

Veritas Storage Foundation™ Installation Guide

HP-UX

6.0

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: 6.0

Document version: 6.0.4

Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support 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/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

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

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

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

www.symantec.com/business/support/contact_techsupp_static.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/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as 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 support 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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

<https://sort.symantec.com/documents>

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Technical Support	4
Section 1 Installation overview and planning	13
Chapter 1 Introducing Storage Foundation	14
About Veritas products	14
About Storage Foundation	14
About Veritas Volume Replicator	15
About Veritas graphical user interfaces	15
About Veritas Operations Manager	15
Chapter 2 Planning to install Storage Foundation	17
About planning for SF installation	17
About installation and configuration methods	18
Downloading the Storage Foundation software	18
Chapter 3 System requirements	20
Release notes	20
Hardware compatibility list (HCL)	20
Supported operating systems	21
Veritas File System requirements	21
Cluster environment requirements	21
Disk space requirements	22
Discovering product versions and various requirement	
information	22
Database requirements	23
Chapter 4 Licensing Veritas products	24
About Veritas SFHA Solutions product licensing	24
Setting or changing the Veritas SFHA Solutions product level for	
keyless licensing	25
Installing Veritas SFHA Solutions product license keys	27

Section 2	Installation of Storage Foundation	28
Chapter 5	Preparing to install Storage Foundation	29
	Installation preparation overview	29
	About using ssh or remsh with the Veritas installer	30
	Creating the /opt directory	31
	Setting environment variables	31
	Mounting the product disc	31
	Assessing the system for installation readiness	32
	Symantec Operations Readiness Tools	32
	Prechecking your systems using the Veritas installer	33
Chapter 6	Installing Storage Foundation using the script-based installer	34
	About installing Storage Foundation on HP-UX	34
	Summary of Storage Foundation installation tasks	35
	About the Veritas installer	35
	Installing Storage Foundation using the installer	36
Chapter 7	Installing Storage Foundation using the Web-based installer	39
	About the Web-based installer	39
	Before using the Veritas Web-based installer	40
	Starting the Veritas Web-based installer	40
	Obtaining a security exception on Mozilla Firefox	41
	Performing a pre-installation check with the Veritas Web-based installer	42
	Installing SF with the Web-based installer	42
Chapter 8	Installing Storage Foundation using other methods	45
	Installing SF using Ignite-UX	45
	Creating the Software Distributor (SD) bundle for SF or the operating system and SF	45
	Using Ignite-UX to perform a standalone SF installation	46
	Using Ignite-UX to install SF and the HP-UX operating system	47

Section 3	Configuration of Storage Foundation	49
Chapter 9	Configuring Storage Foundation	50
	Configuring Storage Foundation using the installer	50
	Configuring Storage Foundation manually	50
	Configuring Veritas Volume Manager	51
	Configuring Veritas File System	56
	Configuring your system after the installation	57
	Configuring the SFDB repository database after installation	58
Section 4	Upgrading Storage Foundation	59
Chapter 10	Preparing to upgrade	60
	About upgrading	60
	About the different ways that you can upgrade	60
	Supported upgrade paths	61
	Preparing to upgrade	62
	Getting ready for the upgrade	62
	Preparing for an upgrade of Storage Foundation	64
	Creating backups	66
	Determining which 5.x release of Veritas File System and Veritas Volume Manager that you have installed	66
	Tasks for upgrading the Storage Foundation for Databases (SFDB)	67
	Pre-upgrade tasks for migrating the SFDB repository database	68
	Pre-upgrade planning for Veritas Volume Replicator	68
	Upgrading the array support	70
Chapter 11	Upgrading Storage Foundation	72
	Upgrading Storage Foundation using the script-based installer	72
	Upgrading from Storage Foundation or Storage Foundation and High Availability 5.0_11iv3, 5.0.1, 5.0.1 RP1, or 5.0.1 RP2, 5.1SP1, or 5.1SP1RP1	72
	Upgrading from previous versions of Storage Foundation on HP-UX 11i v2	74
	Upgrading from Storage Foundation 3.5 on 11i v1 to SF 6.0 on HP-UX 11i v3	75
	Upgrading from VxVM 5.0 on HP-UX 11i v3 to VxVM 6.0 using integrated VxVM 6.0 package for HP-UX 11i v3	76

	Upgrading Storage Foundation using the script-based installer	77
	Upgrading Storage Foundation using the Veritas Web-based installer	79
	Upgrading the HP-UX operating system	80
	Upgrading Veritas Volume Replicator	81
	Upgrading VVR without disrupting replication	81
Chapter 12	Performing post-upgrade tasks	83
	Optional configuration steps	83
	Post upgrade tasks for migrating the SFDB repository database	84
	Migrating from a 5.0 repository database to 6.0	84
	After upgrading from 5.0.x and before migrating SFDB	87
	Recovering VVR if automatic upgrade fails	87
	Upgrading disk layout versions	88
	About upgrading disk layout versions	88
	Upgrading VxFS disk layout versions	89
	When to use vxfsconvert	90
	When to use vxupgrade	90
	Requirements for upgrading to a later disk layout version	90
	Upgrading VxVM disk group versions	90
	Updating variables	91
	Setting the default disk group	91
	Configuring Powerfail Timeout after upgrade	92
	Converting from QuickLog to Multi-Volume support	92
	Verifying the Storage Foundation upgrade	93
Section 5	Verification of the installation or the upgrade	94
Chapter 13	Verifying the installation	95
	Verifying that the products were installed	95
	Installation log files	95
	Using the installation log file	96
	Using the summary file	96
	Starting and stopping processes for the Veritas products	96
	Checking Veritas Volume Manager processes	96
	Checking Veritas File System installation	97
	Command installation verification	97

Section 6	Uninstallation of Storage Foundation	98
Chapter 14	Uninstalling Storage Foundation	99
	Removing the Replicated Data Set	99
	Uninstalling SF depots using the script-based installer	101
	Uninstalling SF with the Veritas Web-based installer	102
	Removing license files (Optional)	103
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product	103
Section 7	Installation reference	105
Appendix A	Installation scripts	106
	About Veritas Storage Foundation and High Availability Solutions installation scripts	106
	Installation script options	107
Appendix B	Response files	112
	About response files	112
	Installing SF using response files	113
	Configuring SF using response files	113
	Upgrading SF using response files	114
	Uninstalling SF using response files	114
	Syntax in the response file	115
	Response file variables to install, upgrade, or uninstall Storage Foundation	115
	Response file variables to configure Storage Foundation	118
	Sample response file for SF upgrade	120
Appendix C	Tunable files for installation	121
	About setting tunable parameters using the installer or a response file	121
	Setting tunables for an installation, configuration, or upgrade	122
	Setting tunables with no other installer-related operations	123
	Setting tunables with an un-integrated response file	124
	Preparing the tunables file	125
	Setting parameters for the tunables file	125
	Tunables value parameter definitions	126

Appendix D	Configuring the secure shell or the remote shell for communications	133
	About configuring secure shell or remote shell communication modes before installing products	133
	Manually configuring and passwordless ssh	134
	Enabling remsh	138
Appendix E	Storage Foundation components	140
	Storage Foundation installation depots	140
	Veritas Storage Foundation obsolete and reorganized installation depots	143
Appendix F	Troubleshooting installation issues	146
	Restarting the installer after a failed connection	146
	What to do if you see a licensing reminder	146
	Incorrect permissions for root on remote system	147
	Resource temporarily unavailable	148
	Inaccessible system	149
	Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 (2184482)	149
	Workaround	150
Appendix G	Compatibility issues when installing Storage Foundation with other products	151
	Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present	151
	Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present	152
	Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present	152
Index		153

Installation overview and planning

- [Chapter 1. Introducing Storage Foundation](#)
- [Chapter 2. Planning to install Storage Foundation](#)
- [Chapter 3. System requirements](#)
- [Chapter 4. Licensing Veritas products](#)

Introducing Storage Foundation

This chapter includes the following topics:

- [About Veritas products](#)
- [About Veritas graphical user interfaces](#)

About Veritas products

The following products are available for this release.

About Storage Foundation

Veritas Storage Foundation by Symantec includes Veritas File System by Symantec (VxFS) and Veritas Volume Manager by Symantec (VxVM) with various feature levels.

Veritas File System is a high-performance journaling file system that provides easy management and quick-recovery for applications. Veritas File System delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity.

