration:
```

```
[bge0]
```

- 5** Enter the Virtual IP address for repository failover.

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

```
xxx.xxx.xxx.xxx
```

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

```
Following information will be used for SFORA HA configuration:
```

```
Public IP address: xxx.xxx.xxx.xxx
```

```
Subnet mask: xxx.xxx.xxx.xxx
```

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

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

**7** Verify that the repository was configured.

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

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

Note: Sfua\_Base group should be online on one node in the cluster.

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

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

## Setting administrative permissions for databases

To allow database administrators to administer a database using Veritas Storage Foundation, you are required to change some permission settings. 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.

# Configuring Veritas Volume Manager

Use the following procedures to configure Veritas Volume Manager. If you have installed and configured VxVM using the product installer, you do not need to complete the procedures in this section.

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

To carry out further tasks such as disk encapsulation or initialization, please see the *Veritas Volume Manager Administrator's Guide*.

In releases of VxVM (Volume Manager) prior to 4.0, a system installed with VxVM 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, VxVM can function without any disk group having been configured. Only when the first disk is placed under 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.

## 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.

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 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:

```
# vxctl 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 `vxctl(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
```

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.

## 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 `vxlist` command on each node to display the shared disk groups.

## Using `vxinstall` to configure Veritas Volume Manager

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 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, refer to the support section.

See “[Supported Solaris operating systems](#)” on page 41.

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 VxVM 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 VxVM 4.0, the default disk group was `rootdg` (the root disk group). For VxVM 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:

```
# vxpdg defaultdg
```

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*.

## 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 one of the following operations:

- 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/vxdmpinq`.

- 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 following options allow you to exclude devices from `vxddmp`:

- 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 VID: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 VID:PID.

Enter the VID:PID combination at the prompt.

```
Enter a VID:PID combination: [<pattern>,all,list,list-  
exclude,q,?]
```

If you selected any of the options, reboot the system for device exclusion to take effect.

## 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.

## Enabling Sun Java System Cluster support in VxVM (optional)

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.

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 Storage Foundation using the pkgadd command”](#) on page 88.
- 3 Initialize VxVM.  
See [“Configuring Veritas Volume Manager”](#) on page 115.
- 4 Configure shared disks.  
See [“Configuring shared disks on Solaris”](#) on page 122.

## Configuring shared disks on Solaris

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.

## 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 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 suggestions can protect 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.  
`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.

## Enabling the Intelligent Storage Provisioning (ISP) feature

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

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

## Enabling cluster support in VxVM (Optional)

This release includes an optional cluster feature that enables VxVM to be used in a 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.

The VxVM cluster feature requires a license, which can be obtained from your Customer Support channel.

### 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.
- 3 Initialize VxVM.
- 4 Start VEA.
- 5 Configure shared disks.

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

## 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.

### To convert existing disk groups to shared disk groups

- 1 Ensure that all systems that are running are part of the same cluster.
- 2 Start the cluster on at least one node.

For a two-node cluster, start the cluster on one node; for a four-node cluster, start the cluster on three nodes.

**3** Configure the disk groups using the following procedure.

To list all disk groups, use the following command:

```
# vxdg list
```

To deport disk groups to be shared, use the following command:

```
# vxdg deport disk_group_name
```

To import disk groups to be shared, use the following command on the master node:

```
# vxdg -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:

```
# vxdg 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 `vxdg list` command on each node to display the shared disk groups. This command displays the same list of shared disk groups displayed earlier.

## 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.

Also, hot-relocation can be configured.

See “[About hot-relocation](#)” on page 116.

## 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:

```
# vxdg list
```

This displays the existing shared disk groups.

## Upgrading in a clustered environment with FastResync set

Upgrading in a clustered environment with FastResync set requires additional steps.

This procedure applies to the following upgrade scenarios:

- Upgrading from VxVM 3.5 to VxVM 5.0MP3
- Upgrading from VxVM 3.5 Maintenance Pack 4 or from VxVM 3.2 Maintenance Pack 5 to VxVM 5.0MP3

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:

```
# vxdctl -c mode
```

- 2 On the master node, list which disk groups are shared by entering:

```
# vxdg -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 Patch 3 or from VxVM 3.2 Maintenance Patch 5, set FastResync to off for each volume:

```
# vxvol -g diskgroup set fastresync=off volume
```

## Configuring Veritas File System

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

```
/etc/vfstab
```

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

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

## Loading and unloading the file system module

On Solaris 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 vxfst_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
```

## 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
```

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 Storage Foundation Cluster File System

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

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

### To configure Storage Foundation Cluster File System

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

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

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

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

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

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

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

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

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

```
Checking system licensing
```

```
SFCFS license registered on system01
```

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

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

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

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

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

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

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

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

All systems are configured to create one cluster.

Enter the unique cluster name and Cluster ID number.

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

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

```
Discovering NICs on host1 ...discovered bge0 bge1 bge2 bge3
```

## 10 Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat link
on host1: [b,?] bge1
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,?] bge2

Would you like to configure a third private heartbeat
link? [y,n,q,b,?] (n) n
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) n
Are you using the same NICs for private heartbeat links
on all systems? [y,n,q,b,?] (y) y
```

---

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

---

Notice that in this example, *bge0* is not selected for use as a private heartbeat NIC because it already in use as the public network interface.

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

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

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

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

See the *Veritas Volume Manager Administrator's Guide*

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

- 13** You are now given the option of specifying the default name of a disk group that is to be assumed by Veritas Volume Manager commands if a disk group is not otherwise specified.

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

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

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

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

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

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

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

- 15** Validate the default disk group information, and press Return.
- 16** You may be prompted to verify the fully qualified hostname of the systems.

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

- 17** Enter **y** to accept the fully qualified domain name.

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

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

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

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

- 19** The Veritas Storage Foundation Cluster File System software is verified and configured.

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

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

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

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

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

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

- Install `VRTSgms` package from the 5.0 MP3 CD
- Install `VRTSodm` package from the 5.0 MP3 CD

- 22** Before installing Oracle binaries (`ORACLE_HOME`), consider these points:

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

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

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

## Configuring Veritas Volume Replicator

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

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

## To configure VVR

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

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

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

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

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

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

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

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

- 5 The script continues the initial system check. The script confirms success by displaying information, such as the OS version, communication with the remote hosts, and whether the required VVR packages are installed. Press Return to continue.
- 6 The script proceeds to verify whether the required licenses are installed. If a valid license for VVR is not present, the script prompts you to enter a license. The script validates whether the current license enables VVR.

See [“Symantec product licensing”](#) on page 22.

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

- 7** The script enables you to choose whether you want to use enclosure-based naming. The enclosure-based naming scheme (rather than disk device naming) is a feature of Veritas Volume Manager. You can reference disks using a symbolic name that is more meaningful than the operating system's normal device access name.

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

See the *Veritas Volume Manager Administrator's Guide*

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

- 8** Specify the default name of a disk group for Veritas Volume Manager commands, if a disk group is not otherwise specified.

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

```
vxdctl defaultdg diskgroup
```

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

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

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

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

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

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

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

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

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

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

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

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

Change the frequency of online statistics collection, if needed.

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

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

```
The maximum number of days for which VVR statistics  
can be retained is set to 3 on system01  
  
Do you want to change the maximum number of days  
for retaining VVR statistics on system01? [y,n,q] (n) n
```

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

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

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

- 15** Repeat steps 11 to 14 for all other systems.
- 16** Verify the fully qualified hostnames of the systems. Press Return to continue.

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

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

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

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

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

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

- 19** The configuration and startup completes automatically.

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

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

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

## Veritas Volume Replicator and Volume Manager setup after installation

VVR is fully integrated with Veritas Volume Manager (VxVM). Before using VVR, you must have the VxVM volumes set up and initialized.

Refer to the Volume Manager documentation for more information.

## Configuring and starting Veritas Enterprise Administrator

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

Optional configuration can also be completed at this time.

## Stopping and starting the VEA server

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

### To start up the VEA server

- 1 Check the state of the VEA server.

```
# /opt/VRTS/bin/vxsvcctl status
```

- 2 Stop the VEA server.

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

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

- 3 Start the VEA server.

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

The VEA server is automatically started on a reboot.

## Starting the VEA client on Windows or Solaris

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 > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator.

## Modifying optional connection access on Solaris

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.

### To modify connection access

- 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
```

## VMSA and VEA co-existence

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

---

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

---

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

```
Veritas VM Storage Administrator Server terminated.
```

```
Stopping Veritas VM Storage Administrator Server
```

```
### Terminated
```

## Configuring Veritas Enterprise Administrator for databases

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

### Configuring Veritas Enterprise Administrator for Oracle

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

#### Adding users to the VEA Service Console Registry for Oracle

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

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

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

```
# 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 [-A] [-f] -n user_name
```

where:

`-a user` adds a user to the registry

`-A` grants the user root access

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

`-n` indicates the name of the user.

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

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

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

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

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

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

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

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

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

where:

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

`-A` grants the user group root access

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

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

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

- 6 Restart the VEA Server.

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

## Removing users from the VEA Service Console Registry for Oracle

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

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

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

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

```
# 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
```

## Setting up Veritas Enterprise Administrator for DB2

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

### Adding users to Veritas Enterprise Administrator for DB2

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

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

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

```
# pkginfo -l VRTSd2gui | egrep STATUS
STATUS: completely installed
```

- 2 Stop the VEA server.

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

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

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

where:

-a user adds a user to the registry

-A grants the user root access

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

-n indicates the name of the user or group.

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

-d specifies the domain to which the user belongs.

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

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

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

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

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

where -a adds a user to the registry.

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

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

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

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

where:

-a `group` adds a group to the registry

-A grants the group root access

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

-n indicates the name of the user or group.

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

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

- 6 Restart the VEA Server.

```
# /opt/VRTS/bin/vxsvcctl restart
```

## Removing users from Veritas Enterprise Administrator for DB2

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

You cannot remove `root` from the AZDB.

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

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

```
# pkginfo -l VRTSd2gui | egrep STATUS  
STATUS: completely installed
```

- 2 Stop the VEA server.

```
# /opt/VRTS/bin/vxsvcctl 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/vxsvcctl restart
```

# Upgrading Storage Foundation

This chapter includes the following topics:

- [Upgrading Storage Foundation or the operating system](#)
- [Planning the upgrade](#)
- [Upgrading Storage Foundation from 5.0 to 5.0MP3 using the product installer or manual steps](#)
- [Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using the product installer or manual steps](#)
- [Upgrading Storage Foundation and/or Solaris using Live Upgrade](#)
- [Upgrading the Solaris OS only](#)
- [Upgrading Veritas Enterprise Administrator clients](#)
- [Upgrading Veritas Volume Replicator](#)
- [Upgrading language packages](#)
- [Post-upgrade tasks](#)
- [Verifying the Veritas Storage Foundation upgrade](#)

## Upgrading Storage Foundation or the operating system

Perform the procedures in the following sections to upgrade Storage Foundation or your operating system, or both. You can perform an upgrade to Storage

Foundation using the Veritas product installer or product installation script if you already have Storage Foundation installed.

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

---

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

---

## Planning the upgrade

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

Complete the following tasks in advance of upgrading:

- 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 meets all the requirements for software and hardware, including any required operating system patches.
- Schedule sufficient outage time for the upgrade, and downtime for any applications using the VxFS file systems or VxVM volumes.
- If using EMC PowerPath, ensure that you are using at least mandatory patch level 2.0.3.  
<http://entsupport.symantec.com/docs/234374>  
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 fallback point.)
- To upgrade on a remote host, `rsh` or `ssh` must be set up.  
See “[Configuring secure shell \(ssh\) or remote shell before installing products](#)” on page 24.
- Determine if the root disk is encapsulated.

See “[Determining if the root disk is encapsulated](#)” on page 149.

- Select the method to upgrade.  
See “[Upgrade paths for Storage Foundation 5.0MP3](#)” on page 149.

## Saving system information before upgrade

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

### To save system information before an upgrade

- 1 Log in as superuser.
- 2 Before upgrading, ensure that you have made backups of all data that you want to preserve.
- 3 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 4 If you are installing the HA version of the Veritas Storage Foundation 5.0MP3 software, follow the guidelines given in the *Veritas Cluster Server Installation Guide* and *Veritas Cluster Server Release Notes* for information on preserving your VCS configuration across the installation procedure.

## Determining if the root disk is encapsulated

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.

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

## Upgrade paths for Storage Foundation 5.0MP3

The procedure you use to upgrade Storage Foundation from a previous release depends on several factors, including the version of the existing Storage Foundation product, and the version of Solaris on which Storage Foundation is installed.

You may choose to upgrade your operating system concurrently with upgrading Storage Foundation. Storage Foundation 5.0MP3 is supported on Solaris versions 8, 9, and 10. You must upgrade the operating system if it is an earlier Solaris version which is not supported on 5.0MP3.

Before you upgrade, you also need to determine if the root disk is encapsulated. In some cases, an encapsulated root disk may affect which upgrade procedure to choose.

See [“Determining if the root disk is encapsulated”](#) on page 149.

---

**Note:** For Solaris 10, all non-global zones must be booted and in running state before using the common product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must upgrade each package in each non-global zone manually.

---

[Table 7-1](#) describes the upgrade procedures for Storage Foundation on Solaris systems.

**Table 7-1** Upgrade paths for Storage Foundation on Solaris

Solaris release	Storage Foundation version	Upgrade procedure
Solaris 2.6, 7, 8, 9	SF 3.5, 3.5 MP4	Upgrade to 5.0MP3 is not supported.  You must uninstall SF using the procedure in the Storage Foundation Installation Guide for your version, then install 5.0MP3.
Solaris 7	SF 4.0	Upgrade OS to at least Solaris 8. Upgrade to 5.0MP3 using the procedure:  See <a href="#">“Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using upgrade scripts (OS upgrade or encapsulated root disk)”</a> on page 171.

**Table 7-1** Upgrade paths for Storage Foundation on Solaris (*continued*)

Solaris release	Storage Foundation version	Upgrade procedure
Solaris 8, 9	SF 4.0	<p>If you do not plan to upgrade the OS, and if the root disk is not encapsulated, upgrade to 5.0MP3 using one of the following procedures:</p> <p>See <a href="#">“Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required”</a> on page 164.</p> <p>See <a href="#">“Upgrading Veritas Storage Foundation with manual steps when OS upgrade is not required”</a> on page 169.</p> <p>If you plan to upgrade the OS, or if the root disk is encapsulated, use the following procedure:</p> <p>See <a href="#">“Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using upgrade scripts (OS upgrade or encapsulated root disk)”</a> on page 171.</p>

**Table 7-1** Upgrade paths for Storage Foundation on Solaris (*continued*)

Solaris release	Storage Foundation version	Upgrade procedure
Solaris 8, 9, 10	SF 4.1	<p>If you do not plan to upgrade the OS, and if the root disk is not encapsulated, upgrade to 5.0MP3 using one of the following procedures:</p> <p>See <a href="#">“Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required”</a> on page 164.</p> <p>See <a href="#">“Upgrading Veritas Storage Foundation with manual steps when OS upgrade is not required”</a> on page 169.</p> <p>If you plan to upgrade the OS, or if the root disk is encapsulated, use the following procedure:</p> <p>See <a href="#">“Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using upgrade scripts (OS upgrade or encapsulated root disk)”</a> on page 171.</p>
Solaris 8, 9, 10	SF 5.0, including Maintenance Packs and Rolling Patches	<p>Upgrade to 5.0MP3 using the procedure:</p> <p>See <a href="#">“Upgrading Storage Foundation from 5.0 to 5.0MP3 using the product installer or manual steps”</a> on page 156.</p> <p>If you plan to upgrade the OS, upgrade to 5.0MP3 and then use the procedure to upgrade the OS only:</p> <p>See <a href="#">“Upgrading the Solaris OS only”</a> on page 184.</p>
Solaris 8, 9	SF 5.0MP3	<p>If you want to upgrade the OS, use the procedure:</p> <p>See <a href="#">“Upgrading the Solaris OS only”</a> on page 184.</p>

## Performing pre-installation checks and configuration

Use the following procedure to prepare for the upgrade.

### To prepare for the upgrade

- 1 Ensure that you have created a valid backup.  
See “[Saving system information before upgrade](#)” on page 149.
- 2 Review the *Veritas Storage Foundation Release Notes*.
- 3 Ensure that you have enough file system space to upgrade Veritas Storage Foundation. Also, identify where you will be copying the distribution and patch tar files. The usual place is `/patches/Veritas` when the root file system has enough space or `/var/tmp/patches` if the `/var` file system has enough space.

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

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

- 4 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 All file systems not on the `root` disk (therefore not required for booting) should 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 Storage Foundation assumes should be in `rootdg` but are not, must be unmounted and the associated entry in `/etc/vfstab` commented out.
- 8 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.

- 9 Make sure file systems are clean.

See [“Verifying that the file systems are clean”](#) on page 154.

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

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

See [“About upgrading disk layout versions”](#) on page 211.

- 11 Upgrade arrays (if required)

See [“Upgrading external ASLs and APMs”](#) on page 155.

## Verifying that the file systems are clean

Prior to upgrading to release 5.0MP3, verify that all file systems have been cleanly unmounted by running the `fsdb` command from the existing release of File System.

### To make sure the file systems are clean

- 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 external ASLs and APMs

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

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

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

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

## Upgrading Storage Foundation from 5.0 to 5.0MP3 using the product installer or manual steps

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

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

- Save system information before the upgrade.  
See [“Saving system information before upgrade”](#) on page 149.
- Perform preinstallation checks and configuration for the upgrade.  
See [“Performing pre-installation checks and configuration”](#) on page 153.
- Upgrade Storage Foundation to 5.0MP3 in one of the following ways:
  - Using the product installer. This is the recommended method of installation.  
See [“Upgrading Storage Foundation software from 5.0 to 5.0MP3 using the product installer”](#) on page 159.
  - Using manual steps (patchadd command).  
See [“Upgrading Storage Foundation from 5.0 to 5.0MP3 with manual steps \(patchadd\)”](#) on page 160.

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

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

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

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

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

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

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

- 2 Determine which external ASL packages are installed:

```
# grep /etc/vx/lib/discovery.d /var/sadm/install/contents | egrep -v VRTS
awk '{ print $10 }'
```

This command lists the packages which installed any files in the ASL directory. Any packages in this directory, other than the base VxVM package, are external ASL packages.

The `VRTSVxvm` package is the base VxVM package.

Example output:

```
# grep /etc/vx/lib/discovery.d /var/sadm/install/contents | egrep -v VRTS
awk '{ print $10 }'
VRTSIBM-DS4xxx-2.0-1.0
```

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

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

Determine if the MSA1500 ASL is installed:

```
# pkginfo | grep "VRTSHP-MSA1500"
```

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

```
# pkgchk -v pkg
```

where *pkg* is the ASL package name.

- 5 Determine if an ASL is 4.x:

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

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

For example:

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

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

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

```
# pkgrm pkg
```

- 7 Determine which external APM packages are installed:

```
# grep /etc/vx/apmkey.d /var/sadm/install/contents | egrep -v "VRTSvxvm" |  
'{ print $10 }'
```

This command displays package names which installed any APM keys.

Any package other than `VRTSvxvm` is an external APM package.

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

```
# pkgchk -v pkg
```

where *pkg* is the APM package.

Look for the following entry:

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

**9** Determine if an APM is 4.x:

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

where *apm* is the name of the APM.

For example:

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

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

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

```
# pkgrm pkg
```

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

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

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

See [“Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0”](#) on page 228.

### To upgrade Storage Foundation 5.0MP3

- 1** Load and mount the disc.  
See [“Mounting a software disc”](#) on page 46.
- 2** Move to the top-level directory on the DVD.