Veritas Volume Manager removes the physical limitations of disk storage. You can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Veritas Volume Manager also provides easy-to-use, online storage management tools to reduce downtime.

You add high availability functionality to Storage Foundation HA by installing Veritas Cluster Server software.

VxFS and VxVM are a part of all Veritas Storage Foundation products. Do not install or update VxFS or VxVM as individual components.

Veritas Storage Foundation has the following products:

- Storage Foundation Standard
- Storage Foundation Standard HA
- Storage Foundation Enterprise
- Storage Foundation Enterprise HA

About Veritas Storage Foundation Basic

Storage Foundation Basic supports all Storage Foundation Standard features, however, there are deployment and technical support limitations.

About Veritas Volume Replicator

Veritas Volume Replicator by Symantec is an optional, separately-licensable feature that is fully integrated with Veritas Volume Manager. This component replicates data to remote locations over any standard IP network to provide continuous data availability.

Volume Replicator is available with Veritas Storage Foundation Standard and Enterprise products.

About Veritas graphical user interfaces

The following are descriptions of Veritas GUIs.

About Veritas Operations Manager

Symantec recommends use of Veritas Operations Manager to manage Storage Foundation and Cluster Server environments.

Veritas Operations Manager provides a centralized management console for Veritas Storage Foundation and High Availability products. You can use Veritas Operations Manager to monitor, visualize, and manage storage resources and generate reports.

You can download Veritas Operations Manager at no charge at <http://go.symantec.com/vom>.

Refer to the Veritas Operations Manager documentation for installation, upgrade, and configuration instructions.

The Veritas Enterprise Administrator (VEA) console is no longer packaged with Storage Foundation products. If you want to continue using VEA, a software version is available for download from

<http://www.symantec.com/operations-manager/support>. Veritas Storage Foundation Management Server is deprecated.

Planning to install Storage Foundation

This chapter includes the following topics:

- [About planning for SF installation](#)
- [About installation and configuration methods](#)
- [Downloading the Storage Foundation software](#)

About planning for SF installation

Before you continue, make sure that you are using the current version of this guide. The latest documentation is available on the Symantec Symantec Operations Readiness Tools (SORT) website.

<https://sort.symantec.com/documents>

Document version: 6.0.4.

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 SF will be installed.

Follow the preinstallation instructions if you are installing Storage Foundation.

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

- Veritas Storage Foundation Basic
- Veritas Storage Foundation (Standard and Enterprise Editions)

Several component products are bundled with each of these SF products.

About installation and configuration methods

You can install and configure SF using Veritas installation programs or using native operating system methods.

Use one of the following methods to install and configure SF:

- The Veritas product installer
 The installer displays a menu that simplifies the selection of installation options.
- The product-specific installation scripts
 The installation scripts provide a command-line interface to install a specific product. The product-specific scripts enable you to specify some additional command-line options. Installing with the installation script is also the same as specifying SF from the installer menu.
- The Web-based Veritas installer
 The installer provides an interface to manage the installation from a remote site using a standard Web browser.
 See [“About the Web-based installer”](#) on page 39.
- Silent installation with response files
 You can use any of the above options to generate a response file. You can then customize the response file for another system. Run the product installation script with the response file to install silently on one or more systems.
 See [“About response files”](#) on page 112.

Downloading the Storage Foundation software

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

For a Trialware download, perform the following. Contact your Veritas representative for more information.

To download the trialware version of the software

- 1 Open the following link in your browser:
<http://www.symantec.com/index.jsp>
- 2 On the bottom of the page, click the **Downloads** link.
- 3 In the Business field, click **Trialware**.
- 4 On the next page near the bottom of the page, click **Business Continuity**.

- 5 Under Cluster Server, click **Download Now**.
- 6 In the new window, click **Download Now**.
- 7 You can use existing credentials to log in or create new credentials.
- 8 Review the terms and conditions, and click **I agree**.
- 9 Find the product that you want to download and select it. Continue with the installation.

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

Note: Trialware is the full product version. The enabled licensing places the product in a demo or a trial state.

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 106.

To download the software

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

The estimated space for download, gunzip, and tar extract is:

4 GB

See [“Disk space requirements”](#) on page 22.

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

- 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. Make sure that different versions exist in different directories.

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

System requirements

This chapter includes the following topics:

- [Release notes](#)
- [Hardware compatibility list \(HCL\)](#)
- [Supported operating systems](#)
- [Veritas File System requirements](#)
- [Cluster environment requirements](#)
- [Disk space requirements](#)
- [Discovering product versions and various requirement information](#)
- [Database requirements](#)

Release notes

The *Release Notes* for each Veritas product contains last minute news and important details for each product, including updates to system requirements and supported software. Review the Release Notes for the latest information before you start installing the product.

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

<https://sort.symantec.com/documents>

Hardware compatibility list (HCL)

The hardware compatibility list contains information about supported hardware and is updated regularly. Before installing or upgrading Storage Foundation and High

Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

For the latest information on supported hardware, visit the following URL:

<http://www.symantec.com/docs/TECH170013>

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

Supported operating systems

For information on supported operating systems, see the *Storage Foundation Release Notes*.

Veritas File System requirements

Complete the tasks in this section before installing Veritas File System.

Before installing Veritas File System, perform the following tasks:

- Review the *Veritas Storage Foundation Release Notes*.
- Ensure that the `/opt` directory exists and has write permissions for `root`.
- The Veritas File System does not support OmniStorage. Do not install VxFS without first retrieving any files archived using OmniStorage.
- Install all the latest required HP-UX patches.

Cluster environment requirements

If you are configuring a cluster, which is a set of hosts that share a set of disks, set up the cluster environment.

To set up a cluster environment

- 1 If you plan to place the root disk group under VxVM control, decide into which disk group you want to configure 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:

```
# vxdg bootdg
```

- 2 Decide on the layout of shared disk groups. There may be one or more shared disk groups. Determine how many you wish to use.
- 3 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* for more information on DRL.
- 4 Install the license that supports the clustering feature on every node in the cluster.

Disk space requirements

Before installing your products, confirm that your system has enough free disk space.

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

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

If you have downloaded SF, you must use the following command:

```
# ./installsf -precheck
```

Discovering product versions and various requirement information

Symantec provides several methods to check the Veritas product you have installed, plus various requirement information.

You can check the existing product versions using the `installer` command with the `-version` option before or after you install. After you have installed the current

version of the product, you can use the `showversion` script in the `/opt/VRTS/install` directory to find version information.

Information the `version` option or the `showversion` script discovers on systems includes the following:

- The installed version of all released Storage Foundation and High Availability Suite of products
- The required depots or patches (if applicable) that are missing
- The available updates (including patches or hotfixes) from Symantec Operations Readiness Tools (SORT) for the installed products

To run the version checker

- 1 Mount the media.
- 2 Start the installer with the `-version` option.

```
# ./installer -version system1 system2
```

Database requirements

The following TechNote identifies the most current information on supported database and operating system combinations:

<http://www.symantec.com/docs/DOC4039>

Note: SF supports running Oracle, DB2, and Sybase on VxFS and VxVM.

SF does not support running SFDB tools with DB2 and Sybase.

Licensing Veritas products

This chapter includes the following topics:

- [About Veritas SFHA Solutions product licensing](#)
- [Setting or changing the Veritas SFHA Solutions product level for keyless licensing](#)
- [Installing Veritas SFHA Solutions product license keys](#)

About Veritas SFHA Solutions product licensing

You have the option to install Veritas products without a license key. Installation without a license does not eliminate the need to obtain a license. A software license is a legal instrument governing the usage or redistribution of copyright protected software. The administrator and company representatives must ensure that a server or cluster is entitled to the license level for the products installed. Symantec reserves the right to ensure entitlement and compliance through auditing.

www.symantec.com/techsupp/

During the installation, you can choose to either:

- Install a license key for the product and features that you want to install.
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.
- Continue to install without a license key.
The installer prompts for the product modes and options that you want to install, and then sets the required product level.
Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server, such as Veritas Operations Manager (VOM). If you do not comply with the above terms,

continuing to use the Symantec product is a violation of your end user license agreement, and results in warning messages.

For more information about keyless licensing, see the following URL:

<http://go.symantec.com/sfhakeyless>

If you encounter problems while licensing this product, visit the Symantec licensing support website.