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

To upgrade the local system, enter the following command:

```
# ./installmp
```

---

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

---

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

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

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

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

- 4 After the initial system checks have completed successfully, press Enter to start the requirement checks for the upgrade.
- 5 After the requirement checks have completed successfully, press Enter to begin upgrading Storage Foundation.
- 6 Reboot each of the nodes on which you upgraded Storage Foundation.
- 7 If you set the value of the `vol_vvr_use_host_byte_order` tunable to 1, reboot the system.
- 8 If you want to use features of Veritas Storage Foundation 5.0 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.
- 9 If you are upgrading a cluster, restore any VCS configuration files as described in the *Veritas Cluster Server 5.0 Installation Guide* and *Veritas Cluster Server 5.0 Release Notes*.

## Upgrading Storage Foundation from 5.0 to 5.0MP3 with manual steps (patchadd)

To upgrade Veritas Storage Foundation from 5.0 to 5.0MP3 using manual steps, use the `patchadd` command to add the patches.

### To upgrade Storage Foundation using the patchadd command

- 1 Load and mount the software disc.  
See [“Mounting a software disc”](#) on page 46.
- 2 Copy the patches to a location to which you can write and then uncompress and untar the packages:

```
# cp -R /mount_point/storage_foundation/patches/* /tmp/patches
```

- 3 Uncompress and extract the packages by using the `gzcat` command:

```
# cd /tmp/patches
# gzcat *.gz | tar xvf -
```

- 4 Run the `dbed_patch_50ga` file if you are manually adding patches to any of the following packages: VRTSdbcom, VRTSdbed, VRTSdb2ed, VRTSorgui, VRTSd2gui, VRTSmapro, VRTSorweb, or VRTSd2web.

- 5 Run the following command to obtain a list of patches to install:

```
./installsf -requiredpkgs
```

See [“Determining the installation order”](#) on page 84.

- 6 Use the `patchadd` command to install the patches from step 5.

The `patchadd` utility also displays the name of the log file where the install status is saved.

This log file specifies the reason that the reboot is required, which could one of the following reasons:

- Some of the drivers failed to load.
- Some of the volumes were open.
- `/root`, `/usr` or `swap` is encapsulated.

If `patchadd` fails then check for the reason in the log file.

If `patchadd` completes and a successful installation message is displayed with a prior message requiring reboot, then reboot the system.

During the patch install (`patchadd`) if any messages appear that request a system reboot then the installation upgrade requires a system reboot.

- 7 Configure the SF installation using the `installsf -configure` command.

See [“Configuring the Storage Foundation products”](#) on page 95.

## Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using the product installer or manual steps

This section describes upgrading a Veritas Storage Foundation product from a release prior to 5.0 to 5.0MP3. 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. Choose the appropriate procedure for your situation.

Choose the appropriate procedure for your situation.

- If the current 4.x Storage Foundation product is installed on an operating system supported by 5.0MP3, you do not need to upgrade the operating system. If the root disk is not encapsulated, and you do not plan to upgrade the operating system, use one of the following upgrade procedures:
  - Upgrade SF but not OS with the product installer. This is the recommended upgrade procedure.  
See [“Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required”](#) on page 164.
  - Upgrade SF but not OS with manual steps (pkgadd and patchadd commands).  
See [“Upgrading Veritas Storage Foundation with manual steps when OS upgrade is not required”](#) on page 169.
- If the root disk is encapsulated, or if you plan to upgrade the operating system, you must perform additional steps to upgrade. If the current 4.x Storage Foundation product is installed on an operating system which is no longer supported by 5.0MP3, you must upgrade the operating system. If the root disk is encapsulated, or if you plan to upgrade the operating system, use the following upgrade procedure:  
See [“Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using upgrade scripts \(OS upgrade or encapsulated root disk\)”](#) on page 171.

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

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

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

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

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

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

- 2 Determine which external ASL packages are installed:

```
# grep /etc/vx/lib/discovery.d /var/sadm/install/contents | \
egrep -v VRTSvxvm | awk '{ print $10 }'
```

This command lists the packages which installed any files in the ASL directory. Any packages in this directory, other than the base VxVM package, are external ASL packages.

The `VRTSvxvm` package is the base VxVM package.

Example output:

```
# grep /etc/vx/lib/discovery.d /var/sadm/install/contents | \
egrep -v VRTSvxvm | awk '{ print $10 }'
VRTSIBM-DS4xxx-2.0-1.0
```

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

- 3 Remove any external ASLs.

```
# pkgrm pkg
```

- 4 Determine which external APM packages are installed:

```
# grep /etc/vx/apmkey.d /var/sadm/install/contents | \
egrep -v "VRTSvxvm" | awk '{ print $10 }'
```

This command lists the packages which installed any APM keys.

Any package other than `VRTSvxvm` is an external APM package.

- 5 Remove these APMs before upgrading Storage Foundation:

```
# pkgrm pkg
```

## Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required

This section describes upgrading to the current Veritas Storage Foundation if the root disk is unencapsulated, and you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 5.0MP3.

See [“Upgrade paths for Storage Foundation 5.0MP3”](#) on page 149.

This procedure can be used to upgrade Veritas Storage Foundation for Oracle, Veritas Storage Foundation for DB2 (on SPARC systems only), and Veritas Storage Foundation for Sybase.

See [“Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0”](#) on page 228.

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

### 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 Make sure the root disk is not encapsulated.

See [“Determining if the root disk is encapsulated”](#) on page 149.

- 3 Unmount any mounted VxFS file systems.

The installer supports the upgrade of multiple hosts, if each host is running the same version of VxVM and VxFS. Hosts must be upgraded separately if they are running different versions.

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.

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

List all service groups:

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

For each service group listed, take it offline:

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

- 5 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation.
- 6 Load and mount the disc.  
See [“Mounting a software disc”](#) on page 46.
- 7 To invoke the common installer, run the `installer` command on the disc as shown in this example:

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

- 8 Enter `I` to install/upgrade and press Return.

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

- 9 When the list of available products is displayed, select Veritas Storage Foundation, enter the corresponding number, and press Return.

Veritas Storage Foundation for Oracle, Veritas Storage Foundation for DB2 (on SPARC systems only), and Veritas Storage Foundation for Sybase can also be installed using this procedure. Select the number corresponding to one of those products, if desired.

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

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

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

**11** The installer checks the existing product license information.

You are prompted to enter additional license information, until all licenses for all systems have been entered. Then reply that you have no additional licenses to enter. License keys for additional product features should also be added at this time

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

**12** You can choose to install required packages or all packages. Optional packages include man pages, for example.

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

For example, you should see output similar to the following:

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

```
1) Install required Veritas Storage Foundation packages - 883 MB required
```

```
2) Install all Veritas Storage Foundation packages - 911 MB required
```

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

**13** You are prompted to confirm that you are ready to upgrade Storage Foundation.

```
Are you sure you want to upgrade Storage Foundation? [y,n,q] (y) y
```

If you select *y*, the installer stops Storage Foundation processes and make some configuration updates before upgrading.

**14** Verify the fully qualified hostname of the systems.

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

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

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

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

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

**16** Reboot the system.

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

**17** After rebooting, you must configure the system.

Start the installer and select 'C' for configure, or run the appropriate installation script with the `-configure` option.

**18** When the list of available products is displayed, select the number corresponding to the product installed in step 9, and press Return.

**19** The installer verifies the system licensing. You are prompted to enter additional license information, until all licenses for all systems have been entered. Then reply that you have no additional licenses to enter.

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

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

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

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

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

**21** The installer displays information about Veritas Volume Manager configuration settings. During a post-upgrade configuration, you cannot change these settings.

If you want to change these settings, you can reconfigure Veritas Volume Manager after the upgrade completes.

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

**22** Verify the fully qualified hostname of the systems.

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

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

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

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

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

**24** The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

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

**25** The configuration completes automatically, and the processes are started.

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

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

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

**26** For Veritas Storage Foundation for DB2 and Veritas Storage Foundation for Oracle, create the repository database.

See [“Creating and configuring the repository database for DB2 and Oracle”](#) on page 111.

Storage Foundation 5.0MP3 recognizes the existing disks and volumes without any intervention from you.

- 27** To take advantage of new features, upgrade the VxVM disk group version to the latest version.

See [“Upgrading VxVM disk group versions”](#) on page 213.

- 28** If you upgraded Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, revalidate the snapshots after upgrading to 5.0MP3.

For Storage Foundation for Oracle, validate the snapshots using the following command:

```
$ /opt/VRTS/bin/dbed_vmchecksnap -S ORACLE_SID  
\ -H ORACLE_HOME -f SNAPPLAN -o validate
```

For more information, refer to the *Veritas Storage Foundation for Oracle Administrator's Guide*

For Storage Foundation for DB2, validate the snapshots using the following command:

```
$ /opt/VRTS/bin/db2ed_vmchecksnap -D DB2DATABASE  
\ -f SNAPPLAN -o validate
```

For more information, refer to the *Veritas Storage Foundation for DB2 Administrator's Guide*

## Upgrading Veritas Storage Foundation with manual steps when OS upgrade is not required

This section describes upgrading from a previous version of Veritas Storage Foundation to the current Veritas Storage Foundation (5.0MP3) if the root disk is unencapsulated, and you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 5.0MP3.

See [“Upgrade paths for Storage Foundation 5.0MP3”](#) on page 149.

### 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 149.

- 4 Unmount any mounted VxFS file systems.

The installer supports the upgrade of multiple hosts, if each host is running the same version of VxVM and VxFS. Hosts must be upgraded separately if they are running different versions.

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/VRTSvcs/bin/hagrp -list
```

For each service group listed, take it offline:

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

- 6 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 5.0MP3.

- 7 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation.
- 8 Load and mount the disc.  
See [“Mounting a software disc”](#) on page 46.
- 9 Remove the Veritas packages from your existing 4.x installation.  
Refer to the *4.x Storage Foundation Installation Guide* for details.
- 10 Run the following command to obtain a list of packages and patches to install:  

```
./installsf -requiredpkgs
```

  
See [“Determining the installation order”](#) on page 84.
- 11 Use the `pkgadd` command to install the packages from step 10.  
Some packages require options with the `pkgadd` command.
- 12 Run the `dbed_patch_50ga` file if you are manually adding patches to any of the following packages: `VRTSdbcom`, `VRTSdbed`, `VRTSdb2ed`, `VRTSorgui`, `VRTSd2gui`, `VRTSmapro`, `VRTSorweb`, or `VRTSd2web`.
- 13 Use the `patchadd` command to install the patches from step 10.
- 14 Configure the SF installation using the `installsf -configure` command.  
See [“Configuring the Storage Foundation products”](#) on page 95.

## Upgrading Veritas Storage Foundation from 4.x to 5.0MP3 using upgrade scripts (OS upgrade or encapsulated root disk)

This section describes upgrading to the current Veritas Storage Foundation if you have an encapsulated root disk, need to upgrade the Solaris version, or both. If the operating system is not at a supported Solaris version, you must follow this procedure.

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).

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

#### To begin the upgrade

- 1 Load and mount the disc.

See “[Mounting a software disc](#)” on page 46.

- 2 Run the `upgrade_start` script to preserve the previous configuration of Volume Manager.

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

- 3 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, restore any other applications, and perform an `init 6` to completely restore the system.

- 4 If the root disk is encapsulated, reboot the machine.

- 5 Remove the existing Storage Foundation packages in one of the following ways:

- using the `uninstallsf` script
- using `pkgrm`

For details, refer to the *Storage Foundation Installation Guide* for the existing Storage Foundation version.

- 6 If you have VxFS file systems specified in the `/etc/vfstab` file, comment them out.

- 7 Reboot the machine.

- 8 If you are upgrading the operating system, do so now.

Refer to the Solaris installation documentation.

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

- 9 Install the Storage Foundation packages in one of the following ways:

- using the common installer

See [“To upgrade the Veritas Storage Foundation packages with the product installer”](#) on page 173.

- using manual steps

See [“To upgrade the Veritas Storage Foundation packages with manual steps”](#) on page 173.

### To upgrade the Veritas Storage Foundation packages with the product installer

- 1 Load and mount the disc.

See [“Mounting a software disc”](#) on page 46.

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

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

- 3 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.

- 4 If you commented out VxFS File System entries in the `/etc/vfstab` file in step 6 of the pre-upgrade procedures, uncomment them.

- 5 Complete the upgrade by restoring the configuration.

See [“Restoring the configuration and completing the upgrade”](#) on page 174.

### To upgrade the Veritas Storage Foundation packages with manual steps

- 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 Load and mount the disc.

See [“Mounting a software disc”](#) on page 46.

- 3 Add packages with the `pkgadd` command.
- 4 Run the `dbed_patch_50ga` file if you are manually adding patches to any of the following packages: `VRTSdbcom`, `VRTSdbed`, `VRTSdb2ed`, `VRTSorgui`, `VRTSd2gui`, `VRTSmapro`, `VRTSorweb`, or `VRTSd2web`.
- 5 Add patches with the `patchadd` command.
- 6 If you commented out VxFS File System entries in the `/etc/vfstab` file in step 6 of the pre-upgrade procedures, uncomment them.
- 7 Complete the upgrade by restoring the configuration.  
See [“Restoring the configuration and completing the upgrade”](#) on page 174.

#### Restoring the configuration and completing the upgrade

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

```
# /mount_point/storage_foundation/scripts/upgrade_finish
```
- 2 Reboot the machine (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.
- 3 Importing a pre-5.0MP3 Veritas Volume Manager disk group does not automatically upgrade the disk group version to the VxVM 5.0MP3 level. You may need to manually upgrade each of your disk groups following a VxVM upgrade.  
See [“Upgrading VxVM disk group versions”](#) on page 213.
- 4 To configure the software, run the installer script with the `-configure` option specified.

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

## Upgrading Storage Foundation 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.

Solaris Live Upgrade upgrades Storage Foundation software in a live environment.

The following 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 Storage Foundation/Solaris and easily switch between them with just one reboot

Storage Foundation Live Upgrade requires an alternate disk to upgrade. This disk can be a mirrored root disk or an independent unused disk which supports booting. Storage Foundation Live Upgrade is mainly done using the Solaris Live Upgrade commands.

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, or if upgrading to a managed host.

## Upgrading Storage Foundation or the Solaris OS or Both Using Live Upgrade

Use Live Upgrade to upgrade Storage Foundation, or the Solaris operating system, or both.

VxVM Live Upgrade supports the following upgrades:

- Upgrading Storage Foundation only (both encapsulated and unencapsulated root disk)
- Upgrading Solaris OS only (both encapsulated and unencapsulated root disk)
- Upgrading both Solaris OS and Storage Foundation (both encapsulated and unencapsulated root disk)
- Installing any Storage Foundation 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.

## Storage Foundation Live Upgrade commands and usage

Storage Foundation Live Upgrade uses the two commands `vxlustart` and `vxlufinish`. These commands are on the Storage Foundation 5.0MP3 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 option 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 `vxdg 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 Storage Foundation only. The command does not prompt for Solaris discs.
- u Option to specify the Solaris version to be upgraded. In case of upgrading only Storage Foundation, 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] [-g diskgroup] [-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.
- g Specify the disk group for the alternate root disk. If the current root disk is encapsulated, the alternate root disk is also encapsulated and placed into the specified disk group. The specified disk group must be unique.
- 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

Use the following procedure to begin the Live Upgrade.

### To begin the Live Upgrade

- 1 If you have not already obtained and installed a Storage Foundation 5.0MP3 license key, do so now.

See [“Supported Solaris operating systems”](#) on page 41.

- 2 Load and mount the software disc.

See [“Mounting a software disc”](#) on page 46.

- 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 /mount_point/storage_foundation/scripts/vxlustart .  
# cp /mount_point/storage_foundation/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:

See [“Upgrading Solaris from the software discs”](#) on page 179.

See [“Upgrading Solaris from network directory path”](#) on page 179.

## 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

Use the following procedure to complete the Solaris upgrade using Live 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 -R VRTSvxvm  
# pkgadd -R VRTSvxvm
```

If required, you can also apply these commands to the `VRTSvxf`s package.

This reinstallation is necessary to install Solaris 10 SMF support for Storage Foundation.

- 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 `halt` command and bring the machine to `ok` prompt to reboot the machine on the alternate root disk.

```
# ok boot alt_boot_disk
```

Where `alt_boot_disk` is the name of the alternate root disk.

If the alternate root disk has a problem rebooting or any upgrade issue, reboot the machine from the older disk.

## Upgrading Storage Foundation

Use the following procedure to upgrade Storage Foundation using Live Upgrade.

If you are using Live Update to install this 5.0MP3 release of Veritas Storage Foundation (including Veritas Storage Foundation for RAC and Storage Foundation Management System version 1.0 MP1), you must first execute the `dbed_patch_50ga`

file located in the `/mount_point/product-directory/scripts` directory before starting Live Update.

### To upgrade Storage Foundation using Live Upgrade

- 1 If only Storage Foundation needs to be upgraded, specify the `-U` option and the current Solaris version to the command. For example, if you want to upgrade Storage Foundation on Solaris 10, use the following command,

```
# vxlustart -u 5.10 -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 VRTSvmmman 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).

Your system may not have all listed packages installed.

Licensing in Storage Foundation 5.0MP3 requires the new `VRTSvlic` package. You do not need to remove the existing `VRTSlic` package.

- 5 If Storage Foundation 4.0 or 4.1 was previously installed, remove the old Storage Foundation 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 VRTSvmmman VRTSscpi
# pkgrm -R /altroot.5.X VRTStep VRTSap VRTSvrdoc \
  VRTSvrw
# pkgrm -R /altroot.5.X VRTSweb VRTSjre VRTSvcsvr \
  VRTSvrpro
# pkgrm -R /altroot.5.X VRTSfspro VRTSalloc
# pkgrm -R /altroot.5.X VRTSvmpo 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).

Your system may not have all listed packages installed.

Licensing in Storage Foundation 5.OMP3 requires the new `VRTSvlic` package.

You do not need to remove the existing `VRTSlic` package.

- 6 If Storage Foundation 5.0 was previously installed, remove the old Storage Foundation packages on the alternate root by entering:

```
# rm /altroot.5.10/var/vx/isis/vxisis.lock
# /opt/VRTSob/bin/vxsvc -k
# pkgrm -R /altroot.5.X VRTSvrdoc
# pkgrm -R /altroot.5.X VRTSvrw
# pkgrm -R /altroot.5.X VRTSweb
# pkgrm -R /altroot.5.X VRTSjrel5
# pkgrm -R /altroot.5.X VRTSjre
# pkgrm -R /altroot.5.X VRTSvcsvr
# pkgrm -R /altroot.5.X VRTSvrpro
# pkgrm -R /altroot.5.X VRTSfsmnd
# pkgrm -R /altroot.5.X VRTSfssdk
# pkgrm -R /altroot.5.X VRTSfsdoc
# pkgrm -R /altroot.5.X VRTSfsman
# pkgrm -R /altroot.5.X VRTSvxfs
# pkgrm -R /altroot.5.X VRTSddlpr
# pkgrm -R /altroot.5.X VRTSalloc
# pkgrm -R /altroot.5.X VRTSfspro
# pkgrm -R /altroot.5.X VRTSvmpro
# pkgrm -R /altroot.5.X VRTSvmdoc
# pkgrm -R /altroot.5.X VRTSvman
# pkgrm -R /altroot.5.X VRTSvxvm
# pkgrm -R /altroot.5.X VRTSobgui
# pkgrm -R /altroot.5.X VRTSob
# pkgrm -R /altroot.5.X VRTSperl
# pkgrm -R /altroot.5.X VRTSvlic
# pkgrm -R /altroot.5.X VRTScpi
```

where the *X* in `/altroot.5.X` corresponds to the Solaris version (8, 9 or 10).

Your system may not have all listed packages installed.

Licensing in Storage Foundation 5.0MP3 requires the new `VRTSvlic` package.

You do not need to remove the existing `VRTSlic` package.

- 7 Install the Storage Foundation 5.0MP3 packages from the distribution media by entering the command:

```
# cd /cdrom/cdrom0/storage_foundation
# ./installsf 'hostname' -rootpath /altroot.5.X
```

- 8 After installing the Storage Foundation software and patches, use the `vxlufinish` command to complete the Live Upgrade process.

```
# vxlufinish -u 5.8
```

The upgrade process is now complete.

- 9 Shut down the machine using `shutdown` or "`init 6`" to reboot the machine on the alternate root disk.

```
# init 6
```

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.

- 10 Configure the Storage Foundation 5.0MP3 packages from the distribution media by entering the following commands:

```
# cd /cdrom/cdrom0/storage_foundation
# ./installsf'hostame' -configure
```

Although you will see messages where VxVM cannot be reconfigured, other portions of Storage Foundation will be configured.

- 11 To verify the upgrade, check the Volume Manager version:

```
# pkginfo -l VRTSvxvm
```

- 12 Verify that all the file systems that were under VxVM control prior to the upgrade are now under VxVM control.

```
# df -k
```

## Upgrading the Solaris OS only

If you are running Storage Foundation 5.0MP3 with an earlier release of the Solaris operating system, you can upgrade the Solaris operating system using the following procedure.

---

**Warning:** You should only use this procedure to upgrade the Solaris operating system if you are running Storage Foundation 5.0MP3.

---

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
```

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 Storage Foundation.

See “[Mounting a software disc](#)” on page 46.

- 3 Change directory:

```
# cd /mount_point/volume_manager/scripts
```

- 4 Run the `upgrade_start` with the `-check` argument to detect any problems that exist which could prevent a successful upgrade. Use the `upgrade_start` script that was supplied with the currently installed SF 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. 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.

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.

- 8 After installing the Solaris operating system, install any Solaris patches required by Veritas Storage Foundation 5.0MP3.  
See the *Veritas Storage Foundation Release Notes*.
- 9 After the system is up with the upgraded Solaris operating system, bring the system down to single-user mode by entering:

```
# init s
```

- 10 Ensure that `/opt` is mounted.
- 11 Load and mount the software disc from the currently installed version of Storage Foundation.  
See [“Mounting a software disc”](#) on page 46.
- 12 If you upgraded to Solaris 10, you must reinstall certain Storage Foundation packages and patches in order to support Solaris 10 functionality.

For Storage Foundation, you must reinstall the `VRTSvxvm`, and `VRTSvxfs` packages.

For Storage Foundation HA, you must reinstall the `VRTSvxvm`, `VRTSvxfs` and `VRTSvcsag` packages.

To reinstall the required packages, follow the steps below:

- Remove the existing packages:

For Storage Foundation:

```
# pkgrm VRTSvxvm VRTSvxfs
```

For Storage Foundation HA:

```
# pkgrm VRTSvxvm VRTSvxfs VRTSgab VRTS11t VRTSvxfen
```

- Change to the directory containing the Storage Foundation packages.

```
# cd /mount_point/storage_foundation/pkg
```

- Uncompress the necessary files. For example, uncompress

```
VRTSvxvm.tar.gz, uncompress VRTSvxfs.tar.gz.
```

- Install the 5.0 packages with the `pkgadd` command.

For Storage Foundation:

```
# pkgadd VRTSvxvm VRTSvxfs
```

For Storage Foundation HA:

```
# pkgadd VRTSvxvm VRTSvxfs VRTS11t VRTSgab VRTSvxfen
```

- Run the following command to obtain a list of patches corresponding to the packages:

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

- Use `patchadd` to apply the Solaris 10 patches that correspond to the packages reinstalled in the previous steps.

For Storage Foundation on Solaris 10, apply the following patches:

```
# patchadd 123202-xx
```

```
# patchadd 122058-xx
```

For Storage Foundation HA on Solaris 10, apply the following patches:

```
# patchadd 123209-xx
```

```
# patchadd 123210-xx
```

```
# patchadd 123211-xx
```

```
# patchadd 123202-xx
```

```
# patchadd 122058-xx
```

- 13 Complete the upgrade from the software disc from the currently installed version of Storage Foundation by entering:

```
# ./upgrade_finish
```

- 14 Perform the reconfiguration reboot:

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

## Upgrading Veritas Enterprise Administrator clients

This section provides information about upgrading VEA clients.

### Upgrading the VEA client on a Microsoft Windows system

To upgrade the VEA client on a Windows system, first uninstall the existing VEA client package as follows:

- 1 Log in as administrator.
- 2 Select **Start > Settings > Control Panel**.
- 3 Double-click **Add/Remove Programs** to display a list of installed products.

- 4 Select **Veritas Enterprise Administrator** from the list, and click the **Remove** button.
- 5 Click **Yes** when a dialog box appears asking you to confirm the removal.
- 6 After removing the existing package, install the new VEA client package by following the instructions for a new installation.  
See [“Installing Veritas Enterprise Administrator”](#) on page 75.

## Upgrading the VEA client on a Solaris system

The current version of Veritas Enterprise Administrator requires only the `VRTSobgui` package. This package is installed when Veritas Storage Foundation products are installed. You only need to use this procedure if you plan to use Veritas Enterprise Administrator on a host other than the server that is used for Veritas Storage Foundation.

**To upgrade the VEA client on a Solaris host (other than the server) using `pkgadd`**

- 1 Log in as superuser (`root`).
- 2 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 to the current working directory on your system.

```
# cp /cdrom/cdrom0/storage_foundation/pkgs/VRTSobgui.tar.gz .
```

- 6 Decompress the packages, and then extract the contents.

```
# /cdrom/cdrom0/gnu/gunzip *.tar.gz  
# tar xvf VRTSobgui.tar
```

- 7 Use the `pkgadd` command to install the package. Answer any questions, as the installation proceeds.

```
# pkgadd -d . VRTSobgui
```

The VEA client package for Solaris is installed.

## Upgrading the VEA Windows client language package

To upgrade the Windows Client language package, you need to uninstall all the existing VEA 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.

See [“Installing the VEA client on Microsoft Windows”](#) on page 77.

## Upgrading Veritas Volume Replicator

This section describes upgrading Veritas Volume Replicator, unless Storage Foundation Cluster File System is present. In that case, refer to the procedures for upgrading Storage Foundation Cluster File System.

See [“Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0”](#) on page 228.

## Supported upgrade methods for Veritas Volume Replicator

To upgrade Veritas Volume Replicator (VVR) from 5.0 (or a 5.0 Maintenance Pack) to 5.0MP3, use the `installmp` procedure.

See [“Upgrading Storage Foundation from 5.0 to 5.0MP3 using the product installer or manual steps”](#) on page 156.

To upgrade Veritas Volume Replicator from a version earlier than 5.0 to VVR 5.0MP3, use one of the procedures in this section.

See [“Upgrading Veritas Volume Replicator using the Veritas product installer”](#) on page 190.

See [“Upgrading VVR without disrupting replication”](#) on page 198.

See [“Upgrading VVR when VCS agents are configured”](#) on page 199.

See [“Upgrading using VVR upgrade scripts”](#) on page 194.

## Upgrading Veritas Volume Replicator using the Veritas product installer

This section describes using the Veritas product installer. Use this method to upgrade VVR, unless you are upgrading in a scenario that is not supported by the Veritas product installer.

---

**Note:** We recommend using the Veritas product installer for upgrading VVR when possible.

---

See [“Supported upgrade methods for Veritas Volume Replicator”](#) on page 189.

To upgrade VVR only, use the Veritas product installer and select the Veritas Volume Replicator option. You can also use the `installvvr` script.

If you have multiple Veritas products, select the option for the appropriate Veritas product. Refer to the corresponding installation section for more details. For example, if you have Veritas Storage Foundation installed, select Veritas Storage Foundation in the Veritas product installer, or use the `installsf` script.

---

**Note:** If you have multiple Veritas products, we strongly recommend using the option to upgrade the entire product suite rather than upgrading each product individually. This ensures that upgrade steps are done in the proper order and product interdependencies are met.

---

For a complete list of Veritas Storage Foundation products, refer to the *Getting Started Guide*.

You may also be required to configure VVR after the upgrade.

See [“Configuring Veritas Volume Replicator”](#) on page 134.

To upgrade VVR, perform the following steps in the order presented:

- [“Preparing to upgrade VVR using the product installer”](#) on page 191.
- [“Upgrading Veritas Volume Replicator using the product installer”](#) on page 191.
- [“Restoring the original VVR configuration using the product installer”](#) on page 193.

- [“Upgrading language packages”](#) on page 194.

## Preparing to upgrade VVR using the product installer

This section describes upgrade preparation prior to using the product installer.

---

**Note:** Use a different set of instructions to upgrade an installation that uses VCS Agents for VVR.

---

See [“Upgrading VVR when VCS agents are configured”](#) on page 199.

### To prepare the upgrade through the product installer

- 1 Make sure that the disk groups that contain RVGs are at least at disk group version 110.

```
# vxdg list diskgroup
```

- 2 Make sure the size of the SRL volume is greater than 110 MB.

Refer to the *Veritas Volume Replicator Administrator's Guide*.

- 3 Stop all the applications involved in replication. For example, if a data volume contains a file system, unmount it.

- 4 Verify that all the Primary RLINKs are up-to-date on all the hosts.

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

---

**Note:** Do not continue until the Primary RLINKs are up-to-date.

---

## Upgrading Veritas Volume Replicator using the product installer

This section describes how to upgrade using the product installer.

### To upgrade the Veritas Volume Replicator with the product installer

- 1 Start the product installer:

```
# cd disc_path
```

```
# ./installer
```

where *disc\_path* is the location where the Veritas software disc is mounted.

- 2 Select Install/Upgrade a Product.

- 3 Select the appropriate product name:
  - If you are upgrading VVR only, use the Veritas Volume Replicator option.
  - If you are upgrading multiple Veritas products, select the appropriate option in the product installer to update all the Veritas products at the same time. Refer to the appropriate installation procedure for detailed instructions.

---

**Note:** If you have multiple Veritas products, we strongly recommend using the option to upgrade the entire product suite rather than upgrading each product individually. This ensures that upgrade steps are done in the proper order and interdependencies are met.

---

- 4 The script detects that an existing installation of VVR is present, and handles upgrade tasks.
- 5 Follow the prompts.

The script displays the name of the directory used for the upgrade. The upgrade directory is created in `/var/tmp` on the host from which the upgrade procedure was begun. The upgrade directory has the name `vvr_upgrade_hostname$timestamp` where the `hostname` is the machine being upgraded, and `$timestamp` is the same digit sequence as the suffix of the log file created for the current session.

For example, the directory `/var/tmp/vvr_upgrade_system01126061743` contains the upgrade files for the host `system01`.

---

**Note:** We strongly recommend you back up the upgrade directory created here, because it is used to restore the configuration.

---

- 6 The script displays the location of the log files.

When the script completes, it displays messages similar to the following:

```
CPI WARNING V-9-11-2246 You have completed upgrading VxVM
on some or all of the systems. Reboot your systems at this time.
```

- 7 Prior to rebooting, copy the VVRTypes.cf from `/etc/VRTSvcs/conf` to:

```
/etc/VRTSvcs/conf/config.
```

- 8 When the upgrade completes, the hosts that are being upgraded must be rebooted.

See [“Restoring the original VVR configuration using the product installer”](#) on page 193.

---

**Note:** If you are upgrading an installation that uses VCS Agents for VVR, do not configure VVR until after you reboot the machine, unfreeze the service groups and restore the original configuration.

See [“Unfreezing the service groups”](#) on page 204.

See [“Restoring the original configuration when VCS agents are configured”](#) on page 205.

---

## Restoring the original VVR configuration using the product installer

You must configure VVR to restore the original configuration and complete the upgrade. Configuring VVR also starts the VVR processes.

If VCS Agents are configured, refer to the procedure in that section.

See [“Restoring the original configuration when VCS agents are configured”](#) on page 205.

### To restore the original configuration through the product installer

- 1 On all Secondary hosts, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup shrinkto volume_name volume_length
```

where `volume_length` is the length of the volume on the Primary.

- 2 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

- 3 Restore the original configuration on each host that has been upgraded, starting with the Secondary hosts. To restore the configuration, configure VVR using one of the following methods:
  - Use the Veritas product installer, select Configure an Installed Product, and then select Veritas Volume Replicator.
  - Use the installation script `installvvr` with the `-configure` option.

The configuration is restored from the configuration files and scripts that were saved in the upgrade directory during the upgrade session.

See “[Configuring Veritas Volume Replicator](#)” on page 134.

- 4 Restart the applications that were stopped.

### If the upgrade fails

If the upgrade fails during the configuration phase, after displaying the VVR upgrade directory, the configuration needs to be restored before the next attempt. Run the scripts in the upgrade directory in the following order to restore the configuration:

```
# restoresrl
# adddcm
# srlprot
# attrlink
# start.rvg
```

After the configuration is restored, the current step can be retried.

### Upgrading language packages

If you are upgrading a language version of VVR, you can upgrade the language packages required by VVR after you have upgraded the base VVR packages. Verify that the English installation is correct before proceeding.

## Upgrading using VVR upgrade scripts

This section describes the procedure to upgrade to VVR 5.0MP3 using upgrade scripts. Use this method only if you are upgrading in a scenario that is not supported by the Veritas product installer. This procedure only applies in a private disk group environment. This procedure does not apply to VVR upgrade in a shared disk group environment.

See “[About upgrading Storage Foundation Cluster File System](#)” on page 227.

---

**Note:** We recommend using the Veritas product installer for upgrading VVR when possible.

---

See “Supported upgrade methods for Veritas Volume Replicator” on page 189.

The upgrade procedure retains the existing VVR configuration. After upgrading, you can use the existing VVR configuration, without running the `vxinstall` command.

To upgrade VVR, perform the following tasks in the order presented:

- [Preparing to upgrade using upgrade scripts](#)
- [Upgrading Veritas packages using upgrade scripts](#)
- [Restoring the original configuration using upgrade scripts](#)

## Preparing to upgrade using upgrade scripts

This section describes how to prepare your configuration for the upgrade using the upgrade scripts:

To prepare your upgrade through the upgrade scripts

- 1 Make sure that the disk groups that contain RVGs are at least at disk group version 110.

```
# vxdg list diskgroup
```

- 2 Make sure the size of the SRL volume is greater than 110 MB.

Refer to the *Veritas Volume Replicator Administrator's Guide*.

- 3 Stop all the applications involved in replication. For example, if a data volume contains a file system, unmount it.

- 4 Verify that all the Primary RLINKs are up-to-date on all the hosts.

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

---

**Caution:** Do not continue until the Primary RLINKs are up-to-date.

---

- 5 Run the `vvr_upgrade_start` script on all hosts to save the original VVR configuration using the following command:

```
# /disc_path/volume_replicator/scripts/vvr_upgrade_start
```

---

**Note:** If the `vvr_upgrade_start` script finds that the SRL size is less than 110 MB, then the script fails and reverts back to the original configuration. It stops with a message that prompts you to modify the SRL size.

---

For information on resizing the SRL, refer to the *Veritas Volume Replicator Administrator's Guide*.

## Upgrading Veritas packages using upgrade scripts

This section describes how to upgrade the packages through the upgrade script.

### On all hosts on which the upgrade is to be performed

- 1 If necessary, upgrade the operating system.
- 2 Upgrade Storage Foundation from the product disc to version 5.0MP3.
- 3 If you have not rebooted the system, reboot it now using the following command:

```
# /usr/sbin/shutdown -y -i6 -g0
```

During the reboot process, ignore the following error messages that appear on the Primary console:

```
VxVM VVR vxrlink ERROR V-5-1-3371 Can not recover rlink_name.  
rvg_name is in PASSTHRU mode
```

```
VxVM VVR vxrlink ERROR V-5-1-3473 Log header I/O error
```

Also ignore the following error message that appears on the Secondary console:

```
WARNING: VxVM VVR vxio V-5-0-278 Rlink rlink_name is stale and  
not replicating
```

- 4 Upgrade the required and optional packages for VVR. Perform the following tasks in the ordered indicated:
  - Remove the old packages.  
See [“Removing the VVR packages”](#) on page 290.
  - Copy the packages from the Veritas software disk to a temporary directory.

```
# cd /disc_path
```

```
# cp -r volume_replicator/pkgs/* /tmp_dir
```

- Unzip the package files.

```
# gunzip VRTS*.gz
```

- Decompress and extract each package.

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

- Use the following command to display the list of VVR packages. The packages must be installed in the order shown.

```
# ./installvvr -installpkgs
```

## 5 Install the new packages using the `pkgadd` command.

Some configurations may require upgrading an installation with VCS Agents for VVR configured.

See “[Unfreezing the service groups](#)” on page 204.

---

**Note:** If you have additional Veritas products to upgrade, refer to the product installation guide for a list of packages to upgrade.

---

## Restoring the original configuration using upgrade scripts

This section describes how to complete the upgrade and restore the original configuration using the upgrade scripts.

### To restore the original configuration through the upgrade script

- 1 On all Secondary hosts, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup shrinkto volume_name volume_length
```

where *volume\_length* is the length of the volume on the Primary

- 2 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

- 3 Issue the following command on all the hosts to complete the upgrade. If a host contains only Secondary RVGs, we recommend that you first run the following command on that host:

```
# /disc_path/volume_replicator/scripts/vvr_upgrade_finish
```

The `vvr_upgrade_finish` script upgrades only the SRL, after which, the RVG cannot work with the earlier versions of VxVM or VVR.

- 4 Restart the applications that were stopped.

## Upgrading VVR without disrupting replication

This section describes the upgrade procedure from an earlier version of VVR to the current version of VVR when replication is in progress, assuming that you do not need to upgrade all the hosts in the RDS simultaneously.

You may also need to set up replication between versions.

See [“Planning an upgrade from the previous VVR version”](#) on page 35.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade can be performed either on the Primary or on the Secondary. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS. This section includes separate sets of steps, for the Primary upgrade and for the Secondary upgrade.

---

**Note:** If you have a cluster setup, you must upgrade all the nodes in the cluster at the same time.

---

See [“Upgrading VVR when VCS agents are configured”](#) on page 199.

### Upgrading VVR on the Secondary

Follow these instructions to upgrade the Secondary hosts.

### To upgrade the Secondary

- 1 Stop replication to the Secondary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvname
```

- 2 Upgrade from VVR 4.1 MP1 to VVR 5.0MP3 on the Secondary.  
See [“Upgrading Veritas Volume Replicator using the Veritas product installer”](#) on page 190.
- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvname sec_hostname
```

### Upgrading VVR on the Primary

After you upgrade the Secondary, use the Veritas product installer to upgrade the Primary.

See [“Upgrading Veritas Volume Replicator using the Veritas product installer”](#) on page 190.

---

**Note:** Reduce application downtime while upgrading by planning your upgrade.

---

See [“Planning an upgrade from the previous VVR version”](#) on page 35.

## Upgrading VVR when VCS agents are configured

This section details the procedure for upgrading VVR when VCS agents for VVR are configured:

- Use the Veritas product installer for upgrading VVR, unless you are upgrading in a scenario that is not supported by the Veritas product installer.