If you upgrade to this release from a prior release of the Veritas software, the product installer does not change the license keys that are already installed. The existing license keys may not activate new features in this release.

If you upgrade with the product installer, or if you install or upgrade with a method other than the product installer, you must do one of the following to license the products:

See [About Veritas Storage Foundation and High Availability Solutions 6.0](#)

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.

See “[Setting or changing the Veritas SFHA Solutions product level for keyless licensing](#)” on page 25.

See the `vxkeyless(1m)` manual page.

- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.

See “[Installing Veritas SFHA Solutions product license keys](#)” on page 27.

See the `vxlicinst(1m)` manual page.

You can also use the above options to change the product levels to another level that you are authorized to use. For example, you can add the replication option to the installed product. You must ensure that you have the appropriate license for the product level and options in use.

Note: In order to change from one product group to another, you may need to perform additional steps.

Setting or changing the Veritas SFHA Solutions product level for keyless licensing

The keyless licensing method uses product levels to determine the Veritas products and functionality that are licensed. In order to use keyless licensing, you must set up a Management Server to manage your systems.

For more information and to download the management server, see the following URL:

<http://go.symantec.com/vom>

When you set the product license level for the first time, you enable keyless licensing for that system. If you install with the product installer and select the keyless option, you are prompted to select the product and feature level that you want to license.

After you install, you can change product license levels at any time to reflect the products and functionality that you want to license. When you set a product level, you agree that you have the license for that functionality.

To set or change the product level

- 1 Show your current working directory:

```
# pwd
```

Output resembles:

```
/opt/VRTSvlic/bin
```

- 2 View the current setting for the product level.

```
# ./vxkeyless -v display
```

- 3 View the possible settings for the product level.

```
# ./vxkeyless displayall
```

- 4 Set the desired product level.

```
# ./vxkeyless set prod_levels
```

where *prod_levels* is a comma-separated list of keywords. The keywords are the product levels as shown by the output of step 3.

If you want to remove keyless licensing and enter a key, you must clear the keyless licenses. Use the NONE keyword to clear all keys from the system.

Warning: Clearing the keys disables the Veritas products until you install a new key or set a new product level.

To clear the product license level

- 1 View the current setting for the product license level.

```
# ./vxkeyless [-v] display
```

- 2 If there are keyless licenses installed, remove all keyless licenses:

```
# ./vxkeyless [-q] set NONE
```

For more details on using the `vxkeyless` utility, see the `vxkeyless(1m)` manual page.

Installing Veritas SFHA Solutions product license keys

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

<code>vxlicinst</code>	Installs a license key for a Symantec product
<code>vxlicrep</code>	Displays currently installed licenses
<code>vxlictest</code>	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.

To install a new license

- ◆ Run the following commands. In a cluster environment, run the commands on each node in the cluster:

```
# cd /opt/VRTS/bin
```

```
# ./vxlicinst -k xxxx-xxxx-xxxx-xxxx-xxxx-xxx
```

Installation of Storage Foundation

- [Chapter 5. Preparing to install Storage Foundation](#)
- [Chapter 6. Installing Storage Foundation using the script-based installer](#)
- [Chapter 7. Installing Storage Foundation using the Web-based installer](#)
- [Chapter 8. Installing Storage Foundation using other methods](#)

Preparing to install Storage Foundation

This chapter includes the following topics:

- [Installation preparation overview](#)
- [About using ssh or remsh with the Veritas installer](#)
- [Creating the /opt directory](#)
- [Setting environment variables](#)
- [Mounting the product disc](#)
- [Assessing the system for installation readiness](#)

Installation preparation overview

[Table 5-1](#) provides an overview of an installation using the product installer.

Table 5-1 Installation overview

Installation task	Section
Obtain product licenses.	See “About Veritas SFHA Solutions product licensing” on page 24.
Download the software, or insert the product DVD.	See “Downloading the Storage Foundation software” on page 18. See “Mounting the product disc” on page 31.
Set environment variables.	See “Setting environment variables” on page 31.

Table 5-1 Installation overview (*continued*)

Installation task	Section
Create the <code>/opt</code> directory, if it does not exist.	See “Creating the /opt directory” on page 31.
Configure the secure shell (ssh) on all nodes.	See “About configuring secure shell or remote shell communication modes before installing products” on page 133.
Verify that hardware, software, and operating system requirements are met.	See “Release notes” on page 20.
Check that sufficient disk space is available.	See “Disk space requirements” on page 22.
Use the installer to install the products.	See “About the Veritas installer” on page 35.

About using ssh or remsh with the Veritas installer

The installer uses passwordless secure shell (`ssh`) or remote shell (`remsh`) communications among systems. The installer uses the `ssh` or `remsh` daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. You then provide the installer with the superuser passwords for the systems where you plan to install. Note that for security reasons, the installation program neither stores nor caches these passwords. The `ssh` or `remsh` communication among the systems is removed when the installation process completes, unless the installation abruptly terminates. If installation terminated abruptly, use the installation script's `-comcleanup` option to remove the `ssh` or `remsh` configuration from the systems.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer can configure `ssh` or `remsh` on the target systems. In the following scenarios, you need to set up `ssh` or `remsh` manually:

- When the root broker is outside of the cluster that you plan to configure.
- When you add new nodes to an existing cluster.
- When the nodes are in a subcluster during a phased upgrade.
- When you perform installer sessions using a response file.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 133.

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.

Setting environment variables

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

After installation, SF commands are in `/opt/VRTS/bin`. SF manual pages are stored in `/opt/VRTS/man`.

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

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

```
$ PATH=$PATH:/usr/sbin:/opt/VRTS/bin
$ MANPATH=/usr/share/man:/opt/VRTS/man:$MANPATH
$ export PATH MANPATH
```

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

```
% set path = ( $path /usr/sbin /opt/VRTS/bin )
% setenv MANPATH /usr/share/man:/opt/VRTS/man:$MANPATH
```

Mounting the product disc

Mounting the product disc

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

To mount the product disc

- 1 Log in as superuser on a system where you want to install SF.
The systems must be in the same subnet.
- 2 Insert the product disc in the appropriate drive on your local system.

- 3 Determine the block device file for the DVD drive:

```
# iocan -fnC disk
```

Make a note of the device file as it applies to your system.

- 4 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom
# mount /dev/rdisk/c0t0d0 /dvdrom
```

- 5 Verify that the disc is mounted:

```
# mount
```

Assessing the system for installation readiness

Symantec provides the following tools for assessing your system, to ensure that the system meets the requirements for installing Storage Foundation 6.0.

Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a Web-based application that is designed to support Symantec enterprise products.

See [“Symantec Operations Readiness Tools”](#) on page 32.

Prechecking your systems using the installer

Performs a pre-installation check on the specified systems. The Veritas product installer reports whether the specified systems meet the minimum requirements for installing Storage Foundation 6.0.

See [“Prechecking your systems using the Veritas installer”](#) on page 33.

Symantec Operations Readiness Tools

[Symantec Operations Readiness Tools \(SORT\)](#) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you

manage your datacenter more efficiently and get the most out of your Symantec products.

Among its broad set of features, SORT lets you do the following:

- Generate server-specific reports that describe how to prepare your servers for installation or upgrade of Symantec enterprise products.
- Access a single site with the latest production information, including patches, agents, and documentation.
- Create automatic email notifications for changes in patches, documentation, and array-specific modules.

To access SORT, go to:

<https://sort.symantec.com>

Prechecking your systems using the Veritas installer

The script-based and Web-based installer's precheck option checks for the following:

- Recommended swap space for installation
- Recommended memory sizes on target systems for Veritas programs for best performance
- Required operating system versions
- Command checks

To use the precheck option

- 1 Start the script-based or Web-based installer.
- 2 Select the precheck option:
 - From the Web-based installer, select the **Perform a Pre-Installation Check** from the Task pull-down menu.
 - In the script-based installer, from root on the system where you want to perform the check, start the installer.