---

**Note:** We recommend using the Veritas product installer for upgrading VVR when possible.

---

See [“Supported upgrade methods for Veritas Volume Replicator”](#) on page 189.

- Use the VVR upgrade scripts only if you are upgrading in a scenario that is not supported by the Veritas product installer.

## Prerequisites

The following lists the VVR upgrade prerequisites with VCS agents:

- Make sure the size of the SRL volume is greater than 110 MB.  
Refer to the *Veritas Volume Replicator Administrator's Guide*.

To upgrade VVR when VCS agents for VVR are configured, perform the following tasks in the order presented:

- [Freezing the service groups and stopping all the applications](#)
- [Preparing for the upgrade when VCS agents are configured](#)
- [Upgrading Veritas Volume Replicator packages when VCS is present](#)
- [Unfreezing the service groups](#)
- [Restoring the original configuration when VCS agents are configured](#)

## Freezing the service groups and stopping all the applications

This section describes how to freeze the service groups and stop all applications.

To freeze the service groups and stop applications

**Perform the following steps for the Primary and Secondary clusters:**

- 1 Log in as the superuser.
- 2 Make sure that `/opt/VRTS/bin` is in your PATH so that you can execute all the product commands.
- 3 Because the upgrade requires a reboot, cleanly shut down all applications as follows:
  - OFFLINE all application service groups that do not contain RVG or RVGShared resources, except the ClusterService, cvm and RVGLogowner groups.
  - If the application resources are part of the same service group as an RVG or RVGShared resource, then OFFLINE only the application resources.

---

**Note:** You must also stop any remaining applications not managed by VCS.

---

- 4 On any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 5 On any node in the cluster, list the groups in your configuration:

```
# hagrps -list
```

- 6 On any node in the cluster, freeze all service groups except the ClusterService group by typing the following command for each group name displayed in the output from step 5.

```
# hagrps -freeze group_name -persistent
```

---

**Note:** Write down the list of frozen service groups for future use.

---

- 7 On any node in the cluster, save the configuration file (`main.cf`) with the groups frozen:

```
# haconf -dump -makero
```

---

**Note:** Continue only after you have performed steps 3 to step 7 for each cluster.

---

- 8 Display the list of service groups that have RVG resources and the nodes on which each service group is online by typing the following command on any node in the cluster:

```
# hares -display -type RVG -attribute State
Resource      Attribute      System      Value
VVRGrp        State          system02    ONLINE
ORAGrp        State          system02    ONLINE
```

---

**Note:** For the resources that are ONLINE, write down the nodes displayed in the System column of the output.

---

- 9 Repeat step 8 for each cluster.
- 10 For private disk groups, determine and note down the hosts on which the disk groups are imported.

See [“Determining the nodes on which disk groups are online”](#) on page 202.

- 11 For shared disk groups, run the following command on any node in the CVM cluster:

```
# vxdctl -c mode
```

Note the master and record it for future use.

- 12 If you are upgrading Storage Foundation Cluster File System, continue with the upgrade steps in that section.

See “[Performing the upgrade \(phased or full\)](#)” on page 234.

### Determining the nodes on which disk groups are online

For private disk groups, determine and note down the hosts on which the disk groups containing RVG resources are imported. This information is required for restoring the configuration after the upgrade.

#### To determine the online disk groups

- 1 On any node in the cluster, list the disk groups in your configuration, and note down the disk group names listed in the output for future use:

```
# hares -display -type RVG -attribute DiskGroup
```

---

**Note:** Write down the list of the disk groups that are under VCS control.

---

- 2 For each disk group listed in the output in step 1, list its corresponding disk group resource name:

```
# hares -list DiskGroup=diskgroup Type=DiskGroup
```

- 3 For each disk group resource name listed in the output in step 2, get and note down the node on which the disk group is imported by typing the following command:

```
# hares -display dg_resname -attribute State
```

The output displays the disk groups that are under VCS control and nodes on which the disk groups are imported.

### Preparing for the upgrade when VCS agents are configured

If you have configured the VCS agents, it is recommended that you take backups of the configuration files, such as `main.cf` and `types.cf`, which are present in the `/etc/VRTSvcs/conf/config` directory.

### To prepare a configuration with VCS agents for an upgrade

- 1 List the disk groups on each of the nodes by typing the following command on each node:

```
# vxdisk -o alldgs list
```

The output displays a list of the disk groups that are under VCS control and the disk groups that are not under VCS control.

---

**Note:** The disk groups that are not locally imported are displayed in parentheses.

---

- 2 If any of the disk groups have not been imported on any node, import them. For disk groups in your VCS configuration, you can import them on any node. For disk groups that are not under VCS control, choose an appropriate node on which to import the disk group. Enter the following command on the appropriate node:

```
# vxdg -t import diskgroup
```

- 3 If a disk group is already imported, then recover the disk group by typing the following command on the node on which it is imported:

```
# vxrecover -bs
```

- 4 Verify that all the Primary RLINKs are up to date.

```
# vxrlink -g diskgroup status rlink_name
```

---

**Note:** Do not continue until the Primary RLINKs are up-to-date.

---

## Upgrading Veritas Volume Replicator packages when VCS is present

After you have performed the steps in the preceding sections, upgrade your Veritas products in one of the following ways:

- [Using the Veritas product installer](#)
- [Using the VVR upgrade scripts](#)

## Using the Veritas product installer

Upgrade your Veritas products by selecting the option for the product suite.

See [“Upgrading Veritas Volume Replicator using the product installer”](#) on page 191.

---

**Note:** You must upgrade VVR on all nodes for the Primary and Secondary cluster, after preparing for your VCS agent configuration.

---

See [“Preparing for the upgrade when VCS agents are configured”](#) on page 202.

## Using the VVR upgrade scripts

Use this option to upgrade only if you are upgrading in a scenario that is not supported by the Veritas product installer.

---

**Note:** We recommend using the Veritas product installer for upgrading VVR when possible.

---

See [“Supported upgrade methods for Veritas Volume Replicator”](#) on page 189.

---

**Note:** To preserve your configuration, you must run the `vvr_upgrade_start` script before upgrading your installation.

---

## To upgrade VVR and preserve the original configuration

- 1 To preserve the original VVR configuration, run the `vvr_upgrade_start` script on the nodes that are to be upgraded, before upgrading your installation.

```
# /disc_path/storage_foundation/volume_replicator/scripts \  
/vvr_upgrade_start
```

- 2 Upgrade your Veritas products.

See [“Using the Veritas product installer”](#) on page 204.

If you do not want to preserve the original VVR configuration, see [Preparing to upgrade VVR using the product installer](#).

## Unfreezing the service groups

This section describes how to unfreeze services groups and bring them online.

### To unfreeze the service groups

- 1 On any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 2 Unfreeze all service groups that were frozen in step 6 of the section [Preparing for the upgrade when VCS agents are configured](#) by typing the following command on any node in the cluster:

```
# hagrps -unfreeze service_group -persistent
```

- 3 Save the configuration on any node in the cluster.

```
# haconf -dump -makero
```

- 4 If you are upgrading in a shared disk group environment, bring online the RVGShared groups with the following commands:

```
# hagrps -online RVGShared -sys masterhost
```

- 5 Bring the respective IP resources online on each node.

See [“Preparing for the upgrade when VCS agents are configured”](#) on page 202.

Type the following command on any node in the cluster.

```
# hares -online ip_name -sys system
```

This IP is the virtual IP that is used for replication within the cluster.

- 6 In shared disk group environment, online the virtual IP resource on the master node that you noted in step 11.

## Restoring the original configuration when VCS agents are configured

This section describes how to restore a configuration with VCS configured agents.

---

**Note:** Restore the original configuration only after you have upgraded VVR on all nodes for the Primary and Secondary cluster.

---

**To restore the original configuration**

- 1 Import all the disk groups in your VVR configuration.

```
# vxdg -t import diskgroup
```

Each disk group should be imported onto the same node on which it was online when the upgrade was performed. The reboot after the upgrade could result in another node being online; for example, because of the order of the nodes in the AutoStartList. In this case, switch the VCS group containing the disk groups to the node on which the disk group was online while preparing for the upgrade.

```
# hagrps -switch grpname -to system
```

- 2 Recover all the disk groups by typing the following command on the node on which the disk group was imported in step 1.

```
# vxrecover -bs
```

- 3 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

- 4 On all nodes that are Secondary hosts of VVR, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup shrinkto volume_name volume_length
```

where *volume\_length* is the length of the volume on the Primary.

---

**Note:** Do not continue until you complete this step on all the nodes in the Primary and Secondary clusters on which VVR is upgraded.

---

- 5 Restore the configuration according to the method you used for upgrade:

If you upgraded with the VVR upgrade scripts

Complete the upgrade by running the `vvr_upgrade_finish` script on all the nodes on which VVR was upgraded. We recommend that you first run the `vvr_upgrade_finish` script on each node that is a Secondary host of VVR.

Perform the following tasks in the order indicated:

- To run the `vvr_upgrade_finish` script, type the following command:

```
# /disc_path/volume_replicator/scripts/vvr_upgrade_finish
```

where *disc\_path* is the location where the Veritas software disc is mounted.

- Attach the RLINKs on the nodes on which the messages were displayed:

```
# vxrlink -g diskgroup -f att rlink_name
```

If you upgraded with the product installer

Use the Veritas product installer and select Configure an Installed Product. Or use the installation script with the `-configure` option.

- 6 Bring online the RVGLogowner group on the master:

```
# hagr -online RVGLogownerGrp -sys masterhost
```

- 7 Start and bring online the cvm group on the remaining host:

```
# hagr -online cvm -sys slave_host
```

- 8 Restart the applications that were stopped.

## 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 71.

## Post-upgrade tasks

The tasks in the following sections must be performed after upgrade, to restore the previous configurations and set up SF 5.0MP3 correctly. Perform the tasks required for the products and features that are relevant to your installation.

### Optional configuration steps

After the upgrade is complete, additional tasks may need performed.

You can perform the following optional configuration steps:

- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.
- If you want to use features of Veritas Storage Foundation 5.0MP3 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.  
See "Upgrading VxFS disk layout versions" on page 211.  
See "Upgrading VxVM disk group versions" on page 213.
- If you upgraded to the current version of Veritas Storage Foundation for DB2, refer to the DB2 configuration procedure.  
See "Configuring Storage Foundation for Databases" on page 106.

## Upgrading to the new repository database for DB2 and Oracle

If you are installing or upgrading Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, you need to either create a new repository database or migrate your old repository database to the new one. To use the `db2ed_update` or `dbed_update` command, you must be the instance owner or database administrator.

### To upgrade your repository database

- 1 Create and configure the new repository database with the `sfua_db_config` command.  

```
# /opt/VRTSdbcom/bin/sfua_db_config
```
- 2 Migrate your old repository information into the new repository database.

- 3 If you are upgrading Veritas Storage Foundation for DB2 in a single-host environment, run the `db2ed_update` command.

```
# /opt/VRTS/bin/db2ed_update -D DB2DATABASE
```

If you are upgrading Veritas Storage Foundation for DB2 in a high availability (HA) environment, run the `db2ed_update` command with the `-G` option.

```
# /opt/VRTS/bin/db2ed_update -D DB2DATABASE -G service_group
```

If you are upgrading Veritas Storage Foundation for Oracle in a single-host environment, run the `dbed_update` command.

```
# /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

If you are upgrading Veritas Storage Foundation for Oracle in a high availability (HA) environment, run the `dbed_update` command with the `-G` option.

```
# /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME \  
-G service_group
```

- 4 After the upgrade, the old repository database will be marked with a hidden file name, such as `/etc/vx/vxdba/.instance_name`, to prevent further updates. If you need to perform an additional upgrade, the file must be removed.

## Changing permissions for Storage Foundation for Databases

After installing the Veritas Storage Foundation 5.0MP3 patches, follow these post-installation steps to ensure Veritas Storage Foundation for Oracle and Veritas Storage Foundation for DB2 commands work correctly.

This does not apply to Sybase.

---

**Note:** Do not recursively change permissions, groups, or owners

---

### To change permissions

- 1 Change permissions for the following directory, depending on which product you have installed:

For Veritas Storage Foundation for Oracle:

```
# chmod 550 /opt/VRTSdbed
```

For Veritas Storage Foundation for DB2:

```
# chmod 550 /opt/VRTSdb2ed
```

- 2 Reset owner and group settings to the appropriate owner and group for the database administrators on your system.

For example, in Veritas Storage Foundation for Oracle, to change owner to the user oracle and the group dba, run the following command:

```
# chown oracle:dba /opt/VRTSdbed
```

In Veritas Storage Foundation for DB2, for example, to change owner to the user db2 and the group db2grp, run the following command:

```
# chown db2:db2grp /opt/VRTSdb2ed
```

- 3 Upgrade the repository.
- 4 In a standalone instance, run `sfua_db_config` once:

```
# /opt/VRTSdbcom/bin/sfua_db_config
```

This step completes the upgrade of the repository in a standalone configuration.

- 5 In a cluster environment, complete the remaining steps.
- 6 Unconfigure the SFUA repository from the VCS configuration.

```
# /opt/VRTSdbcom/bin/sfua_db_config -o unconfig_cluster
```

- 7 Mount the repository file system manually.
- 8 Run the repository upgrade command again with no options.

```
# /opt/VRTSdbcom/bin/sfua_db_config
```

## About upgrading disk layout versions

You must upgrade your older disk layout versions to make use of the extended features available in the Veritas File System 5.0MP3 release.

See the *Veritas Storage Foundation Release Notes 5.0MP3* for information on new features.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 7 as described in the following sections.

### Upgrading VxFS disk layout versions

File System disk layouts may need to be upgraded.

VxFS 5.0MP3 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 5.0MP3. In the next major release of VxFS, support for disk layouts Version 4 and 5 may be dropped.

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 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 5.0MP3. You can upgrade these layout versions online before installing VxFS 5.0MP3, or upgrade them using `vxfsconvert` after installing VxFS 5.0MP3.

To upgrade VxFS Release 3.5 MP2 or lower with disk layout versions 1, 2, 3, or 5, use `vxupgrade` to upgrade to disk layout version 4 or version 5.

To upgrade VxFS Release 4.0 or higher with disk layout versions 1 or 2, use `vxfsconvert` to upgrade to disk layout version 4.

To upgrade VxFS Release 4.0 or higher with disk layout versions 4 or 5, use `vxupgrade` 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.

---

**Warning:** The contents of intent logs created on previous disk layout versions cannot be used after the disk layout version is upgraded.

---

## Requirements for upgrading to disk layout Version 7

Converting a previous disk layout to a Version 7 disk layout requires adequate free space. The space and time required to complete the upgrade increases with the number of files, extended attributes, and hard links in the file system. Typical maximum space is at least two additional inodes with one block for every inode. Allow at least ten minutes to upgrade for every million inodes in the file system.

## Migrating from `/etc/vx/vxdba` to `/var/vx/vxdba` for DB2 and Oracle

If you are upgrading Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, you can migrate to `/var/vx/vxdba` to save space under the root partition. Migrating to `/var/vx/vxdba` is optional. However, if you do not perform this migration, you cannot remove any file or directory from `/etc/vx/vxdba` to ensure proper operation. This procedure can be done at any time.

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 -rpf /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 VxVM disk group versions

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk groups.

After upgrading from Storage Foundation 4.x to 5.0MP3, you must upgrade any existing disk groups which are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

For 5.0MP3, the Veritas Volume Manager disk group version is the same as it was for the VxVM 5.0 release. Upgrading the disk group version is only required if you upgraded from a version earlier than 5.0.

Use the following command to find the version of a disk group:

```
# vxdg list diskgroup
```

To upgrade a disk group to the current disk group version, use the following command:

```
# vxdg upgrade diskgroup
```

For more information about disk group versions, see the *Veritas Volume Manager Administrator's Guide*.

## Updating variables

In /etc/profile, update the PATH and MANPATH variables as needed.

MANPATH could include `/opt/VRTS/man` and PATH `/opt/VRTS/bin`.

## Setting the default disk group

In releases prior to Volume Manager 4.0, the default disk group was `rootdg` (the root disk group). For Volume Manager to function, the `rootdg` disk group had to exist and it had to contain at least one disk.

This requirement no longer exists; however, you may find it convenient to create a system-wide default disk group. The main benefit of creating a default disk group is that VxVM commands default to the default disk group. You do not need to use the `-g` option.

You can set the name of the default disk group after installation by running the following command on a system:

```
# vxctl defaultdg diskgroup
```

See the *Veritas Volume Manager Administrator's Guide*.

If you want to confirm that the root disk is encapsulated, enter the command:

```
# vxvg bootdg
```

## Upgrading the Array Support Library

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software packages.

You can obtain ASL packages from the following locations:

- The VxVM release package
- The disk array provided by the vendor
- The Symantec Technical Support site

### 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 5.0MP3. 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 5.0MP3. 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 5.0MP3, 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**

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 5.0MP3, 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 5.0MP3, 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 5.0MP3, you must remove the suppression of the subpaths and controllers so that DMP can determine the association between EMCpower metadevices and `c#t#d#` disk devices.

In the following scenarios, you may need to unsuppress DMP subpaths and controllers:

- Converting a foreign disk  
See [“Converting a foreign disk to auto:simple”](#) on page 216.
- Converting a defined disk  
See [“Converting a defined disk to auto:simple”](#) on page 219.
- Converting a powervxvm disk  
See [“Converting a powervxvm disk to auto:simple”](#) on page 222.

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 `vxddladm 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 5.0MP3.

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 5.0MP3 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 metadvice, `emcpower10s2`, enter the following command:

```
# vxdmadm 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 5.0MP3.

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 40Sep 24 17:59 /dev/vx/rmp/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 5.0MP3 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:

```
# vxddm adm getsubpaths dmpnodename=emcpower10s2
```

- 7 Import the disk group and start the volumes:

```
# vxdg 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
```

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 5.0MP3.

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  TUTILO  PUTILO
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 5.0MP3 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
emcpower0s2	auto:simple	-	-	online

**7 Import the disk group and start the volumes.**

```
# vxvg import ppdg
# vxvol -g ppdg startall
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c6t0d12s2	auto:sliced	-	-	online
emcpower0s2	auto:simple	ppdisk01	ppdg	online

## Converting from QuickLog to Multi-Volume support

The 4.1 release of the Veritas File System is the last major release to support QuickLog. The Version 6 or Version 7 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 or Version 7 disk layout is used.

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.

```
# qloddetach -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 or Version 7 disk layout.

See [“About upgrading disk layout versions”](#) on page 211.

For example:

```
# 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
```

## Verifying the Veritas Storage Foundation upgrade

Refer to the section about verifying the installation to verify the upgrade.

See [“Verifying that the products were installed”](#) on page 257.

# Upgrading Storage Foundation Cluster File System

This chapter includes the following topics:

- [About upgrading Storage Foundation Cluster File System](#)
- [Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0](#)
- [Upgrading Storage Foundation Cluster File System to 5.0MP3 on a Storage Foundation 5.0 system](#)

## About upgrading Storage Foundation Cluster File System

Perform the procedures in the following sections to upgrade Storage Foundation Cluster File System. You can perform an upgrade to Storage Foundation using the Veritas product installer or product installation script if you already have Storage Foundation Cluster File System installed.

See [“Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0”](#) on page 228.

See [“Upgrading Storage Foundation Cluster File System to 5.0MP3 on a Storage Foundation 5.0 system”](#) on page 245.

# Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0

This section contains procedures for the Veritas Storage Foundation Cluster File System upgrade.

## Planning the upgrade

Complete the following tasks in advance of upgrading:

- Review the *Veritas Storage Foundation Release Notes* for any late-breaking information on upgrading your system.
- Be sure that the administrator doing the upgrade has root access and a working knowledge of system administration.
- Schedule sufficient outage time for the upgrade.
- Make sure you have upgraded all the file systems to disk layout Version 6, before you upgrade SFCFS to 5.0MP3. Disk layout Version 7 is the recommended version for SFCFS 5.0MP3.  
See `vxupgrade(1M)`, `vxconvert(1M)`, and `fsadm(1M)` manual pages.  
See the *Veritas File System Administrator's Guide*.
- Verify all the file systems are working fine and data is intact.  
See the `cfsmount(1M)` manual page.

## Preparing the system and backing up files before upgrading

Before upgrading an installed Veritas Storage Foundation Cluster File System, preserve the existing configuration information.

To preserve the existing configuration information, perform the following actions:

- Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/vfstab` file. You will need to recreate these entries in the `/etc/vfstab` file on the freshly installed system.
- Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/etc/vfstab`. You should also run the `vxlicrep`, `vxdisklist`, and `vxprint-ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.
- Use the `vxlicrep` command to make a record of the currently installed Veritas licenses.
- Back up the configuration files.

## Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0

```
# cp /etc/VRTSvcs/conf/ backupdirectory/
```

- If Veritas Volume Replicator (VVR) is configured, do the following steps in the order shown:
  - Make sure that the disk groups that contain RVGs are at least at disk group version 110.

```
# vxdg list diskgroup
```

- Make sure the size of the SRL volume is greater than 110 MB. Refer to the *Veritas Volume Replicator Administrator's Guide*.
- Stop all the applications involved in replication. For example, if a data volume contains a file system, unmount it.
- Verify that all of the Primary RLINKs are up to date:

```
# vxrlink -g diskgroup status rlink_name
```

---

**Note:** Do not continue until the Primary RLINKs are up to date.

---

## Upgrade paths for Veritas Storage Foundation Cluster File System 5.0MP3

[Table 8-1](#) shows the upgrade paths for Veritas Storage Foundation Cluster File System.

**Table 8-1** Upgrade paths

From	Upgrade to	Tasks
Storage Foundation Cluster File System 4.0	Storage Foundation Cluster File System 5.0MP3	See <a href="#">“Overview of procedures”</a> on page 230.
Storage Foundation Cluster File System 4.0 MP1	Storage Foundation Cluster File System 5.0MP3	Not Supported. You must upgrade SFCFS 4.0 MP1 to 4.0 MP2 to upgrade to 5.0MP3.

**Table 8-1** Upgrade paths (*continued*)

From	Upgrade to	Tasks
Storage Foundation Cluster File System 4.0 MP2	Storage Foundation Cluster File System 5.0MP3	See <a href="#">“Overview of procedures”</a> on page 230.
Storage Foundation Cluster File System 4.1	Storage Foundation Cluster File System 5.0MP3	See <a href="#">“Overview of procedures”</a> on page 230.
Any Storage Foundation Cluster File System 4.1 MP release	Storage Foundation Cluster File System 5.0MP3	See <a href="#">“Overview of procedures”</a> on page 230.

## Overview of procedures

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation Cluster File System: phased and full.

---

**Note:** If VVR is configured, phased upgrade is not supported. We recommend that the secondary cluster be upgraded before the primary cluster in the RDS.

---

The upgrade procedures apply to both the phased and full upgrade procedures unless otherwise noted. Occasionally, steps differ between the two procedures. Screen output is also common between both procedures unless otherwise noted.

---

**Note:** Both procedures automatically uninstall the previous version of the software.

---

### Phased upgrade

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time. Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

The stages of the phased upgrade procedure are:

- Modify the configuration information in the `main.cf` file.
- Select one or more nodes to upgrade.

- Install the new version.
- Shut down VCS on remaining non-upgraded nodes and ensure the file systems are clean.
- Reboot the upgraded nodes.
- Install the new version on each remaining node and reboot them.

---

**Note:** A phased upgrade should not be performed from one of the nodes in the cluster.

---

## Full upgrade

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

The stages of the full upgrade procedure are:

- Ensure the file systems are clean.
- Modify the configuration information in the `main.cf` file.
- Install the new version.
- Reboot the upgraded nodes.

## Ensuring the file systems are clean (full only)

Before upgrading to SFCFS 5.0MP3, ensure that the file systems are clean. To ensure that the logs have been replayed and the file systems are marked clean:

### To ensure the file systems are clean

- 1 Log in as superuser onto any node in the cluster.
- 2 Offline the group on each node of the cluster:

```
# hagrps -offline group -sys system01
# hagrps -offline group -sys system02
# hagrps -offline group -sys system03
# hagrps -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

Repeat this step for each SFCFS service group.

- 3 Check and repair each file system:

```
# fsck -F vxfs /dev/vx/dsk/diskgroup/volume
```

Repeat this step for each file system.

## Modifying the main.cf file (phased or full)

Save a copy of the `main.cf` file and modify the configuration information in the `main.cf` file.

### To modify the main.cf file

- 1 On any node, make a copy of the current `main.cf` file. For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
  /etc/VRTSvcs/conf/main.save
```

- 2 Choose one node from the cluster to execute step 3 through step 9.
- 3 On the node you selected in step 2, enter:

```
# haconf -makerw
# hares -unlink vxfsckd qlogckd
# hares -unlink qlogckd cvm_clus
# hares -link vxfsckd cvm_clus
# hares -delete qlogckd
# haconf -dump -makero
```

4 On all the nodes in the cluster, enter:

```
# ps -ef | grep qlogckd
# kill -9 pid_of_qlogckd
# modinfo | grep -i qlog
# modunload -i module_id_of_qlog
```

5 On the node you selected in step 2, stop VCS on all nodes:

```
# /opt/VRTS/bin/hastop -all -force
```

6 On the node you selected in step 2 and if you have configured the VCS Cluster Manager (Web Console), complete the following to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

■ Remove VRTSweb:

```
Process VRTSweb (
    PathName = "/opt/VRTSvcs/bin/haweb"
    Arguments = "10.129.96.64 8181"
)
```

■ Replace it with:

```
VRTSWebApp VCSweb (
    Critical = 0
    AppName = vcs
    InstallDir = "/opt/VRTSweb/VERITAS"
    TimeForOnline = 5
)
```

■ Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named `csgnic` and the public NIC device is `hme0`, add:

```
NIC csgnic (
    Device = hme0
```

■ Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMngr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
ntfr requires csgnic
webip requires csgnic
```

- 7 On the node you selected in step 2, remove `qlogckd` from the `/etc/VRTSvcs/conf/config/main.cf` file. For example:

```
CFSQlogckd qlogckd (  
    Critical = 0  
)
```

Make sure you remove all dependencies on `qlogckd` from the `main.cf` file.

- 8 On the node you selected in step 2, verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file:

```
# cd /etc/VRTSvcs/conf/config  
# /opt/VRTS/bin/hacf -verify .
```

- 9 On the node you selected in step 2, start VCS:

```
# /opt/VRTS/bin/hastart
```

- 10 On the remaining nodes in the cluster, start VCS:

```
# /opt/VRTS/bin/hastart
```

- 11 If VVR is configured, freeze the service groups and stop the applications.

See [“Freezing the service groups and stopping all the applications”](#) on page 200.

## Performing the upgrade (phased or full)

This section describes how to upgrade for phased or full.

If you are performing a phased upgrade, select one or more nodes to upgrade.

### To perform the upgrade

- 1 Log in as superuser.
- 2 Insert the appropriate media disc into your system’s DVD-ROM drive.
- 3 If volume management software is running on your system, the software disc automatically mounts as `/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

```
# mount -F hfs -o ro /dev/dsk/c0t6d0s2 /cdrom
```

where `c#t#d#` is the location of the CD drive.

- 4 Change to the top-level directory on the disc:

```
# cd /cdrom
```

- 5 Verify there are no VxFS file systems mounted on the nodes being upgraded:

```
# mount -p | grep vxfs
```

If any VxFS file systems are mounted, offline the group on each node of the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

Repeat this step for each SFCFS service group.

- 6 If performing a phased upgrade, start the installation procedure from the node selected in the cluster. In the following example the phased upgrade is performed on one node in a four-node cluster.

Ensure that the HAD daemon of VCS is running on all nodes selected for the upgrade. Enter the following command, and then press **Return**.

```
# ./installsfcfs system01 system02
```

If performing a full upgrade, start the installation from any node in the cluster. Enter the following command, and then press **y** to upgrade the cluster configuration.

```
# ./installsfcfs
```

- 7 Press **Return** to begin installing infrastructure packages.
- 8 Press **Return** to begin license verification.
- 9 Press **Return** to begin the uninstall. The uninstall utility checks the system's uninstallation requirements and identifies packages, patches and dependencies.

The output can be extensive and exceed 100 lines, depending on your configuration.

- 10 If VCS is running you are prompted to upgrade and reset the VCS password. To ensure security, passwords are not written to the install log.

```
installsfcfs must now make configuration updates and stop  
the cluster before upgrading VCS packages.
```

```
Are you ready to begin the Cluster Server upgrade at this  
time? [y,n,q]
```

- 11 Enter **y**.
- 12 At the prompt, enter your new password.
- 13 Reenter your new password.
- 14 Output shows information that Cluster Server must be stopped on a running system. Enter **y** to continue.
- 15 Press **Return** to begin removing the previous packages and installing the new.
- 16 Press **Return** again for summary information about logs and reboots.  
Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.
- 17 If performing a phased upgrade, proceed to shut down VCS.  
See [“Shutting down VCS \(phased only\)”](#) on page 236.  
If performing a full upgrade, proceed to updating the configuration.  
See [“Updating the configuration and confirm startup \(phased or full\)”](#) on page 238.

## Shutting down VCS (phased only)

Shutdown VCS on remaining nodes that are not being upgraded to preventing them from rejoining the cluster.

### To shut down the cluster

- 1 Separate the nodes that are not being upgraded from those that are.
- 2 Check to see if there are frozen CVM and SFCFS groups, enter

```
# /opt/VRTSvcs/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

## Upgrading Storage Foundation Cluster File System to 5.0MP3 from a release prior to 5.0

- Make configuration read/write:

```
# /opt/VRTSvcs/bin/haconf -makerw
```

- Unfreeze the group:

```
# /opt/VRTSvcs/bin/hagrp -unfreeze group -persistent
```

- Save the configuration:

```
# /opt/VRTSvcs/bin/haconf -dump -makero
```

- 3 Offline the group on each of the remaining nodes of the cluster:

```
# hagrp -offline group -sys system01
# hagrp -offline group -sys system02
# hagrp -offline group -sys system03
# hagrp -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

Repeat this step for each SFCFS service group.

- 4 Check and repair each file system:

```
# fsck -F vxfs /dev/vx/dsk/diskgroup/volume
```

Repeat this step for each file system.

- 5 Complete the following steps on each node.

On each node that is not being upgraded, shutdown VCS, enter the following commands:

- Enter the following commands:

```
# /opt/VRTSvcs/bin/hastop -local
# /etc/init.d/vxfen stop
# /opt/VRTS/bin/fsclustadm cfsdeinit
```

- Check to see if the qllog module is loaded:

```
# /usr/sbin/modinfo | grep qllog
```

- If the qllog module is loaded, then unload it:

```
# /usr/sbin/modunload -i <module_id>
```

- Enter the following commands to stop GAB and LLT:

```
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- 6 Proceed to updating the configuration.

See [“Updating the configuration and confirm startup \(phased or full\)”](#) on page 238.

## Updating the configuration and confirm startup (phased or full)

Perform the following steps on each upgraded node.

### To update the configuration and confirm startup

- 1 Remove the `/etc/VRTSvcs/conf/config/.stale` file, if it exists.

```
# rm -f /etc/VRTSvcs/conf/config/.stale
```

- 2 Reboot the upgraded nodes.

```
# reboot
```

- 3 After the nodes reboot, verify that LLT is running:

```
# lltconfig
LLT is running
```

- 4 Verify GAB is configured:

```
# gabconfig -l | grep 'Driver.state' | \
grep Configured
Driver state : Configured
```

- 5 Verify VxVM daemon is started and enabled:

```
# /opt/VRTS/bin/vxdctl mode
mode: enabled
```

- 6 Confirm all upgraded nodes are in a running state.

```
# /opt/VRTSvcs/bin/hasys -state | grep RUNNING | \
/usr/bin/wc -l
```

```
1
```

- 7 Log in as superuser.
- 8 Insert the appropriate media disc into your system's CD-ROM drive.
- 9 If volume management software is running on your system, the software disc automatically mounts as `/cdrom`.

If volume management software is not available to mount the CD-ROM, you must mount it manually, enter:

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
```

where `c##d##` is the location of the CD drive.

- 10 Change to the top-level directory on the disc:

```
# cd /cdrom
```

- 11 Run `installsfdfs` from any node in the cluster:

```
# ./installsfdfs -configure system01 system02
```

- 12 After the configuration is complete, the CVM and SFCFS groups may come up frozen. To find out the frozen CVM and SFCFS groups, enter the following command:

```
# /opt/VRTS/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

- Make the configuration read/write.

```
# /opt/VRTS/bin/haconf -makerw
```

- Unfreeze the group.

```
# /opt/VRTS/bin/hagrp -unfreeze group -persistent
```

- Save the configuration.

```
# /opt/VRTS/bin/haconf -dump -makero
```

- 13** If VVR is configured, and the CVM and SFCFS groups are offline, bring the groups online in the following order:

Bring online the CVM groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group -sys system01
# /opt/VRTS/bin/hagrp -online group -sys system02
```

where *group* is the VCS service group that has the CVMVoldg resource.

Bring online the RVGShared groups and the virtual IP on the master node using the following commands:

```
# hagrp -online RVGShared -sys masterhost
# hares -online ip_name masterhost
```

Bring online the SFCFS groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group -sys system01
# /opt/VRTS/bin/hagrp -online group -sys system02
```

where *group* is the VCS service group that has the CFSMount resource.

If the SFCFS service groups do not come online then your file system could be dirty.

See [“Making the file systems clean”](#) on page 243.

- 14** If performing a phased upgrade, upgrade the remaining nodes.

See [“Upgrading remaining nodes \(phased only\)”](#) on page 240.

If performing a full upgrade, proceed to upgrade the remaining nodes.

See [“Upgrading remaining nodes \(full only\)”](#) on page 242.

## Upgrading remaining nodes (phased only)

This section describes how to upgrade the remaining nodes.

### To upgrade remaining nodes

- 1 Log in as superuser.
- 2 Insert the appropriate media disc into your system’s CD-ROM drive.
- 3 Change to the `storage_foundation_cluster_file_system` directory:

```
# cd /cdrom/storage_foundation_cluster_file_system
```

- 4 Install SFCFS on the remaining nodes. Type:

```
# ./installsfcfs system03 system04
```

When upgrading the remaining nodes, you may be prompted that Cluster Server should have been running during the initial upgrade.

- 5 Press **y** to continue and follow all the prompts.
- 6 Reboot the recently upgraded nodes.

```
# reboot
```

Following the reboot, VCS configuration files are automatically updated to match the files of the first group of upgraded nodes.

- 7 Log in as superuser.
- 8 Insert the appropriate media disc into your system's CD-ROM drive.
- 9 If volume management software is running on your system, the software disc automatically mounts as `/cdrom`.

If volume management software is not available to mount the CD-ROM, you must mount it manually, enter:

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
```

where `c##t##d##` is the location of the CD drive.

- 10 Change to the top-level directory on the disc:

```
# cd /cdrom
```

- 11 Run `installsfcfs` from any node in the cluster:

```
# ./installsfcfs -configure system03 system04
```

- 12 Check cluster status. Type:

```
# hastatus -summary
```

- 13 If you are configuring SFCFS for a fenced environment.

See the *Veritas Cluster Server Administrator's Guide*.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide*.

- 14 To verify the cluster protocol version, enter the following command that displays the same on a running node:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

If the cluster protocol version is less than 70, then it needs to be upgraded to 70 for SFCFS 5.0MP3.

The cluster protocol version can only be updated on the master node.

Upgrade the entire cluster using the following command on the master node.

```
# /opt/VRTS/bin/vxdctl upgrade
```

- 15 Type the following command on one upgraded node to enable membership:

```
# gabconfig -xc
```

## Upgrading remaining nodes (full only)

This section describes how to upgrade the remaining nodes.

### To upgrade remaining nodes

- 1 If you are configuring SFCFS for a fenced environment.

See the *Veritas Cluster Server Administrator's Guide*.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide*.

- 2 To verify the cluster protocol version, enter the following command that displays the same on a running node:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

If the cluster protocol version is less than 70, then it needs to be upgraded to 70 for SFCFS 5.0MP3.

The cluster protocol version can only be updated on the master node.

Upgrade the entire cluster using the following command on the master node.

```
# /opt/VRTS/bin/vxdctl upgrade
```

- 3 Type the following command on one upgraded node to enable membership:

```
# gabconfig -xc
```

## Making the file systems clean

If you upgrade to SFCFS 5.0MP3 and the file systems are dirty, you have to deport the shared disk group and import it as non-shared. After the import, run `fsck`. `fsck` should succeed. Then deport the disk group and import it back as shared.

### To make the file systems clean

- 1 Log in as superuser onto the CVM master node.
- 2 If performing a full upgrade, offline the group on all the nodes in the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

If performing a phased upgrade, offline the group:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

- 3 Deport the disk group:

```
# vxdg deport diskgroup
```

where *diskgroup* is the SFCFS disk group.

- 4 Import the disk group:

```
# vxdg -C import diskgroup
```

- 5 Start all the volumes in the disk group:

```
# vxvol -g diskgroup startall
```

- 6 Check and repair each file system:

```
# fsck -F vxfs /dev/vx/dsk/diskgroup/volume
```

Repeat this step for each file system.

- 7 Deport the disk group:

```
# vxdg deport diskgroup
```

**8** Import the disk group:

```
# vxdg -s import diskgroup
```

**9** Start all the volumes in the disk group:

```
# vxvol -g diskgroup startall
```

**10** If VVR is configured, upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

**11** If performing a full upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01  
# hares -clear resource -sys system02  
# hares -clear resource -sys system03  
# hares -clear resource -sys system04
```

If performing a phased upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01  
# hares -clear resource -sys system02
```

**12** If performing a full upgrade, online the group on all the nodes in the cluster:

```
# hagrps -online group -sys system01  
# hagrps -online group -sys system02  
# hagrps -online group -sys system03  
# hagrps -online group -sys system04
```

If performing a phased upgrade, online the group:

```
# hagrps -online group -sys system01  
# hagrps -online group -sys system02
```

**13** If VVR is configured, bring online the RVGLogowner group on the master:

```
# hagrps -online RVGLogownerGrp -sys masterhost
```

Restart the applications that were stopped.

## Upgrading language packages

Language packages are upgraded using the `install_lp` script. If you are upgrading a language version you must do so after upgrading all product packages. Verify that the original installation is correct before proceeding.

### To upgrade language packages

- ◆ Insert the language disc and launch the `install_lp` script. If you are using volume management software the disc is automatically mounted as `/cdrom`.

```
# cd cdrom
# ./install_lp
```

# Upgrading Storage Foundation Cluster File System to 5.0MP3 on a Storage Foundation 5.0 system

Use this procedure to upgrade to 5.0MP3 from 5.0, or from an earlier 5.0 Maintenance Pack.

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation Cluster File System: phased and full.

See [“Phased upgrade for a Maintenance Pack”](#) on page 246.

See [“Full upgrade for a Maintenance Pack”](#) on page 252.

An upgrade requires stopping cluster failover functionality during the entire procedure. The upgrade is performed in a number of stages depending on the type of upgrade you are performing.