```
# ./installer
```

In the Task Menu, press the p key to start the precheck.

- 3 Review the output and make the changes that the installer recommends.

Installing Storage Foundation using the script-based installer

This chapter includes the following topics:

- [About installing Storage Foundation on HP-UX](#)
- [Summary of Storage Foundation installation tasks](#)
- [About the Veritas installer](#)
- [Installing Storage Foundation using the installer](#)

About installing Storage Foundation on HP-UX

This release of Storage Foundation requires a specific release of the HP-UX operating system.

See the *Storage Foundation Release Notes* section on supported operating systems.

If you are not running one of the listed releases of HP-UX, upgrade HP-UX on your system before you install the new Veritas software.

For an initial installation on a new system, you can use one of the installation procedures described in this section. If you have an existing installation of Storage Foundation that you are upgrading, you must perform an upgrade to move to the 6.0 versions of the Veritas products.

Summary of Storage Foundation installation tasks

Installation of Storage Foundation consists of the following tasks:

- Obtain a license key, if required.
- If the operating system is not at the required OS fusion level, upgrade the operating system to the latest release.
The operating system is bundled with Veritas Volume Manager and Veritas File System. If the Veritas Volume Manager or Veritas File System is in use, follow the steps in the upgrade chapter to upgrade the Storage Foundation and the operating system.
- If patches for the operating system are required, install the patches before upgrading the product.
- Mount the disk.
- Install Storage Foundation 6.0.
Start the installer and select 'I' for install, or run the appropriate installation script.
- Reboot the system.

```
# /usr/sbin/shutdown -r now
```
- Configure the Veritas software.
Start the installer and select 'C' for configure, or run the appropriate installation script with the `-configure` option.

About the Veritas installer

The installer enables you to install and 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 Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 106.

At most points during the installation you can type the following characters for different actions:

- Use `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.

- Use `Control+c` to stop and exit the program if an installation procedure hangs. After a short delay, the script exits.
- Use `q` to quit the installer.
- Use `?` to display help information.
- Use the Enter button to accept a default response.

See [“Installation script options”](#) on page 107.

Installing Storage Foundation using the 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 Set up the systems so that the commands execute on remote machines without prompting for passwords or confirmations with remote shell or secure shell communication utilities.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 133.

- 2 Load and mount the software disc. If you downloaded the software, navigate to the top level of the download directory and skip the next step.

See [“Mounting the product disc”](#) on page 31.

- 3 Move to the top-level directory on the disc.
- 4 From this directory, type the following command to start the installation on the local system. Use this command to install on remote systems if secure shell or remote shell communication modes are configured:

```
# ./installer
```

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

- 7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement as
specified in the storage_foundation/EULA/lang/
EULA_SF_Ux_version.pdf file present on the media? [y,n,q,?] y
```

```
Do you agree with the terms of the End User License Agreement
as specified in the storage_foundation_high_availability/EULA/
lang/EULA_SFHA_Ux_version.pdf file present on the media? [y,n,q,?] y
```

- 8 Select from one of the following installation options:

- Minimal depots: installs only the basic functionality for the selected product.
- Recommended depots: installs the full feature set without optional depots.
- All depots: installs all available depots.

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

- 9 You are prompted to enter the system names where you want to install the software. Enter the system name or names and then press Enter.

```
Enter the system names separated by spaces:
[q,?] host1
```

- 10 After the system checks complete, the installer displays a list of the depots to be installed. Press Enter to continue with the installation.
- 11 You need to synchronize the system clocks of your application servers or have them point to an NTP server. After the system check, if the nodes have time difference, the installer prompts:

```
Do you want to synchronize system clock with NTP server(s)?
[y,n,q] (y)
```

- 12 The installer can configure remote shell or secure shell communications for you among systems, however each system needs to have RSH or SSH servers installed. You also need to provide the superuser passwords for the systems. Note that for security reasons, the installation program neither stores nor caches these passwords.
- 13 The installer may prompt for previous Veritas Volume Manager configurations.

- 14 Choose the licensing method. Answer the licensing questions and follow the prompts.

Note: The keyless license option enables you to install without entering a key. However, you still need a valid license to install and use Veritas products. Keyless licensing requires that you manage the systems with a Management Server.

See [“About Veritas SFHA Solutions product licensing”](#) on page 24.

- 15 The installer prompts you to configure SFHA. You can continue with configuration if you answer **y**.

See "Configuration for Storage Foundation High Availability" for more information.

- 16 You are prompted to enter the Standard or Enterprise product mode.

```
1) SF Standard
2) SF Enterprise
b) Back to previous menu
```

```
Select product mode to license: [1-2,b,q,?] (2) 1
```

- 17 At the prompt, specify whether you want to send your installation information to Symantec.

```
Would you like to send the information about this installation to
Symantec to help improve installation in the future? [y,n,q,?] (y) y
```

Check the log file, if needed, to confirm the installation and configuration. Follow the prompts and reboot as necessary.

Installing Storage Foundation using the Web-based installer

This chapter includes the following topics:

- [About the Web-based installer](#)
- [Before using the Veritas Web-based installer](#)
- [Starting the Veritas Web-based installer](#)
- [Obtaining a security exception on Mozilla Firefox](#)
- [Performing a pre-installation check with the Veritas Web-based installer](#)
- [Installing SF with the Web-based installer](#)

About the Web-based installer

Use the Web-based installer interface to install Veritas products. The Web-based installer can perform most of the tasks that the script-based installer performs.

You use the `webinstaller` script to start and stop the Veritas XPortal Server `xprt1wid` process. The `webinstaller` script can also be used to check the status of the XPortal Server.

When the `webinstaller` script starts the `xprt1wid` process, the script displays a URL. Use this URL to access the Web-based installer from a Web browser such as Internet Explorer or FireFox.

The Web installer creates log files whenever the Web installer is operating. While the installation processes are operating, the log files are located in a session-based

directory under the `/var/tmp` directory. After the install process completes, the log files are located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep these files for auditing, debugging, and future use.

The location of the Veritas XPortal Server configuration file is `/var/opt/webinstaller/xprtlwid.conf`.

See [“Before using the Veritas Web-based installer”](#) on page 40.

See [“Starting the Veritas Web-based installer”](#) on page 40.

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 7-1 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Veritas products.	Must be a supported platform for Storage Foundation 6.0.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must use the same operating system as the target systems and must be at one of the supported operating system update levels.
Administrative system	The system where you run the Web browser to perform the installation.	Must have a Web browser. Supported browsers: ■ Internet Explorer 6, 7, and 8 ■ Firefox 3.x and later

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprt1wid`, on the installation server:

```
# ./webinstaller start
```

The webinstaller script displays a URL. Note this URL.

Note: If you do not see the URL, run the command again.

The default listening port is 14172. If you have a firewall that blocks port 14172, use the `-port` option to use a free port instead.

- 2 On the administrative server, start the Web browser.
- 3 Navigate to the URL that the script displayed.
- 4 Certain browsers may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

When prompted, enter `root` and root's password of the installation server.

- 5 Log in as superuser.

Obtaining a security exception on Mozilla Firefox

You may need to get a security exception on Mozilla Firefox.

The following instructions are general. They may change because of the rapid release cycle of Mozilla browsers.

To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **Add Exception** button.
- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.
- 5 Click **Confirm Security Exception** button.
- 6 Enter `root` in User Name field and root password of the web server in the Password field.

Performing a pre-installation check with the Veritas Web-based installer

This section describes performing a pre-installation check with the Veritas Web-based installer.

To perform a pre-installation check

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 40.
- 2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.
- 3 Select the Storage Foundation from the **Product** drop-down list, and click **Next**.
- 4 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Next**.
- 5 The installer performs the precheck and displays the results.
- 6 If the validation completes successfully, click **Next**. The installer prompts you to begin the installation. Click **Yes** to install on the selected system. Click **No** to install later.
- 7 Click **Finish**. The installer prompts you for another task.

Installing SF with the Web-based installer

This section describes installing SF with the Veritas Web-based installer.

To install SF using the Web-based installer

- 1 Perform preliminary steps.
See [“Performing a pre-installation check with the Veritas Web-based installer”](#) on page 42.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 40.
- 3 Select **Install a Product** from the **Task** drop-down list.
- 4 Select **Storage Foundation** from the Product drop-down list, and click **Next**.
- 5 On the License agreement page, read the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 6 Choose minimal, recommended, or all depots. Click **Next**.

- 7 Indicate the systems where you want to install. Separate multiple system names with spaces. Click **Next**.
- 8 If you have not yet configured a communication mode among systems, you have the option to let the installer configure ssh or remsh. If you choose to allow this configuration, select the communication mode and provide the superuser passwords for the systems.
- 9 After the validation completes successfully, click **Next** to install SF on the selected system.
- 10 After the installation completes, you must choose your licensing method.

On the license page, select one of the following tabs:

- Keyless licensing

Note: The keyless license option enables you to install without entering a key. However, in order to ensure compliance you must manage the systems with a management server.

For more information, go to the following website:

<http://go.symantec.com/sfhakeyless>

Complete the following information:

- Choose whether you want to install Standard or Enterprise mode.
- Choose whether you want to enable Veritas Replicator.

Click **Register**.

- Enter license key

If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

- 11 For Storage Foundation, click **Next** to complete the configuration and start the product processes.

Note that you are prompted to configure only if the product is not yet configured.

If you select n, you can exit the installer. You must configure the product before you can use SF.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

- 12 To configure Veritas Storage Foundation, start the Web-based installer and select **Configure a product**. Click the **OK** button. The installer checks for updates. Click the **Next** button.

The installer displays the save location for the task log files, summary file, and response file.

Click **Finish** button. A message displays requesting a reboot. Execute the command to reboot the system.

```
/usr/sbin/shutdown -r now
```

- 13 If prompted, select the checkbox to specify whether you want to send your installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future?

Click **Finish**. The installer prompts you for another task.

Installing Storage Foundation using other methods

This chapter includes the following topics:

- [Installing SF using Ignite-UX](#)

Installing SF using Ignite-UX

You can install SF or the HP-UX operating system and SF using Ignite-UX.

The following procedures describe:

- See [“Creating the Software Distributor \(SD\) bundle for SF or the operating system and SF”](#) on page 45.
- See [“Using Ignite-UX to perform a standalone SF installation”](#) on page 46.
- See [“Using Ignite-UX to install SF and the HP-UX operating system”](#) on page 47.

Creating the Software Distributor (SD) bundle for SF or the operating system and SF

You can use the installer to create SD bundles.

You must run the following commands from an Ignite-UX Server. The `-ignite` option cannot run with other installation options.

Note: When you create the SD bundle for SF, the Veritas product disc must be mounted on the Ignite-UX Server.

To create an SD bundle using the installer

- 1 Log in to a configured and running Ignite-UX Server and mount the Veritas installation disc.
- 2 From the prompt, run the **installer** command with the **-ignite** option.

```
# installer -ignite
```

- 3 Select the product to create its SD bundle.
- 4 The installer prompts you for the directory name to place the bundle.

```
Enter the file directory to create the VCS bundle:
(/var/opt/ignite/depots)
Checking the free space of file system ..... Done
Enter a name for the bundle which holds all the VCS depots:
(VCS60_bundle)
```

- 5 Accept the default bundle name or give the bundle a new name.
- 6 The installer copies the depots of the selected product from the disc to the Ignite-UX Server and creates the bundle. It then generates configuration files for the bundle.
- 7 The bundle is ready for a standalone installation of the specific product. To quit the installer choose the last option, **None of the above**.

Continue to the next step if you plan to create an SD bundle for both the operating system and SF.

- 8 The installer checks the `/var/opt/ignite/data/INDEX` file to determine if the HP-UX operating system configuration files are available on the Ignite-UX Server. If the file is available, the installer prompts you to add the newly created bundle cfg into the HP-UX operating system cfg clause. You need to add it so that you can choose the bundle during the HP-UX operating system installation.

Answer **y** to add the bundle cfg into the HP-UX operating system cfg clause.

Using Ignite-UX to perform a standalone SF installation

You can use Ignite-UX to install SF on a standalone system.

To use Ignite-UX to install SF

- 1 Make sure that the following bundles or packages are not on the target systems; remove them if they exist.
 - Operating system bundles:

- Base-VxTools-50
 - Base-VxVM-50
 - B3929FB
 - Base-VxFS-50
 - Base-VxVM
 - Operating system bundle packages:
 - AVXTOOL
 - AVXVM
 - AONLINEJFS
 - OnlineJFS01
 - AVXFS
- 2 Create the SD bundle. You should be able to install this bundle to HP-UX systems on your network.
- See [“Creating the Software Distributor \(SD\) bundle for SF or the operating system and SF”](#) on page 45.
- 3 On the system where you want to install the Veritas product, run the following command.
- ```
swinstall -x autoreboot=true -s \
ignite_server_ipadd:/var/opt/ignite/depots/
product_bundle product_bundle
```
- Where *ignite\_server\_ipadd* is the IP address of the Ignite-UX Server and where */var/opt/ignite/depots* is the directory path.
- For example:
- ```
# swinstall -x autoreboot=true -s \  
10.198.92.81:/var/opt/ignite/depots/SF60_bundle SF60_bundle
```
- 4 After you install the bundle, reboot the system.
- 5 Configure the product. See the configuration chapter of this guide.

Using Ignite-UX to install SF and the HP-UX operating system

You can use Ignite-UX to install SF and the operating system.

To use Ignite-UX to install SF and the operating system

- 1 Create the SD bundle. You should be able to install this bundle to HP-UX systems on your network.

See [“Creating the Software Distributor \(SD\) bundle for SF or the operating system and SF”](#) on page 45.
- 2 Install the operating system. See the appropriate HP-UX documentation for details.
- 3 If you use the Ignite-UX screen GUI, switch to the **Software** tab on the configuration page of the operating system installation. On the **Software** tab, select and enable the Veritas product bundle that you want to install.
- 4 On the **Software** tab, deselect the following operating system bundles if they are selected:
 - Base-VxTools-50
 - Base-VxVM-50
 - B3929FB
 - Base-VxFS-50
 - Base-VxVM
- 5 After you have installed the bundle, you need to configure the product. See the configuration chapter of this guide.

Configuration of Storage Foundation

- [Chapter 9. Configuring Storage Foundation](#)

Configuring Storage Foundation

This chapter includes the following topics:

- [Configuring Storage Foundation using the installer](#)
- [Configuring Storage Foundation manually](#)
- [Configuring your system after the installation](#)
- [Configuring the SFDB repository database after installation](#)

Configuring Storage Foundation using the installer

You can use the installer to configure Storage Foundation, although it requires minimal configuration. You do need to start it.

To start Storage Foundation

- ◆ Run the installer command with the configure option.

```
# ./installsf -configure
```

Configuring Storage Foundation manually

You can manually configure different products within Storage Foundation.

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 Storage Foundation Administrator's Guide*.

Converting to a VxVM root disk

It is possible to select VxVM as a choice for your root disk when performing a new installation using Ignite-UX. Alternatively, you can use the following procedure to achieve VxVM rootability by cloning your LVM root disk using the `vxcp_lvmroot` command.

To convert to a VxVM root disk

- 1 Select the disk to be used as your new VxVM root disk. It is recommended that this disk is internal to the main computer cabinet. If this is currently an LVM disk, then it must be removed from LVM control as follows:
 - Use the `lvremove` command to remove any LVM volumes that are using the disk.
 - Use the `vgreduce` command to remove the disk from any LVM volume groups to which it belongs.
 - Use the `pvremove` command to erase the LVM disk headers

If the disk to be removed is the last disk in the volume group, use the `vgremove` command to remove the volume group, and then use `pvremove` to erase the LVM disk headers.

If the disk is not currently in use by any volume or volume group, but has been initialized by `pvcreate`, you must still use the `pvremove` command to remove LVM disk headers.

If you want to mirror the root disk across multiple disks, make sure that all the disks are free from LVM control.

- 2 While booted on the newly upgraded LVM root disk, invoke the `vxcp_lvmroot` command to clone the LVM root disk to the disk(s) you have designated to be the new VxVM root disks. In the following example, `c1t0d0` is used for the target VxVM root disk:

```
# /etc/vx/bin/vxcp_lvmroot -v c1t0d0
```

To additionally create a mirror of the root disk on `c2t0d0`:

```
# /etc/vx/bin/vxcp_lvmroot -v -m c2t0d0 c1t0d0
```

Use of the `-v` (verbose) option is highly recommended. The cloning of the root disk is a lengthy operation, and this option gives a time-stamped progress indication as each volume is copied, and other major events.

- 3 Use the `setboot` (1M) command to save the hardware path of the new VxVM root disk in the system NVRAM. The disk hardware paths can be found using this command:

```
# ioscand -kfnC disk
```

- 4 Reboot from the new VxVM root disk. If you created a mirrored root disk, then there is nothing more to do. The LVM root disk safely co-exists with your VxVM root disk, and provides a backup boot target.
- 5 If desired, you can convert the original LVM root disk into a mirror of your VxVM root disk by using the following commands:

```
# /etc/vx/bin/vxdestroy_lvmroot -v c2t0d0  
# /etc/vx/bin/vxrootmir -v c2t0d0
```

Once this operation is complete, the system is running on a completely mirrored VxVM root disk.