You must have superuser (root) privileges to install the Veritas software.

You should also review the *Veritas Storage Foundation Release Notes* for important release information.

---

**Caution:** A phased upgrade procedure results in a system PANIC on configurations where LLT is configured over UDP. This issue is fixed in 5.0 MP1. This issue is specific to configurations where LLT is configured over UDP and not present in usual LLT Ethernet configurations. The full upgrade procedure should be used for upgrading from SFCFS 5.0 or SFCFS 5.0 RP1 to SFCFS 5.0 MP1 on configurations where LLT is configured over UDP.

---

## Preparing to upgrade to the Maintenance Pack

If you are upgrading an installed Veritas Storage Foundation 5.0 version or from an earlier 5.0 Maintenance Pack, preserve the existing configuration information.

To preserve the existing configuration information, perform the following actions:

- Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/vfstab` file. You will need to recreate these entries in the `/etc/vfstab` file on the freshly installed system.
- Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/etc/vfstab`. You should also run the `vxlicrep`, `vxdisklist`, and `vxprint-ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.
- Use the `vxlicrep` command to make a record of the currently installed Veritas licenses.

## Phased upgrade for a Maintenance Pack

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time.

Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

Each phase of the phased upgrade should be performed on more than one node of the cluster.

The stages of the phased upgrade procedure include the following steps:

- Freeze service group operations and stop cluster failover operations.
- Select a two or more nodes to upgrade, and leave a group of one or more nodes running.
- Take the selected group of nodes offline and prepare them for the upgrade.
- Upgrade the Veritas Storage Foundation Cluster File System software on the selected group of nodes.
- Take the second group of nodes offline.
- Bring the first group of nodes online.
- Upgrade the second group of nodes.
- Bring the second group of nodes online and restart cluster failover services. The cluster is fully restored.

## Performing the phased upgrade for a Maintenance Pack

This section describes how to perform a phased upgrade for a Maintenance Pack.

### To freeze service group operations and stop cluster failover

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so that you can execute all product commands.
- 3 From any node in the cluster, make the cluster configuration writable.

```
# haconf -makerw
```

- 4 Enter the following command to freeze high availability service group operations on each node:

```
# hasys -freeze -persistent node_name
```

- 5 Make the configuration read-only

```
# haconf -dump -makero
```

### To select the nodes for the upgrade

- 1 Select one or more nodes to upgrade first.
- 2 Leave a group of one or more nodes running.

## Upgrading the Veritas Storage Foundation Cluster File System software to a Maintenance Pack

After the selected group of nodes is offline, the Veritas Storage Foundation Cluster File System software can be upgraded, using `installmp` for the Maintenance Pack.

### To take the selected group of nodes offline and prepare them for the upgrade

- 1 Stop cluster operations on each node in the group being upgraded, by entering the following command:

```
# hastop -local
```

- 2 Check if each node's root disk is under VxVM control by running this command:

```
# df -v /
```

- 3 The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

- 4 On each node, use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df | grep vxfs
```

- 5 On each node in the cluster, unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

- 6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

To avoid data corruption, do not proceed until all RLINKs are up-to-date.

- 7 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

- 8 If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 9 If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.
- 10 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
- 11 On each node, stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

#### To upgrade the Veritas Storage Foundation Cluster File System software

- 1 Insert the appropriate media disc into your system's DVD-ROM drive.
- 2 If volume management software is running on your system, the software disc automatically mounts as `/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
```

where `c#t#d#` is the location of the CD drive.

- 3 Change to the top-level directory on the disc:

```
# cd /cdrom
```

- 4 To upgrade the Storage Foundation Cluster File System, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:

- To install on the local system only, enter the following command:

```
# ./installmp
```

- To install on more than one system using secure shell (ssh) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ...
```

- To install on more than one system using remote shell (rsh) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ... -rsh
```

- 5 After the initial system checks are complete, press **Return** to start the requirement checks.
- 6 After the requirement checks are complete, press **Return** to start upgrading the packages. If you are upgrading multiple nodes, you have the option of upgrading them simultaneously. You will be prompted after the upgrade is complete.
- 7 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
- 8 Reboot the system (or systems).

## Upgrading the remaining nodes

This section describes how to upgrade the remaining nodes.

Take the second group of nodes offline.

Bring the first group (with the newly installed patches) online to restart cluster failover services.

Upgrade the second group of nodes.

### To take the second group of nodes offline

- ◆ Stop cluster operations on each node in the second group being upgraded, by entering the following command:

```
# hastop -local
```

### To bring the first group of nodes online

- ◆ Bring the first group of nodes online.

See [“Bringing the upgraded group of nodes online”](#) on page 251.

**To upgrade the second group of nodes**

- 1 To upgrade the second group of nodes, perform the upgrade of the Veritas Storage Foundation Cluster File System software on the second group of nodes.

See [“Upgrading the Veritas Storage Foundation Cluster File System software to a Maintenance Pack”](#) on page 247.

- 2 Then bring the second group of nodes online.

See [“Bringing the upgraded group of nodes online”](#) on page 251.

**Bringing the upgraded group of nodes online**

Use the following procedure to bring the upgraded group of nodes online.

**To bring the upgraded group of nodes online**

- 1 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.
- 2 If necessary, reinstate any missing mount points in the `/etc/vfstab` file on each node.
- 3 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.
- 4 Make the VCS configuration writable again from any node in the upgraded group:

```
# haconf -makerw
```

- 5 Enter the following command on each node in the upgraded group to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent node_name
```

- 6 Make the configuration read-only:

```
# haconf -dump -makero
```

- 7 Bring the CVM service group online on each node in the upgraded group:

```
# hagrps -online cvm -sys node_name
```

- 8 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 9 If you have stopped any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

- 10 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem  
# mount /checkpoint_name
```

- 11 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

- 12 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

## Full upgrade for a Maintenance Pack

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

The stages of the full upgrade procedure are:

- Freeze service group operations and stop cluster failover operations.
- Take all nodes in the cluster offline and install the software patches.
- Bring all the nodes (with the newly installed patches) online to restart cluster failover services. The cluster is fully restored.

### Performing the full upgrade to a Maintenance Pack

This section describes how to perform a full upgrade to a Maintenance Pack.

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

**To prepare for a full upgrade to a Maintenance Pack**

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
- 3 Stop high-availability cluster operations. This command can be executed from any node in the cluster, and stops cluster operations on all the nodes.

```
# hstop -all
```

- 4 Check if each node's root disk is under VxVM control by running this command:

```
# df -v /
```

- 5 The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

- 6 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df | grep vxfs
```

**7** Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

**8** If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvlg stop` command to stop each RVG individually:

```
# vxrvlg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

To avoid data corruption, do not proceed until all RLINKs are up-to-date.

**9** Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

**10** If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```

**11** If there are still disk groups that are imported at this time then proceed with the remaining steps. Otherwise, skip to the procedure to upgrade the Veritas software.

**12** Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

**13** Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

- 14 To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

- 15 Continue to the procedure to upgrade the Veritas Storage Foundation Cluster File System software.

### To upgrade the Veritas Storage Foundation Cluster File System software

- 1 Insert the appropriate media disc into your system's DVD-ROM drive.
- 2 If volume management software is running on your system, the software disc automatically mounts as `/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

```
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
```

where `c#t#d#` is the location of the CD drive.

- 3 Change to the top-level directory on the disc:

```
# cd /cdrom
```

- 4 To upgrade the Storage Foundation Cluster File System, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:

- To install on the local system only, enter the following command:

```
# ./installmp
```

- To install on more than one system using secure shell (SSH) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ...
```

- To install on more than one system using remote shell (RSH) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ... -rsh
```

- 5 After the initial system checks are complete, press **Return** to start the requirement checks.

- 6 After the requirement checks are complete, press **Return** to start upgrading the packages. If you are upgrading multiple nodes, you have the option of upgrading them simultaneously. You will be prompted after the upgrade is complete.
- 7 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
- 8 (Optional) If you are going to upgrade your operating system, then upgrade your operating system, and patch it to a kernel version.
- 9 Shut down and reboot the system.

**To bring the upgraded cluster online and restore components**

- 1 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.
- 2 If necessary, reinstate any missing mount points in the `/etc/vfstab` file on each node.
- 3 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.
- 4 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 5 If you have stopped any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

- 6 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem  
# mount /checkpoint_name
```

- 7 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

- 8 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

# Verifying the Storage Foundation installation

This chapter includes the following topics:

- [Verifying that the products were installed](#)
- [Installation log files](#)
- [Checking Volume Manager processes](#)
- [Checking Veritas File System installation](#)
- [Verifying the configuration files for Storage Foundation Cluster File System](#)
- [Verifying agent configuration for Storage Foundation Cluster File System](#)
- [Synchronizing time on Cluster File Systems](#)
- [Configuring VCS for Storage Foundation Cluster File System](#)

## Verifying that the products were installed

Verify that the Veritas Storage Foundation products are installed.

Use the `pkginfo` command to check which packages have been installed.

```
# pkginfo -l VRTSvlic package_name package_name ...
```

Use the following sections to further verify the product installation.

## Installation log files

After every product installation, the installer creates three text files:

- Installation log file
- Response file
- Summary file

The name and location of each file is displayed at the end of a product installation, and are always located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

## Using the installation log file

The installation log file contains all commands executed during the procedure, their output, and errors generated by the commands. This file is for debugging installation problems and can be used for analysis by Veritas Support.

## Using the response file

The response file contains the configuration information that you entered during the procedure. You can use the response file for future installation procedures by invoking an installation script with the `responsefile` option. The response file passes arguments to the script to automate the installation of that product. You can edit the file to automate installation and configuration of additional systems.

## Using the summary file

The summary file contains the results of the installation by the common product installer or product installation scripts. The summary includes the list of the packages and patches, and the status (success or failure) of each package or patch. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

# Checking Volume Manager processes

Use the following procedure to verify that Volume Manager processes are running.

**To confirm that key Volume Manager processes are running**

- ◆ Type the following command:

```
# ps -e | grep vx
```

Entries for the `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxsmd`, `vxpall`, `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 Veritas File System 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 load and then unload the file system module to complete the process.

See [“Loading and unloading the file system module”](#) on page 128.

### Verifying command installation

**Table 9-1** lists the directories with Veritas File System commands.

**Table 9-1** 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 [“Setting environment variables”](#) on page 23.

## Verifying the configuration files for Storage Foundation Cluster File System

You can inspect the contents of the configuration files that were installed and modified after a successful installation process. These files reflect the configuration based on the information you supplied.

### To verify the configuration files

- 1 Log in as superuser to any system in the cluster.
- 2 Set up your environment `PATH` variable.

```
# export PATH=$PATH:/sbin:/usr/sbin:/opt/VRTS/bin
```

## Low Latency Transport configuration files

The following files are required by the VCS communication services for Low Latency Transport (LLT).

### `/etc/llhosts`

The file `llhosts(4M)` is a database, containing one entry per system, that links the LLT system ID (in the first column) with the LLT host name. This file is identical on each system in the cluster.

For example, the file `/etc/llhosts` contains entries that resemble:

```
0    system01
1    system02
```

## **/etc/llttab**

The file `llttab(4M)` contains information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the network links that correspond to the specific system.

For example, the file `/etc/llttab` contains entries that resemble:

```
set-node system01
set-cluster 100

link lan1 lan:1 - ether - -
link lan2 lan:2 - ether - -
```

The first line identifies the local system name. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines, beginning with the `link` command, identify the two network cards used by the LLT protocol.

See the `llttab(4M)` manual page.

## Checking Low Latency Transport operation

Use the `lltstat` command to verify that links are active for LLT. This command returns information about the links for LLT for the system on which it is typed. See the `lltstat(1M)` manual page.

In the following example, `lltstat -n` is typed on each system in the cluster.

### To check LLT operation

#### 1 Log into system01.

```
# lltstat -n
```

Output resembles:

```
LLT node information:
Node                               State           Links
*  0  system01                      OPEN            2
   1  system02                      OPEN            2
```

#### 2 Log into system02.

```
# lltstat -n
```

Output resembles:

```
LLT node information:
Node                               State           Links
   0  system01                      OPEN            2
*  1  system02                      OPEN            2
```

Each system has two links and that each system is in the OPEN state. An asterisk (\*) denotes the system on which the command is typed.

With LLT configured correctly, the output of `lltstat -n` shows all of the systems in the cluster and two links for each system. If the output shows otherwise, you can use the verbose option of `lltstat`. For example, type `lltstat -nvv | more` on a system to view additional information about LLT. In the following example, `lltstat -nvv | more` is typed on a system in a two-node cluster.

**3** Log into system01.

```
# lltstat -nvv | more
```

Output resembles:

Node	State	Link	Status	Address	
*0	system01	OPEN	lan1	UP	08:00:20:93:0E:34
			lan2	UP	08:00:20:93:0E:34
1	system02	OPEN	lan1	UP	08:00:20:8F:D1:F2
			lan2	DOWN	08:00:20:8F:D1:F2
2	CONNWAIT				
			lan1	DOWN	
			lan2	DOWN	
.					
.					
.					
31	CONNWAIT				
			lan1	DOWN	
			lan2	DOWN	

The output lists 32 nodes. It reports on the two cluster nodes, system01 and system02, plus non-existent nodes. For each correctly configured system, the information shows a state of OPEN, a status for each link of UP, and an address for each link. However, in the example above, the output shows that for node system02, the private network may have failed earlier, or the information in `/etc/llttab` may be incorrect.

To obtain information about the ports open for LLT, type `lltstat -p` on any system. In the following example, `lltstat -p` is typed on one system in the cluster.

#### 4 Log into system01.

```
# lltstat -p
```

Output resembles:

```
LLT port information:
Port      Usage      Cookie
0         gab        0x0
          opens:      0 1 3 4 5 6 7 8 9 10 11 12 13...
          connects: 0 1
```

The two systems with node ID's 0 and 1 are connected.

See “[/etc/llthosts](#)” on page 260.

## Group Membership and Atomic Broadcast configuration files

The following files are required by the VCS communication services for Group Membership and Atomic Broadcast (GAB).

### **/etc/gabtab**

After installation, the file `/etc/gabtab` contains a `gabconfig(1M)` command that configures the GAB driver for use.

The file `/etc/gabtab` contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

where the `-c` option configures the driver for use and `-nN` specifies that the cluster will not be formed until at least  $N$  systems are ready to form the cluster.  $N$  is the number of systems in the cluster.

## Checking Group Membership and Atomic Broadcast operation

This section describes how to check GAB operation.

**To check GAB operation**

- ◆ Enter the following command on each node in the cluster.

```
# gabconfig -a
```

If GAB is operational, the following output displays with GAB port membership information:

```
GAB Port Memberships
=====
Port a gen 1bbf01 membership 01
Port b gen 1bbf06 membership 01
Port f gen 1bbf0f membership 01
Port h gen 1bbf03 membership 01
Port v gen 1bbf0b membership 01
Port w gen 1bbf0d membership 01
```

If GAB is not operational, the following output display with no GAB port membership information:

```
GAB Port Memberships
=====
```

See the *Veritas Cluster Server User's Guide*.

## Checking cluster operation

This section describes how to check cluster operation.

**To check cluster operation**

- 1 Enter the following command on any system:

```
# hastatus -summary
```

The output for an SFCFS HA installation resembles:

```
-- SYSTEM STATE
-- System                               State                               Frozen

A  system01                             RUNNING                            0
A  system02                             RUNNING                            0

-- GROUP STATE
-- Group      System      Probed AutoDisabled  State

B  cvm         system01   Y      N                  ONLINE
B  cvm         system02   Y      N                  OFFLINE
```

If the State value is running, VCS is successfully installed and running on that node. The group state lists the cvm group, which is online on system01 and offline on system02.

See the `hastatus(1M)` manual page.

See the *Veritas Cluster Server User's Guide*.

- 2 Enter the following command on any systems:

```
# hasys -display
```

See the *Veritas Cluster Server User's Guide*.

For more information on the `hasys -display` command, see the `hasys(1M)` manual page.

The example shows the output of system01. The list continues with similar information for system02 (not shown) and any other systems in the cluster. The output should be similar on each system.

[Table 9-2](#) shows on each system, the output should be similar:

**Table 9-2** System sample output for `hasys -display`

#System	Attribute	Value
system01	AgentsStopped	0
system01	AvailableCapacity	1

**Table 9-2** System sample output for `hasys -display` (continued)

#System	Attribute	Value
system01	Capacity	1
system01	ConfigBlockCount	54
system01	ConfigChecksum	29776
system01	ConfigDiskState	CURRENT
system01	ConfigFile	/etc/VRTSvcs/conf/config
system01	ConfigInfoCnt	0
system01	ConfigModDate	Tues June 25 23:00:00 2006
system01	CurrentLimits	
system01	DiskHbStatus	
system01	DynamicLoad	0
system01	Frozen	0
system01	GUIIPAddr	
system01	LLTNodeId	0
system01	Limits	
system01	LoadTimeCounter	1890
system01	LoadTimeThreshold	600
system01	LoadWarningLevel	80
system01	MajorVersion	2
system01	MinorVersion	0
system01	NodeId	0
system01	OnGrpCnt	1
system01	ShutdownTimeout	60
system01	SourceFile	./main.cf
system01	SysName	system01

**Table 9-2** System sample output for `hasys -display` (continued)

#System	Attribute	Value
system01	SysState	RUNNING
system01	SystemLocation	
system01	SystemOwner	
system01	TFrozen	0
system01	TRSE	0
system01	UpDownState	Up
system01	UserInt	0
system01	UserStr	

## Verifying agent configuration for Storage Foundation Cluster File System

This section describes how to verify the agent configuration.

### To verify the agent configuration

- ◆ Enter the cluster status command from any node in the cluster:

```
# cfscluster status
```

Output resembles:

```
Node           : system01
Cluster Manager : running
CVM state      : running
No mount point registered with cluster configuration
```

```
Node           : system02
Cluster Manager : running
CVM state      : running
No mount point registered with cluster configuration
```

## Synchronizing time on Cluster File Systems

SFCFS requires that the system clocks on all nodes are synchronized using some external component such as the Network Time Protocol (NTP) daemon. If the nodes are not in sync, timestamps for change (`ctime`) and modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

## Configuring VCS for Storage Foundation Cluster File System

Configuring VCS means conveying to the VCS engine the definitions of the cluster, service groups, resources, and resource dependencies. VCS uses two configuration files in a default configuration:

- The `main.cf` file defines the entire cluster.
- The `types.cf` file defines the resource types.

By default, both files reside in the directory `/etc/VRTSvcs/conf/config`.

Additional files similar to `types.cf` may be present if agents have been added, such as `Oracletypes.cf`.

In a VCS cluster, the first system to be brought online reads the configuration file and creates an internal (in-memory) representation of the configuration. Systems brought online after the first system derive their information from systems running in the cluster. You must stop the cluster while you are modifying the files from the command line. Changes made by editing the configuration files take effect when the cluster is restarted. The node on which the changes were made should be the first node to be brought back online.

### main.cf file

The VCS configuration file `main.cf` is created during the installation procedure. After installation, the `main.cf` file contains the base definitions of the cluster and its nodes. Additionally, the file `types.cf` listed in the include statement defines the bundled agents for VCS resources.

See the *Veritas Cluster Server User's Guide*.

A typical VCS configuration file for SFCFS file resembles:

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTTypes.cf"
```

```
cluster sfcfs_1 (
    HacliUserLevel = COMMANDROOT
)

system thor150 (
)

system thor151 (
)

group cvm (
    SystemList = { thor150 = 0, thor151 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { thor150, thor151 }
)

CVMCluster cvm_clus (
    CVMClustName = sfcfs_1
    CVMNodeId = { thor150 = 0, thor151 = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

cvm_clus requires cvm_vxconfigd

// resource dependency tree
//
//     group cvm
//     {
//         CVMCluster cvm_clus
//         {
//             CVMVxconfigd cvm_vxconfigd
//         }
//     }
// }
```

## Storage Foundation Cluster File System HA Only

If you configured VCS Cluster Manager (Web Console), a service group, "ClusterService," was created that includes IP, Process, and Notifier resources. These resources were configured according to information you provided during the installation procedure. A resource dependency was also created.