- 6 If later required, you can use the `vxres_lvmroot` command to restore the LVM root disk.

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
```

This message indicates that `vxconfigd` is running, but not enabled:

```
mode: disabled
```

To start the `vxconfigd` daemon, enter the following command:

```
# vxconfigd
```

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 no_of_daemons
```

where the number of daemons ranges from 1 to 16. Symantec recommends that at least one `vxiod` daemon should be run for each CPU in the system.

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

Enabling optional cluster support in VxVM

An optional cluster feature enables you to use VxVM 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.

Converting existing VxVM disk groups to shared disk groups

Use this procedure if you are upgrading from VxVM 3.x to VxVM 6.0 (or Storage Foundation 3.x to a Storage Foundation product at the 6.0 level) and you want to convert existing disk groups to shared disk groups.

If you want to convert existing private disk groups to shared disk groups, use the following procedure. Use these steps if you are moving from a single node to a cluster, or if you are already in a cluster and have existing private disk groups.

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 all of the nodes on which you are converting the disk groups.

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
```

Make sure that CVM is started. To check the master node:

```
# vxdctl -c mode
```

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 disk_group_name
```

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.

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 -s list
```

This displays the existing shared disk groups.

Configuring Veritas File System

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

```
/etc/fstab
```

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

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

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: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 your system after the installation

Use the following procedure to configure your system after installation.

To configure your system after the software upgrade

- 1 Reinstall the mount points in the `/etc/fstab` file that you recorded in the preparation steps.
- 2 Reboot the upgraded systems.
- 3 Restart all the volumes by entering the following command for each disk group:

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

Optional configuration steps

Perform the following optional configuration steps:

- If you want to use Storage Foundation 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.
 - Stop the cluster, restore the VCS configuration files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.
 - To create root volumes that are under VxVM control after installation, use the `vxcp_lvmroot` command.
See [“Converting to a VxVM root disk”](#) on page 51.
See the *Veritas Volume Manager Administrator's Guide*.
 - To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
See [“Upgrading VxFS disk layout versions”](#) on page 89.
See [“Upgrading VxVM disk group versions”](#) on page 90.
- 4 After you complete the installation procedure, proceed to initializing (where required), setting up, and using Veritas Storage Foundation.

Configuring the SFDB repository database after installation

If you want to use the Storage Foundation for Databases (SFDB) tools, you must set up the SFDB repository after installing and configuring SF and Oracle. For SFDB repository set up procedures:

See *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases*

Upgrading Storage Foundation

- [Chapter 10. Preparing to upgrade](#)
- [Chapter 11. Upgrading Storage Foundation](#)
- [Chapter 12. Performing post-upgrade tasks](#)

Preparing to upgrade

This chapter includes the following topics:

- [About upgrading](#)
- [About the different ways that you can upgrade](#)
- [Supported upgrade paths](#)
- [Preparing to upgrade](#)

About upgrading

There are many types of upgrades available. Before you start to upgrade, review the types of upgrades for the Veritas products.

See [“About the different ways that you can upgrade”](#) on page 60.

Review the supported upgrade paths that are available for the different methods of upgrading.

See [“Supported upgrade paths”](#) on page 61.

After you determine the type of upgrade that you want to perform and its upgrade paths, review the steps to prepare for the upgrade.

About the different ways that you can upgrade

Symantec offers you several different ways to upgrade. You need to decide which upgrade method best suits your environment, your expertise, and the downtime required.

Table 10-1 Review this table to determine how you want to perform the upgrade

Upgrade types and considerations	Methods available for upgrade
Typical upgrades—use a Veritas provided tool or you can perform the upgrade manually. Requires some server downtime.	Script-based—you can use this to upgrade for the supported upgrade paths Web-based—you can use this to upgrade for the supported upgrade paths Manual—you can use this to upgrade from the previous release Response file—you can use this to upgrade from the previous release
Native operating system upgrade—use the upgrade software that comes with the operating system. Note that not all operating systems support native upgrades.	Operating system specific methods Operating system upgrades

Supported upgrade paths

[Table 10-2](#) describes upgrading to 6.0 on HP-UX.

Table 10-2 HP-UX upgrades using the script- or Web-based installer

Veritas software versions	HP-UX 11.11	HP-UX 11.23	HP-UX 11.31
3.5	Upgrade the operating system to 11.23, then upgrade it to 11.31. Use the installer to upgrade the Veritas product to 6.0.SFCFS requires additional manual changes."	N/A	N/A
3.5_11iv2	N/A	Upgrade to 4.1. Upgrade the operating system to 11.31. Use the installer to upgrade to 6.0.	N/A

Table 10-2 HP-UX upgrades using the script- or Web-based installer (*continued*)

Veritas software versions	HP-UX 11.11	HP-UX 11.23	HP-UX 11.31
4.1	N/A	Upgrade the operating system to 11.31. Use the installer to upgrade to 6.0.	N/A
4.1 MP1			
4.1 MP2			
5.0			
5.0 MP1			
5.0 MP2			
5.0 MP2 RPx			
5.0_11iv3	N/A	N/A	Use the installer to upgrade to 6.0.
5.0.1			
5.0.1 RPx			
5.1 SP1			
5.1 SP1 RPx			

Preparing to upgrade

Before you upgrade, you need to prepare the systems and storage. Review the following procedures and perform the appropriate tasks.

Getting ready for the upgrade

Complete the following tasks before you perform the upgrade:

- Review the *Veritas Storage Foundation Release Notes* for any late-breaking information on upgrading your system.
- Review the Symantec Technical Support website for additional information: <http://www.symantec.com/techsupp/>
- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.

- Ensure that you have enough file system space to upgrade. Identify where you want to copy the depots, for example `/packages/Veritas` when the root file system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.
Do not put the files on a file system that is inaccessible prior to running the upgrade script.
You can use a Veritas-supplied disc for the upgrade as long as modifications to the upgrade script are not required.
- For any startup scripts in `/etc/rcS.d`, comment out any application commands or processes that are known to hang if their file systems are not present.
- Make sure that the current operating system supports version 6.0 of the product. If the operating system does not support it, plan for a staged upgrade.
- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Veritas products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/fstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/fstab` and not mounted during the upgrade. Active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure the file systems are clean before upgrading.
- Upgrade arrays (if required).
- If replication using VVR is configured, Symantec recommends that the disk group version is at least 110 prior to upgrading.

```
# vxdbg list diskgroup
```

- If replication using VVR is configured, make sure the size of the SRL volume is greater than 110 MB.
Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.
- If replication using VVR is configured, 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.

- If VCS is used to manage VVR replication, follow the preparation steps to upgrade VVR and VCS agents.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. Removing the disk in this manner offers a fallback point.

Preparing for an upgrade of Storage Foundation

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/fstab`. You should also run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.

To prepare for the Veritas software upgrade

- 1 Log in as superuser.
- 2 Perform any necessary preinstallation checks and configuration.
See [“About planning for SF installation”](#) on page 17.
- 3 Use the `vxlicrep` command to make a record of the currently installed Veritas licenses. Print the output or save it on a different system.
- 4 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes.
- 5 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -F vxfs
```

- 6 Unmount all Storage Checkpoints and non-system VxFS file systems:

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

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

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | 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.

- 8 (Optional) If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem
# mount -F vxfs filesystem mountpoint
# umount mountpoint
```

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

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

```
file system device busy
```

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.
```

- 9 (Optional) 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.
- 10 (Optional) Repeat step 6 to verify that the unclean file system is now clean.
- 11 Stop all the volumes by entering the following command for each disk group:

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

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

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

- 12 Comment out the non-system local VxFS mount points from the `/etc/fstab`. Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to uncomment these entries in the `/etc/fstab` file on the upgraded system.
- 13 If Veritas Volume Replicator (VVR) is configured, do the following steps in the order shown:

- Verify that all of the Primary RLINKs are up to date:

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

- Detach the RLINKs.

- Disassociate the SRL.

Creating backups

Save relevant system information before the upgrade.

To create backups

- 1 Log in as superuser.
- 2 Before the upgrade, ensure that you have made backups of all data that you want to preserve.
- 3 Copy the `fstab` file to `fstab.orig`:

```
# cp /etc/fstab /etc/fstab.orig
```
- 4 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 5 If you are installing the high availability version of the Veritas Storage Foundation 6.0 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 which 5.x release of Veritas File System and Veritas Volume Manager that you have installed

If you are upgrading to this release and have a previously-installed release of Veritas File System (VxFS) and Veritas Volume Manager (VxVM), you must determine which release you have installed. Determining which 5.x release that you have installed can be difficult due to the binary path names being the same for both releases. Use the following procedures to determine which release you have installed.

See [“Discovering product versions and various requirement information”](#) on page 22.

To determine which release of VxFS that you have installed

- ◆ To determine which release of VxFS that you have installed, enter the following command:

```
# swlist -l product VRTSvxfs
```

If you have the 5.0 release installed, the command output includes the following information:

```
VRTSvxfs          5.0.31.0          VERITAS File System
```

If you have the 5.0.1 release installed, the command output includes the following information:

```
VRTSvxfs          5.0.31.5          VERITAS File System
```

If you have the 5.1 SP1 release installed, the command output includes the following information:

```
VRTSvxfs          5.1.100.000       VERITAS File System
```

To determine which release of VxVM that you have installed

- ◆ To determine which release of VxVM that you have installed, enter the following command:

```
# swlist -l product VRTSvxvm
```

If you have the 5.0 release installed, the command output includes the following information:

```
VRTSvxvm          5.0.31.1          Veritas Volume Manager by Symantec
```

If you have the 5.0.1 release installed, the command output includes the following information:

```
VRTSvxvm          5.0.31.5          Veritas Volume Manager by Symantec
```

If you have the 5.1 SP1 release installed, the command output includes the following information:

```
VRTSvxvm          5.1.100.000       Veritas Volume Manager by Symantec
```

Tasks for upgrading the Storage Foundation for Databases (SFDB)

Tasks for upgrading SFDB tools to version 6.0:

- Preparing to migrate the repository database before upgrading from 5.0x or earlier to 6.0
See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 68.
- Migrating the repository database after upgrading from 5.0.x or earlier to 6.0
See [“Post upgrade tasks for migrating the SFDB repository database”](#) on page 84.

Caution: If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 6.0: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 6.0.

Pre-upgrade tasks for migrating the SFDB repository database

If you plan to continue using Database Storage Checkpoints or SmartTier for Oracle policies you created with a 5.0x or earlier version of Storage Foundation for Oracle, you must prepare to migrate the SFDB repository database to 6.0 before upgrading to Storage Foundation or Storage Foundation for Oracle RAC 6.0.

If you are upgrading from 5.1 or 5.1SP1 to 6.0, no upgrade steps are required for the SFDB tools. The migration process is automatic.

Note: The Sfua_Base repository resource group will be removed from the main.cf file. It is not required as a separate service group for SF 6.0.

Perform the following before upgrading SF.

To prepare to migrate the repository database

- ◆ Resynchronize all existing snapshots before upgrading. As Oracle user, enter:

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

Warning: The Database Flashsnap clone database will not be able to be carried over after upgrading. You must create a new Database Flashsnap clone database after upgrading to 6.0.

Pre-upgrade planning for Veritas Volume Replicator

Before installing or upgrading Veritas Volume Replicator (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.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for more information.

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

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 sites. Symantec recommends that the Secondary hosts be upgraded before the Primary host in the Replicated Data Set (RDS).

See the *Storage Foundation Release Notes* for information regarding VVR support for replicating across Storage Foundation versions

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 Replicated Volume Groups (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.

Also, if you specify TCP as the network protocol, the VVR versions on the Primary and Secondary determine whether the checksum is calculated. As shown in [Table 10-3](#), if either the Primary or Secondary are running a version of VVR prior to 6.0, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 6.0, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 10-3 VVR versions and checksum calculations

VVR prior to 6.0 (DG version <= 140)	VVR 6.0 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes
Primary and Secondary		Yes
	Primary and Secondary	No

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.

Planning and upgrading VVR to use IPv6 as connection protocol

Storage Foundation supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

- VVR continues to support replication between IPv4-only nodes with IPv4 as the internet protocol
- VVR supports replication between IPv4-only nodes and IPv4/IPv6 dual-stack nodes with IPv4 as the internet protocol
- VVR supports replication between IPv6-only nodes and IPv4/IPv6 dual-stack nodes with IPv6 as the internet protocol
- VVR supports replication between IPv6 only nodes
- VVR supports replication to one or more IPv6 only nodes and one or more IPv4 only nodes from a IPv4/IPv6 dual-stack node
- VVR supports replication of a shared disk group only when all the nodes in the cluster that share the disk group are at IPv4 or IPv6

Upgrading the array support

The Storage Foundation 6.0 release includes all array support in a single depot, VRTSaslapm. The array support depot includes the array support previously included in the VRTSvxvm depot. The array support depot also includes support previously packaged as external array support libraries (ASLs) and array policy modules (APMs).

See the 6.0 Hardware Compatibility List for information about supported arrays.

See [“Hardware compatibility list \(HCL\)”](#) on page 20.

When you upgrade Storage Foundation products with the product installer, the installer automatically upgrades the array support. If you upgrade Storage

Foundation products with manual steps, you should remove any external ASLs or APMs that were installed previously on your system. Installing the VRTSvxvm depot exits with an error if external ASLs or APMs are detected.

After you have installed Storage Foundation 6.0, Symantec provides support for new disk arrays through updates to the `VRTSaslapm` depot.

For more information about array support, see the *Veritas Storage Foundation Administrator's Guide*.

Upgrading Storage Foundation

This chapter includes the following topics:

- [Upgrading Storage Foundation using the script-based installer](#)
- [Upgrading Storage Foundation using the Veritas Web-based installer](#)
- [Upgrading the HP-UX operating system](#)
- [Upgrading Veritas Volume Replicator](#)

Upgrading Storage Foundation using the script-based installer

You can use the script-based installer to upgrade Storage Foundation.

Upgrading from Storage Foundation or Storage Foundation and High Availability 5.0_11iv3, 5.0.1, 5.0.1 RP1, or 5.0.1 RP2, 5.1SP1, or 5.1SP1RP1

This procedure describes upgrading from Storage Foundation or Storage Foundation and High Availability 5.0_11iv3, 5.0.1, or 5.0.1 RP1 on HP-UX 11i v3 to 6.0 on HP-UX 11i v3.

After successful completion of the upgrade, disk groups that you created in previous versions are accessible.

To upgrade from Storage Foundation or Storage Foundation and High Availability on HP-UX 11i v3

- 1 Perform the necessary pre-upgrade tasks such as resynchronizing existing database snapshots.
- 2 Upgrade the HP-UX operating system to the latest available HP-UX 11i v3 fusion release.

See [“Upgrading the HP-UX operating system”](#) on page 80.
- 3 If patches to HP-UX 11i v3 are required, apply the patches before upgrading the product.
- 4 Install Storage Foundation 6.0 for HP-UX 11i v3 using the installer script.

```
# ./installer
```

- 5 Enter **G** to upgrade and press **Return**.
- 6 You are prompted to enter the system names on which the software is to be installed. Enter the system name or names and then press Return.

Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
- 7 You are prompted to agree with the End User License Agreement. Enter **y** and press Return.

```
Do you agree with the terms of the End User License Agreement as
specified in the storage_foundation/EULA/en/EULA_SF_Ux_6.0.pdf
file present on media? [y,n,q,?] y
```

- 8 The installer lists the packages to install or upgrade. You are prompted to confirm that you are ready to stop SF processes.

```
Do you want to stop SF processes now? [y,n,q,?] (y) y
```

If you select y, the installer stops the product processes and makes some configuration updates before upgrading.

- 9 The installer uninstalls and reinstalls the listed packages.
- 10 Uncomment the entries in the `/etc/fstab` file which were commented as part of the pre-upgrade steps.
- 11 Reboot all the nodes.

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

- 12 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

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

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

Upgrading from previous versions of Storage Foundation on HP-UX 11i v2

Use this procedure to upgrade from Storage Foundation that uses a previous version on HP-UX 11i v2. You can upgrade to Storage Foundation 6.0 from Storage Foundation 4.1, 4.1 MP1, 4.1 MP2, 5.0, 5.0 MP1, or 5.0 MP2 on HP-UX 11i v2.