## Veritas Cluster Server application failover services

If you installed SFCFS HA, you can begin implementing the application monitoring failover services provided by the Veritas Cluster Server. Information about setting up VCS services is beyond the scope of this document.

See the *Veritas Cluster Server* documentation.



# Uninstalling Storage Foundation

This chapter includes the following topics:

- [About removing Veritas Storage Foundation](#)
- [Dropping the repository database for DB2 and Oracle](#)
- [Shutting down cluster operations](#)
- [Shutting down Veritas Volume Manager](#)
- [Uninstalling Veritas Storage Foundation](#)
- [Uninstalling Storage Foundation Cluster File System](#)
- [Uninstalling the VCS agents for VVR](#)
- [Uninstalling Veritas Volume Replicator \(VVR\)](#)
- [Uninstalling language packages](#)
- [Uninstalling Storage Foundation using the pkgrm command](#)
- [Removing the Veritas Enterprise Administrator client](#)

## About removing Veritas Storage Foundation

This section covers uninstallation requirements and steps to uninstall the Veritas software.

### Uninstallation requirements for Solaris

Review the uninstallation requirements before removing the Veritas software.

## Remote uninstallation

For information on remote uninstallations, refer to the information about remote shells.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

## Preparing to remove Veritas Volume Manager

This section describes the steps you need to take before removing Veritas Volume Manager (VxVM) to preserve the contents of the volumes.

---

**Warning:** Failure to follow the preparations in this section might result in unexpected behavior.

---

### Moving volumes from an encapsulated root disk

Use the following procedure to move volumes from an encapsulated root disk.

### To uninstall VxVM if `root`, `swap`, `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
```

---

**Warning:** 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 the following 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.  
See [“Moving volumes to disk partitions”](#) on page 275.  
Otherwise, shut down VxVM.  
See [“Shutting down Veritas Volume Manager”](#) on page 284.

### Moving volumes to disk partitions

Use the following procedure to move volumes incrementally 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 VxVM 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 VxVM 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 VxVM control.

To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.

```
# vxprint -F '%snum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VxVM control.

```
# vxdg rmdisk diskname
# vxdisk rm devname
```

Use the free space created for adding the data from the next volume you want to remove.

- 11** After you successfully convert all volumes into disk partitions, reboot the system.
- 12** After the reboot, make sure none of the volumes are open by using the `vxprint` command.

```
# vxprint -Aht -e v_open
```

- 13** If any volumes remain open, repeat the steps listed above.

### Example of moving volumes to disk partitions on Solaris

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 `vol101` and `disk3` is a free disk. The data on `vol101` is copied to `disk3` using `vxevac`.

These are the contents of the disk group `voldg` before the data on `vol101` 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  voll1          -                ENABLED  ACTIVE   4196448  ROUND   -        fsgen
pl pl1           voll1           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         DATAVOL     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      NVOLLLAYR      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  voll1          -                ENABLED  ACTIVE   4196448  ROUND   -        fsgen
pl pl1           voll1           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	DATAVOL	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 voll	-	ENABLED	ACTIVE	4196448	ROUND	-	fsgen
pl pl1	voll	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 VxVM 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

# vxdg rmdisk disk1
# vxdg rmdisk disk2
# vxdisk rm c1t12d0
# vxdisk rm c1t14d0
```

Verify that the evacuated disks have been removed from VxVM 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 voll
/voll on /dev/vx/dsk/voldg/voll
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 (clt12d0s1)`.

```
# 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. clt3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
2. clt9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
3. clt10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
4. clt11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
5. clt12d0 <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[wm]:
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 VxVM.

```
# vxedit -rf rm /dev/vx/dsk/voldg/vol101
```

To complete the procedure, follow the remaining steps.

## Preparing to remove 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.

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
```

---

**Warning:** Do not use the `-force` option when executing `hastop`. This will leave all service groups online and shut down VCS, causing undesired results during uninstallation of the packages.

---

## Shutting down Veritas Volume Manager

Use the following procedure to shut down Veritas Volume Manager.

**To shut down Veritas Volume Manager**

- ◆ Enter the `vxctl` and `vxiod` commands as follows:

```
# vxctl stop  
# vxiod -f set 0
```

## 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.

If you installed the Chinese version of Veritas Volume Manager, the uninstallation script does not remove the 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 VxVM package (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

See [“Preparing to remove Veritas Volume Manager”](#) on page 274.

- 4 Stop the VEA Service.

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 5 In a stand-alone configuration, if you are uninstalling Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, stop the repository database and unmount the database repository volume.

```
# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb  
# /opt/VRTSdbcom/config/sfua_rep_mount stop
```

- 6 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 7 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
```

- 8 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall Storage Foundation, for example, `host1`:

```
Enter the system names separated by spaces from which to
uninstall Storage Foundation: host1
```

- 9 The uninstall script prompts you to confirm the uninstall. If you respond yes, the processes are stopped and the packages are uninstalled.

The uninstall script creates log files and displays the location of the log files.

- 10 Most packages have kernel components. In order to ensure complete removal, a system reboot is recommended after all packages have been removed.

- 11 To verify the removal of the packages, use the `pkginfo` command.

```
# pkginfo | grep VRTS
```

## Uninstalling Storage Foundation Cluster File System

If you need to uninstall SFCFS software. Use the `uninstallsfcfs` script.

### To uninstall SFCFS HA

- 1 Log in as superuser.

Do not use the `hastop -force` command to stop VCS.

- 2 Change directory to `/opt/VRTS/install`:

```
# cd /opt/VRTS/install
```

- 3 Run the `uninstallsfcfs` command to uninstall SFCFS. The `uninstallsfcfs` script uses `ssh` to communicate with remote nodes as default:

```
# ./uninstallsfcfs
```

If you want to use `rsh` you must specify on the command line:

```
# ./uninstallsfcfs -rsh
```

- 4 Enter the system names to uninstall SFCFS.

```
Enter the system names separated by spaces on which to  
uninstall SFCFS: system01 system02
```

- 5 Enter `y` to uninstall SFCFS.

```
Are you sure you want to uninstall SFCFS? [y,n,q] (y)
```

## Uninstalling the VCS agents for VVR

To uninstall the VCS Agents for VVR, you must first disable the agents.

If VCS Agents for VVR are not installed on your system, go to [Uninstalling Veritas Volume Replicator \(VVR\)](#).

### Disabling the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

### To disable the agents

- 1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
# hagrps -state service_group -sys system_name
```

If none of the service groups is online, skip to 3.

- 2 If the service group is online, take it offline.

To take the service group offline without bringing it online on any other system in the cluster, enter:

```
# hagrps -offline service_group -sys system_name
```

- 3 Stop the agent on the system by entering:

```
# haagent -stop agent_name -sys system_name
```

When you get the message `Please look for messages in the log file, check the file /var/VRTSvcs/log/engine_A.log for a message confirming that each agent has stopped.`

You can also use the `ps` command to confirm that the agent is stopped.

- 4 Remove the system from the `SystemList` of the service group. If you disable the agent on all the systems in the `SystemList`, you can also remove the service groups and resource types from the VCS configuration.

Read information on administering VCS from the command line.

Refer to the *Veritas Cluster Server User's Guide*.

## Uninstalling Veritas Volume Replicator (VVR)

This section describes how to uninstall Volume Replicator.

---

**Note:** If you are upgrading Veritas Volume Replicator, do not remove the Replicated Data Set, but only remove the VVR packages.

---

Uninstalling Veritas Volume Replicator (VVR) involves performing the following tasks in the order indicated:

- [Removing the Replicated Data Set](#)
- [Removing the VVR packages](#)

For more information about VVR commands, refer to the *Veritas Volume Replicator Administrator's Guide*.

## Removing the Replicated Data Set

This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

### To remove the Replicated Data Set

- 1 Verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

If the Secondary is not required to be up-to-date, proceed to [2](#) and stop replication using the `-f` option with the `vradmin stoprep` command.

- 2 Stop replication to the Secondary by issuing the following command on any host in the RDS:

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

```
# vradmin -g diskgroup stoprep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 3 Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

```
# vradmin -g diskgroup delsec local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 4 Remove the Primary from the RDS by issuing the following command on the Primary:

```
# vradmin -g diskgroup delpri local_rvgname
```

When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

The RDS is removed. Go on to uninstalling Volume Manager to uninstall VVR.

- 5 If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

```
# vxedit -r -g diskgroup rm srl_name
```

- 6 Uninstall the VVR packages.

See [“Removing the VVR packages”](#) on page 290.

## Removing the VVR packages

Use the uninstall program to remove the VVR software packages.

### To remove the VVR packages

- 1 Insert the software disc, mount it, and enter the following commands:

```
# cd /disc_path/pkgs
```

```
# ./installer
```

- 2 Select Uninstall from the menu.

- 3 Select VVR.

The program prompts you to confirm whether you want to remove the packages that are being used by other Veritas products.

- 4 Answer the set of questions depending on your requirements. Note that if you uninstall the `VRTSvxvm` package you will not be able to use the Veritas Volume Manager functionality.

The program asks you to confirm that you want to remove VVR and then removes all the packages except the infrastructure packages. If open volumes exist, the program prompts you to stop the open volumes and unmount the file systems.

The output is similar to the following:

```
uninstallvvr is now ready to uninstall VVR packages.  
All VVR processes that are currently running will be stopped.  
Are you sure you want to uninstall VVR packages? [y,n,q] (y)
```

- 5 Press Return to continue.
- 6 Confirm the packages have been removed.

```
# pkginfo | grep VRTS
```

If you do not have any other Veritas products installed on the system, you can remove the `/etc/vx` directory, the `/usr/lib/vxvm` directory, and the `/opt/VRTS*` directories.

## Additional ways to remove VVR packages

Before removing the packages, determine whether any other Veritas products are installed on your system. Other products might depend on the packages you may be removing. A warning appears when you try to remove packages that are being used by other products. If the language packages are installed, we recommend that you remove the corresponding language packages before removing the VVR packages.

---

**Note:** If `/opt` is on its own partition, you must mount `/opt` manually before you run the `pkgrm` command.

---

### To remove the VVR packages using the `pkgrm` command

- 1 Use the `pkgrm` command to remove the installed Veritas Volume Replicator software packages. Remove the packages in the order shown:

```
# pkgrm VRTSvmdoc VRTSvrdoc VRTSvmman VRTSvcsvr VRTSap VRTStep
```

You can also include `VRTSvlic` in the removal line, if you have not installed any other packages that use `VRTSvlic`.

- 2 Remove the Veritas Provider Packages, Veritas Virtual Disk Management Services Provider and Veritas Volume Replicator Management Services Provider, using the following commands:

```
# pkgrm VRTSvmpro
```

```
# pkgrm VRTSvrpro
```

- 3 Remove the Veritas Enterprise Administrator packages using the following commands:

```
# pkgrm VRTSob
```

```
# pkgrm VRTSobgui
```

- 4 Remove the Windows Client software. Perform the following tasks in the order indicated:

- Click **Start > Settings > Control Panel > Add/Remove Software**
- Choose Veritas Enterprise Administrator for removal.

- 5 Remove the Veritas Volume Replicator Web GUI (VRW) Application package:

```
# pkgrm VRTSvrw
```

---

**Note:** The Veritas Web GUI Engine, `VRTSweb` is used by other Veritas products, such as GCM or QuickStart, that have Web GUIs. Do not perform 6 if you have other Veritas products with Web GUIs installed on your system.

---

- 6 Remove the Veritas Web GUI Engine `VRTSweb` by entering the following command:

```
# pkgrm VRTSweb
```

- 7 Remove `VRTSvxvm`.

For instructions, see the *Veritas Storage Foundation Installation Guide*.

## 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 Veritas software.

See “[Uninstalling Veritas Storage Foundation](#)” on page 284.

The product installer and uninstallation scripts do not have an option to remove only language packages.

## Uninstalling language packages using the `pkgrm` command

If you would like to remove only the language packages, you can do so with the `pkgrm` command.

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 ...
```

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
```

# Uninstalling Storage Foundation using the `pkgrm` command

Use the following procedure to uninstall Storage Foundation using the `pkgrm` command.

If you are uninstalling Veritas Storage Foundation using the `pkgrm` command, the packages must be removed in a specific order, or else the uninstallation will fail. Removing the packages out of order will result in some errors, including possible core dumps, although the packages will still be removed.

## To uninstall Storage Foundation

- 1 Unmount all VxFS file systems and Storage Checkpoints, and close all VxVM volumes.

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 and Storage Checkpoints.

```
# umount /mount_point
```

- 3 Stop all applications from accessing VxVM volumes, and close all VxVM volumes.

- 4 Stop various daemons, if applicable.

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c stop
# /opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c stop
# /opt/VRTSobc/pal33/bin/vxpalctrl -a gridnode -c stop
# /opt/VRTS/bin/vxsvcctrl stop
```

- 5 Remove the packages in the following order:

```
# pkgrm VRTSmapro VRTSgapms VRTSvxmsa VRTSfasdc \
VRTSffas VRTSvail VRTSfsmnd VRTSfssdk VRTSfsdoc \
VRTSfsman VRTSvxfs VRTSvrdoc VRTSvrw VRTSweb VRTSjrel5 \
VRTSjre VRTSvcsvr VRTSvrpro VRTSddlpr VRTSvdid \
VRTSvsvc VRTSvmpo VRTSalloc VRTSdcli VRTSvmdoc \
VRTSvman VRTSfspro VRTSdsa VRTSvxvm SYMClma VRTSspt \
VRTSaa VRTSmh VRTSccg VRTSobgui VRTSob VRTSobc33 \
VRTSat VRTSsmf VRTSspbv VRTSicsco VRTSperl
```

# Removing the Veritas Enterprise Administrator client

You should also remove the client software from any machines you used to access the Veritas software.

## To remove the VEA client 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.



# Installation scripts

This appendix includes the following topics:

- [About installation scripts](#)
- [Installation script options](#)

## About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.0MP3 provides several installation scripts.

To install the Veritas Storage Foundation products 5.0MP3 on a system that already has Veritas Storage Foundation 5.0, including maintenance packs and rolling patches, use the `installmp` script.

To install a fresh installation on a system, or to upgrade from Veritas Storage Foundation version prior to 5.0, the recommended installation method is to use the common product installer. To use the common product installer, run the `installer` command.

See [“About the common product installer”](#) on page 47.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from an electronic download site, which does not include the common product installer, use the appropriate product installation script.

The following product installation scripts are available:

Veritas Cluster Server (VCS)	<code>installvcs</code>
Veritas Volume Replicator (VVR)	<code>installvvr</code>
Veritas Storage Foundation (SF)	<code>installsf</code>

Veritas Storage Foundation for Oracle (SFORA)	<code>installsfora</code>
Veritas Storage Foundation for DB2 (SFDB2)	<code>installsfdb2</code>
Veritas Storage Foundation for Sybase (SFSYB)	<code>installsfsyb</code>
Veritas Storage Foundation Cluster File System (SFCFS)	<code>installsfcfs</code>
Symantec Product Authentication Service (AT)	<code>installat</code>
Veritas Volume Manager.	<code>installvm</code>
Veritas File System.	<code>installfs</code>

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

## Installation script options

[Table A-1](#) shows command line options for the product installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts, except where otherwise noted.

See [“About installation scripts”](#) on page 297.

**Table A-1** Available command line options

Command Line Option	Function
<code>system1 system2...</code>	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.
<code>-configure</code>	Configures the product after installation.
<code>-enckeyfile encryption_key_file</code>	Specifies the location of a file containing the key to decrypt encrypted passwords stored in response files. See the <code>-responsefile</code> and the <code>-encrypt</code> options.

**Table A-1** Available command line options (*continued*)

Command Line Option	Function
<code>-encrypt <i>password</i></code>	Encrypts <i>password</i> using the encryption key provided with the <code>-enckeyfile</code> option so that the encrypted password can be stored in response files.
<code>-hostfile <i>full_path_to_file</i></code>	Specifies the location of a file that contains a list of hostnames on which to install.
<code>-installpkgs</code>	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 <code>requiredpkgs</code> option.
<code>-installonly</code>	Installs packages, but does not configure the product.
<code>-jumpstart <i>dir_path</i></code>	Produces a sample finish file for Solaris JumpStart installation. The <i>dir_path</i> indicates the path to the directory in which to create the finish file.
<code>-keyfile <i>ssh_key_file</i></code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i <i>ssh_key_file</i></code> to every SSH invocation.
<code>-license</code>	Registers or updates product licenses on the specified systems.
<code>-logpath <i>log_path</i></code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
<code>-noextrapkgs</code>	Additional packages can be installed so that you can upgrade to another Symantec product simply by installing a new license. The <code>noextrapkgs</code> option bypasses installation of extra product packages to simplify future maintenance updates.

**Table A-1** Available command line options (*continued*)

Command Line Option	Function
-nohapkgs	Limits the list of Storage Foundation packages to exclude the Veritas Cluster Server packages.  This option only applies to the <code>installsf</code> script when one of the following options is specified: <ul style="list-style-type: none"> <li>■ -installpkgs</li> <li>■ -requiredpkgs</li> <li>■ -jumpstart</li> </ul>
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-nooptionalpkgs	Bypasses installation of optional product packages such as manual pages.
-nostart	Bypasses startup of the product following installation and configuration.
-osversion	Displays only the packages and the patches which apply to the specified OS version. Valid values are: <code>sol8</code> , <code>sol9</code> , or <code>sol10</code>  This option only applies when one of the following options is specified: <ul style="list-style-type: none"> <li>■ -installpkgs</li> <li>■ -requiredpkgs</li> <li>■ -jumpstart</li> </ul>
-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 by 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 by all specified installation systems.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.

**Table A-1** Available command line options (*continued*)

Command Line Option	Function
-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> ]	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. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.  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.
-rootpath <i>root_path</i>	Specifies an alternative root directory on which to install packages.  On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code> command.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.  See “ <a href="#">Configuring secure shell (ssh) or remote shell before installing products</a> ” on page 24.
-security	Enable or disable Symantec Product Authentication Service in a VCS cluster that is running. Install and configure Root Broker for Symantec Product Authentication Service. You can specify this option with the <code>installvcs</code> , <code>installsf</code> or <code>installsfdfs</code> scripts.  For more information about Symantec Product Authentication Service in a VCS cluster, see the <i>Veritas Cluster Server Installation Guide</i> .

**Table A-1** Available command line options (*continued*)

Command Line Option	Function
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-timeout <i>timeout_value</i>	Specifies the timeout (in seconds) that the installer uses for each command it issues during the installation. The default timeout is set to 600 secs. Use the -timeout option to override the default value.
-tmppath <i>tmp_path</i>	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.
-verbose	Displays details during installation of product packages. By default, the installation displays only a progress bar.

# Storage Foundation and High Availability components

This appendix includes the following topics:

- [Veritas Storage Foundation installation packages](#)
- [Obsolete packages in Storage Foundation](#)

## Veritas Storage Foundation installation packages

[Table B-1](#) shows the package name and contents for each English language package for Veritas Storage Foundation, Veritas Storage Foundation High Availability, Veritas Storage Foundation Cluster File System, and Veritas Storage Foundation for databases.

**Table B-1** Storage Foundation packages

package	Contents	Required/Optional
Veritas Volume Manager		
VRTSalloc	Veritas Volume Manager Veritas Intelligent Storage Provisioning  Provides the volume tagging features, which is required for dynamic storage tiering (DST).	Required

**Table B-1** Storage Foundation packages (*continued*)

<b>package</b>	<b>Contents</b>	<b>Required/Optional</b>
VRTSdcli	Veritas Distributed Command Line Interface	Required
VRTSddlpr	Veritas Device Discovery Layer Services Provider  Provides the necessary management backend required to administer VxVM DMP features and objects like enclosures, controllers, and paths from the GUI.	Required
VRTSvdid	Veritas Device Identification API	Required
VRTSvmman	Veritas Volume Manager Manual Pages	Optional
VRTSvmpro	Veritas Volume Manager Management Services Provider  Provides the necessary management backend required to administer VxVM from the GUI.	Required
VRTSvxvm	Veritas Volume Manager binaries	Required
Veritas File System		
VRTSfsman	Veritas File System Manual Pages	Optional
VRTSfsmnd	Veritas File System Software Developer Kit Manual Pages	Optional
VRTSfspro	Veritas File System Management Services Provider  Provides the necessary management that is required to administer VxFS and other platform filesystems to manage from the GUI. Also, provides dynamic storage tiering (DST) capability that allows users to do policy based control for data placement.	Required

**Table B-1** Storage Foundation packages (*continued*)

package	Contents	Required/Optional
VRTSfssdk	Veritas File System Software Developer Kit  For VxFS APIs, the package contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.	Required
VRTSvxfs	Veritas File System binaries  Required for VxFS file system support.	Required
Storage Foundation Cluster File System		
VRTScavf	Veritas Cluster Server Agents for Storage Foundation Cluster File System	Required
VRTSglm	Veritas Group Lock Manager for Storage Foundation Cluster File System	Required
VRTSgms	Veritas Group Messaging Services for Storage Foundation Cluster File System	Required
Databases		
VRTSd2gui	Veritas Storage Foundation for DB2 Graphical User Interface	Required (for Storage Foundation for DB2)
VRTSdb2ed	Veritas Storage Foundation for DB2	Required (for Storage Foundation for DB2)
VRTSdbcom	Veritas Storage Foundation Common Utilities for Databases	Required (for Storage Foundation for databases)
VRTSdbed	Veritas Storage Foundation for Oracle	Required (for Storage Foundation for Oracle)

**Table B-1** Storage Foundation packages (*continued*)

<b>package</b>	<b>Contents</b>	<b>Required/Optional</b>
VRTSodm	ODM Driver for VxFS Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager allows Oracle9i and 10g to improve performance and manage system bandwidth.	Required (for Storage Foundation for Oracle)
VRTSorgui	Veritas Storage Foundation for Oracle Graphical User Interface	Required (for Storage Foundation for Oracle)
VRTSvxmsa	Veritas Mapping Service, Application Libraries	Required (for DB2 and Oracle products)
Veritas Enterprise Administrator		
VRTSaa	Veritas Enterprise Administrator Action Agent	Required
VRTSccg	Veritas Enterprise Administrator Central Control Grid	Required
VRTSob	Veritas Enterprise Administrator	Required
VRTSobc33	Veritas Enterprise Administrator Core	Required
VRTSobweb	Veritas Enterprise Administrator Web Console	Required
VRTSobgui	Veritas Enterprise Administrator	Optional
Infrastructure		

**Table B-1** Storage Foundation packages (*continued*)

package	Contents	Required/Optional
VRTSat	Symantec Product Authentication Service  Installs the Symantec Product Authentication Service, which provides authentication services to other Symantec products. This package contains a server and client component. The server provides services for a root broker, authentication broker, or both. The client allows Symantec products to communicate with the brokers.	Required
VRTSgapms	Veritas Generic Array Plugin	Required
VRTSicsco	Symantec Infrastructure Core Services Common	Required
VRTSvail	Veritas Array Integration Layer	Required
High Availability	Note: some of these packages are also required for Storage Foundation Cluster File System.	
VRTSaclib	Veritas Application Competency Center Library  VRTSaclib is a set of Perl modules that many cluster server agents use.	Required Depends on VRTSvcS.
VRTScmccc	Veritas Cluster Server Management Console Cluster Connector	Optional
VRTScmcs	Veritas Cluster Management Console for single cluster environments	Optional
VRTScscm	Veritas Cluster Server Cluster Manager	Required Depends on VRTSvcS and VRTSjre15.

**Table B-1** Storage Foundation packages (*continued*)

<b>package</b>	<b>Contents</b>	<b>Required/Optional</b>
VRTSscsw	Veritas Cluster Server configuration wizards	Required Depends on VRTSvcsag and VRTSjre15.
VRTSscsocw	Veritas agent for Oracle and SF Oracle RAC configuration wizards.	Optional for VCS. Required to use VCS with the high availability agent for Oracle.
VRTSscsim	Veritas Cluster Server Simulator	Optional
VRTScutil	Veritas Cluster Server Utilities	Required Depends on VRTSvcs.
VRTSgab	Veritas Cluster Server group membership and atomic broadcast services	Required Depends on VRTSllt.
VRTSjre	Veritas Java Runtime Environment Redistribution	Required
VRTSjre15	Veritas Java Runtime Environment Redistribution  This package installs the Java Runtime Environment for all Symantec products that require Java.	Required
VRTSllt	Veritas Cluster Server low-latency transport	Required
VRTSvcs	Veritas Cluster Server	Required Depends on VRTSut, VRTSperl, VRTSvxfen, VRTSgab, and VRTSllt.
VRTSvcsag	Veritas Cluster Server Bundled Agents	Required Depends on VRTSvcs.
VRTSvcsdb	Veritas High Availability Agent for DB2	Optional for VCS. Required to use VCS with the high availability agent for DB2. Depends on VRTSvcs.

**Table B-1** Storage Foundation packages (*continued*)

package	Contents	Required/Optional
VRTSvcsmsg	Veritas Cluster Server English message catalogs	Required Depends on VRTSvcs.
VRTSvcsmn	Manual Pages for Veritas Cluster Server	Optional
VRTSvcsor	Veritas High Availability Agent for Oracle	Optional for VCS. Required to use VCS with the high availability agent for Oracle. Depends on VRTSvcs.
VRTSvcssy	Veritas High Availability Agent for Sybase	Optional for VCS. Required to use VCS with the high availability agent for Sybase. Depends on VRTSvcs.
VRTSvxfen	Veritas I/O Fencing	Required Depends on VRTSgab.
VRTSweb	Symantec Web Server	Required
Veritas Volume Replicator		
VRTSvcsvr	Veritas Cluster Server Agents for VVR	Required
VRTSvrpro	Veritas Volume Replicator Client Extension and Provider for Veritas Enterprise Administrator	Required
VRTSvrw	Veritas Volume Replicator Web Console	Required
Other packages		
SYMClma	Symantec License Inventory Agent	Required
VRTSdbms3	Veritas Shared DBMS	Required
VRTSdsa	Veritas Datacenter Storage Agent	Required

**Table B-1** Storage Foundation packages (*continued*)

<b>package</b>	<b>Contents</b>	<b>Required/Optional</b>
VRTSmapro	Veritas Storage Foundation GUI for Mapping	Required
VRTSmh	Veritas Centralized Management for Storage Foundation Managed Host	Required
VRTSpbx	<p>Symantec Private Branch Exchange</p> <p>This package installs the Symantec Private Branch Exchange, which allows other Symantec products to share a common well-known port for publishing services and communicating.</p>	<p>Required</p> <p>If VRTSpbx is removed, Symantec products that use it are unable to communicate, which can cause the products to stop working.</p> <p>If VRTSat is configured to work with VRTSpbx, and VRTSpbx is removed, VRTSat continues to work. However, the Symantec Product Authentication Service remote administration functionality are not available. Removing VRTSat can affect Symantec products that use the Symantec Product Authentication Service remote administration feature, such as VEA.</p>
VRTSperl	Perl 5.8.8 for Veritas	Required
VRTSsmf	Symantec Service Management Framework	Required
VRTSspt	Veritas Software Support Tools	Required
VRTSvlic	<p>Veritas License Utilities</p> <p>Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.</p>	<p>Required</p> <p>If VRTSvlic is removed, the Storage Foundation products may not be able to access their license information, The products may fail to start or fail to work properly.</p>
VRTScweb	Symantec Web Server	Required
VRTSdcp	Veritas Disk Correlator Provider	Required

**Table B-1** Storage Foundation packages (*continued*)

package	Contents	Required/Optional
VRTSdsm	Veritas Datacenter Storage Manager	Required
VRTSgcscha	Veritas GCS High Availability Agents	Required
VRTSgcspr	Veritas SAN Global Configuration Server Object Bus provider	Required
windows/vrtsobgui.msi	Veritas Enterprise Administrator for Windows	Optional

## Obsolete packages in Storage Foundation

The following packages were included in previous releases of Storage Foundation but are now obsolete:

SYMClma  
VRTSsmf  
VRTScmcm  
VRTSjre  
VRTSvsvc  
VRTSfsdoc  
VRTSvmdoc  
VRTSvrdoc  
VRTSvcsdc  
VRTSdbdoc  
VRTScsdoc  
VRTScfsdc



# Troubleshooting information

This appendix includes the following topics:

- [Troubleshooting information](#)
- [Storage Foundation Cluster File System installation issues](#)
- [Storage Foundation Cluster File System problems](#)

## Troubleshooting information

The `VRTSspt` package provides a group of tools for troubleshooting a system and collecting information on its configuration. The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. The tools are not required for operation of any Veritas product, and they may adversely impact system performance if not used correctly. Veritas provides these tools to analyze systems if you suspect that there are performance problems. The tools should be used only under the direction of a Veritas Technical Support Engineer.

## Storage Foundation Cluster File System installation issues

If you encounter any issues installing SFCFS, refer to the following paragraphs for typical problems and their solutions.

## Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Checking ssh communication with system01 ..... permission denied
installer requires that ssh commands used between systems execute without
prompting for passwords or confirmations. Please run installer again with
the ssh configured for password free logins, or configure rsh and use the
-rsh option.
```

**Suggested solution:** You need to set up the systems to allow remote access using `ssh` or `rsh`.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 24.

---

**Note:** Remove remote shell permissions after completing the SFCFS installation and configuration.

---

## Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
Checking communication with system01 ..... FAILED
System not accessible : system01
```

**Suggested solution:** Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

## Storage Foundation Cluster File System problems

If there is a device failure or controller failure to a device, the file system may become disabled cluster-wide. To address the problem, unmount file system on all the nodes, then run a full `fsck`. When the file system check completes, mount all nodes again.

## Unmount failures

The `umount` command can fail if a reference is being held by an NFS server. Unshare the mount point and try the unmount again.

## Mount failures

Mounting a file system can fail for the following reasons:

- The file system is not using disk layout Version 6 or 7.
- The mount options do not match the options of already mounted nodes.
- A cluster file system is mounted by default with the `qio` option enabled if the node has a Quick I/O for Databases license installed, even if the `qio` mount option was not explicitly specified. If the Quick I/O license is not installed, a cluster file system is mounted without the `qio` option enabled. So if some nodes in the cluster have a Quick I/O license installed and others do not, a cluster mount can succeed on some nodes and fail on others due to different mount options. To avoid this situation, ensure that Quick I/O licensing is uniformly applied, or be careful to mount the cluster file system with the `qio/noqio` option appropriately specified on each node of the cluster.

See the `mount(1M)` manual page.

- A shared CVM volume was not specified.
- The device is still mounted as a local file system somewhere on the cluster. Unmount the device.
- The `fsck` or `mkfs` command is being run on the same volume from another node, or the volume is mounted in non-cluster mode from another node.
- The `vxfsckd` daemon is not running. This typically happens only if the `CFSfsckd` agent was not started correctly.

- If `mount` fails with an error message:

```
vxfs mount: cannot open mnttab
/etc/mnttab is missing or you do not have root privileges.
```

- If `mount` fails with an error message:

```
vxfs mount: device already mounted, ...
```

The device is in use by `mount`, `mkfs` or `fsck` on the same node. This error cannot be generated from another node in the cluster.

- If this error message displays:

```
mount: slow
```

The node may be in the process of joining the cluster.

- If you try to mount a file system that is already mounted without `-o cluster` option (that is, not in shared mode) on another cluster node,

```
# mount -F vxfs /dev/vx/dsk/share/vol01 /vol01
```

The following error message displays:

```
vxfs mount: /dev/vx/dsk/share/vol01 is already mounted,  
/vol01 is busy, allowable number of mount points exceeded,  
or cluster reservation failed for the volume
```

## Command failures

This section describes command failures.

- Manual pages not accessible with the `man` command. Set the `MANPATH` environment variable appropriately. See [“Setting environment variables”](#) on page 23.
- The `mount`, `fsck`, and `mkfs` utilities reserve a shared volume. They fail on volumes that are in use. Be careful when accessing shared volumes with other utilities such as `dd`, it is possible for these commands to destroy data on the disk.
- Running some commands, such as `vxupgrade -n 7/vol02`, can generate the following error message:

```
vxfs vxupgrade: ERROR: not primary in a cluster file system
```

This means that you can run this command only on the primary, that is, the system that mounted this file system first.

## Performance issues

Quick I/O File system performance is adversely affected if a cluster file system is mounted with the `qio` option enabled, but the file system is not used for Quick I/O files. Because `qio` is enabled by default, if you do not intend to use a shared file system for Quick I/O, explicitly specify the `noqio` option when mounting.

## High availability issues

This section describes high availability issues.

### Network partition/jeopardy

Network partition (or split brain) is a condition where a network failure can be misinterpreted as a failure of one or more nodes in a cluster. If one system in the cluster incorrectly assumes that another system failed, it may restart applications

already running on the other system, thereby corrupting data. CFS tries to prevent this by having redundant heartbeat links.

At least one link must be active to maintain the integrity of the cluster. If all the links go down, after the last network link is broken, the node can no longer communicate with other nodes in the cluster. Thus the cluster is in one of two possible states. Either the last network link is broken (called a network partition condition), or the last network link is okay, but the node crashed, in which case it is not a network partition problem. It is not possible to identify whether it is the first or second state, so a kernel message is issued to indicate that a network partition may exist and there is a possibility of data corruption.

Jeopardy is a condition where a node in the cluster has a problem connecting to other nodes. In this situation, the link or disk heartbeat may be down, so a jeopardy warning may be displayed. Specifically, this message appears when a node has only one remaining link to the cluster and that link is a network link. This is considered a critical event because the node may lose its only remaining connection to the network.

---

**Warning:** Do not remove the communication links while shared storage is still connected.

---

## Low memory

Under heavy loads, software that manages heartbeat communication links may not be able to allocate kernel memory. If this occurs, a node halts to avoid any chance of network partitioning. Reduce the load on the node if this happens frequently.

A similar situation may occur if the values in the `/etc/llttab` files on all cluster nodes are not correct or identical.



# Index

## A

- agents
  - disabling 287
- applications, stopping 202

## B

- bootdg 120

## C

- CFS
  - mount and unmount failures 315
  - synchronization 269
  - troubleshooting 314
- cluster functionality
  - enabling 122, 125
  - environment requirements 30
  - shared disks 122, 126
- clusters
  - verifying operation 265
- command failures 316
- commands
  - hacf 234
  - hastatus 266
  - lltconfig 261
  - lltstat 261
  - pkgrm 292
  - vxctl 269
  - vxrootmir 275
- configuration
  - restoring the original 205
- configuration daemon (vxconfigd)
  - starting 116
- configuring
  - new disks 117
  - shared disks 122, 126

## D

- data
  - protecting 123
- default disk group 119

- defaultdg 119
- deleting VVR packages 292
- devices
  - suppress devices 120
- disabling the agents 287
- disk groups
  - bootdg 120
  - default 119
  - nodg 120
  - root 119
  - rootdg 115, 119
- disk space requirements
  - requirements for disk space 43
- disks
  - placing in other disk group 123
- DMP
  - prevent multipathing 120

## F

- Fibre Channel fabric 32
- freezing service groups 202

## G

- gabconfig command
  - in gabtab file 264
- gabtab file
  - verifying after installation 264

## H

- hastatus -summary command 266
- high availability issues 317
  - how memory 317
  - network partition 317
- hot-relocation 116

## I

- I/O daemon (vxiod)
  - starting 117
- Installation Menu
  - product installer 65

- installing VEA
  - planning 34
- installing VVR
  - using the product installer 65

**J**

- jeopardy 317

**L**

- language packages
  - removal 293
- Links
  - private network 261
- LLT
  - verifying 261
- lltconfig command 261
- llhosts file
  - verifying after installation 260
- lltstat command 261
- llttab file
  - verifying after installation 261
- localized environment settings for using VVR
  - settings for using VVR in a localized environment 36

**M**

- manual pages
  - potential problems 316
  - troubleshooting 316
- mirroring
  - root disk 124
- mount command
  - potential problems 315

**N**

- network partition 317
- new disks
  - configuring 117
- nodg 120
- NTP
  - network time protocol daemon 269

**O**

- original configuration
  - restoring the 205

**P**

- packages for VVR
  - decompressing 71
  - removing 290, 292
- pkgm command 292
- planning to upgrade VVR 34
- preinstallation 34
- preparing to upgrade VVR 191, 202
- Prevent Multipathing/Suppress Devices from VxVMbsxd5 s view 120
- preventive maintenance 123
- problems
  - accessing manual pages 316
  - executing file system commands 316
  - mounting and unmounting file systems 315
- product installer
  - using 65
- protecting data 123

**Q**

- Quick I/O
  - performance on CFS 316

**R**

- removing
  - the Replicated Data Set 289
  - VEA for VVR 292
  - VRTSweb 292
  - VVR packages 290, 292
- Replicated Data Set
  - removing the 289
- requirements for disk space
  - disk space requirements 43
- restoring the original configuration 205
- root disk
  - alternate 124
- root disk group 115, 119
- rootdg 119

**S**

- SAN
  - see Storage Area Network 32
- service groups
  - freezing 202
  - unfreezing 204
- settings for using VVR in a localized environment
  - localized environment settings for using VVR 36
- shared disks, configuring 122, 126

- split brain 317
- starting vxconfigd configuration daemon 116
- starting vxiod daemon 117
- stopping
  - applications 202
- Storage Area Network 32
- suppress devices 120

## T

- troubleshooting
  - accessing manual pages 316
  - executing file system commands 316
  - mounting and unmounting file systems 315

## U

- unfreezing service groups 204
- uninstallvvr program 290
- upgrade
  - planning 148
- upgrading
  - clustered environment 127
- upgrading VVR
  - from 4.1 35
  - from releases prior to VVR 3.5mp2 194
  - planning 34, 191
  - preparing 202
  - using upgrade scripts 194
  - when VCS is present 199
  - without using upgrade scripts 199

## V

- VEA
  - client, starting 139
- VEA installation
  - planning 34
- verifying installation
  - kernel component 259
- Veritas Enterprise Administrator
  - removing 292
- vradmin
  - delpri 290
  - stoprep 289
- VRTSweb
  - removing 292
- VVR 4.1
  - planning an upgrade from 35
- vvr\_upgrade\_finish script 206

- vxconfigd configuration daemon
  - starting 116
- vxctl command 269
- vxctl mode command 116
- vxinstall program 118–120
- vxinstall program, running 118
- vxiod I/O daemon
  - starting 117
- vxrootmir command 275