After successful completion of the upgrade, any disk groups that were created in Storage Foundation 4.1, 4.1 MP1, 4.1 MP2, 5.0, 5.0 MP1, or 5.0 MP2 are accessible by Storage Foundation 6.0.

Note: If you are upgrading from Storage Foundation for Oracle:

See [“Tasks for upgrading the Storage Foundation for Databases \(SFDB\)”](#) on page 67.

To upgrade from Storage Foundation on HP-UX 11i v2 or Storage Foundation for Oracle on HP-UX 11i v2

- 1 Perform the necessary preupgrade tasks such as resynchronizing existing database snapshots.
- 2 If you have any external Array Policy Modules (APMs) installed, uninstall the APMs. The following warning message displays during the OS upgrade and also when you issue an administrative command for HP-UX kernel modules after the upgrade, until SF 5.0 on HP-UX 11i v3 is installed:

```
WARNING: The file '/usr/conf/mod/dmpXXX.1' does not
contain valid kernel code. It will be ignored.
```

This message can be ignored and does not affect the functionality of SF.

- 3 Upgrade the HP-UX operating system to the latest available HP-UX 11i v3 fusion release.

See [“Upgrading the HP-UX operating system”](#) on page 80.

- 4 If patches to HP-UX 11i v3 are required, apply the patches before upgrading the product.

- 5 Install Storage Foundation 6.0 for HP-UX 11i v3 using the installer script.

```
# ./installer
```

- 6 Enter **G** to upgrade and press **Return**.
- 7 You are prompted to enter the system names on which the software is to be installed. Enter the system name or names and then press Return.

Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
- 8 You are prompted to agree with the End User License Agreement. Enter **y** and press Return.

```
Do you agree with the terms of the End User License Agreement as
specified in the storage_foundation/EULA/en/EULA_SF_Ux_6.0.pdf
file present on media? [y,n,q,?] y
```

- 9 The installer lists the packages that will be installed or updated. You are prompted to confirm that you are ready to stop SF processes.

```
Do you want to stop SF processes now? [y,n,q,?] (y) y
```

If you select y, the installer stops the product processes and makes some configuration updates before upgrading.

- 10 The installer uninstalls and reinstalls the listed packages.
- 11 Reboot all the nodes.

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

Upgrading from Storage Foundation 3.5 on 11i v1 to SF 6.0 on HP-UX 11i v3

This procedure describes upgrading Storage Foundation 3.5 on HP-UX 11i v1 to Storage Foundation 6.0 on HP-UX 11i v3. Upgrading from HP-UX 11i v1 requires an intermediate upgrade to HP-UX 11i v2 for Storage Foundation and Storage Foundation Clustered File System.

Veritas Volume Manager 3.5 and Veritas Volume Manager 6.0 both support disk group version 90. Therefore, any disk groups with version 90 are accessible by Storage Foundation 6.0 after the upgrade. However, certain features in Storage Foundation 6.0 may require the latest disk group version. Therefore, we recommend upgrading the disk group.

To upgrade from Storage Foundation 3.5 on HP-UX 11i v1 to SF 6.0 on HP-UX 11i v3

- 1 Stop activity to all Storage Foundation 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.
- 2 Upgrade from HP-UX 11i v1 to HP-UX 11i v2 using practices recommended by HP. HP-UX 11i v2 includes VxVM 4.1 by default. All disk groups created using Storage Foundation 3.5 on HP-UX 11i v1 would be accessible.
- 3 Upgrade from HP-UX 11i v2 to the latest available HP-UX 11i v3 fusion release, using practices recommended by HP. The HP-UX 11i v3 fusion release includes VxVM 5.0 by default.
- 4 If patches to HP-UX 11i v3 are required, apply the patches before upgrading the product.
- 5 Install SF 6.0 for HP-UX 11i v3.
- 6 Reboot the systems.

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

Upgrading from VxVM 5.0 on HP-UX 11i v3 to VxVM 6.0 using integrated VxVM 6.0 package for HP-UX 11i v3

You can upgrade from VxVM 5.0 on HP-UX 11i v3 to VxVM 6.0 on HP-UX 11i v3. Use the integrated VxVM 6.0 package for HP-UX 11i v3 from the Ignite depot.

To upgrade using the integrated VxVM 6.0 package from the Ignite depot

- 1 Use HP recommended steps to integrate VxVM 6.0 package with the latest 11i v3 fusion.
- 2 Stop activity to all Storage Foundation 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.
- 3 Upgrade the HP-UX 11i v3 operating system to the latest fusion using the integrated VxVM 6.0 package. This process installs VxVM 6.0 package with the operating system. All disk groups that were created using VxVM 5.0 11i v3 are accessible.

Upgrading Storage Foundation using the script-based installer

Perform the following procedure to upgrade Storage Foundation. The operating system must be at a supported level for this upgrade.

After successful completion of the upgrade, any disk groups that were created in Storage Foundation High Availability 5.0 or 6.0 are accessible by Storage Foundation High Availability 6.0.

Note: If you are upgrading from Storage Foundation for Oracle:

See [“Tasks for upgrading the Storage Foundation for Databases \(SFDB\)”](#) on page 67.

To upgrade SF

- 1 Offline all the VCS service groups.


```
# hagrps -offline servicegroup -sys node_name
```
- 2 Stop VCS if it is already running. On any node, run the following command:


```
# /opt/VRTS/bin/hastop -all
```
- 3 If fencing is configured with VCS, you must disable fencing before proceeding to upgrade.

 To disable fencing, perform the following steps:
 - If the cluster-wide attribute “UseFence” is set to SCSI3, then reset the value to NONE in the `/etc/VRTSvcs/conf/config/main.cf` file
 - On each node, edit the `/etc/vxsfenmode` file to configure vxsfen in disabled mode.

```
# cat /etc/vxfenmode
vxfen_mode=disabled
```

- Stop I/O fencing on each node:

```
# /sbin/init.d/vxfen stop

# /sbin/vxfenconfig -U
```

- 4 Upgrade the HP-UX operating system to the latest available HP-UX 11i v3 Fusion release.
- 5 If patches to HP-UX 11i v3 are required, apply the patches before upgrading the product.
- 6 Install SF 6.0 for HP-UX 11i v3 using the installer script.

```
# ./installer
```

- 7 Enter **G** to upgrade and press the **Return** key.
- 8 Enter the names of the systems that you want to upgrade and press the **Return** key.

Various messages and prompts appear. Answer the prompts appropriately.

- 9 Review the End User License Agreement, and enter **y** if you agree with it. Press the **Return** key.

```
Do you agree with the terms of the End User License Agreement
as specified in the storage_foundation_high_availability/
EULA/lang/EULA_SFHA_Ux_6.0.pdf file present on media?
[y,n,q,?] y
```

- 10 The installer lists the packages that it will install or update. Confirm that you are ready to stop SF processes.

```
Do you want to stop SF processes now? [y,n,q,?] (y)y
```

If you select **y**, the installer stops the product processes and makes some configuration updates.

- 11 The installer uninstalls and reinstalls the listed packages.
- 12 Enable I/O fencing if required. Follow the below steps to enable the fencing.
 - Execute the following steps on all the nodes:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
```

- Set the clusterwide attribute "UseFence" to use SCSI3. Add the following line to the `/etc/VRTSvcs/conf/config/main.cf` file:

```
UseFence=SCSI3
```

- Verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file by running the following commands:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

13 Reboot all the nodes.

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

14 Start Storage Foundation 6.0 for HP-UX 11i v3 using the following command:

```
# /opt/VRTS/install/installsf -start
```

Note the location of the log files, summary file, and response file.

15 Disk groups that were created using VxVM 5.0 can be imported after upgrading to VxVM 5.1.100. However, we recommend upgrading the VxVM disk groups to the latest version.

See ["Upgrading VxVM disk group versions"](#) on page 90.

Upgrading Storage Foundation using the Veritas Web-based installer

This section describes upgrading SF with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems.

To upgrade SF

- 1 Perform the required steps to save any data that you wish to preserve. For example, make configuration file backups.
- 2 Start the Web-based installer.

See ["Starting the Veritas Web-based installer"](#) on page 40.

- 3 On the Select a task and a product page, select **Upgrade a Product** from the Task drop-down menu.

The installer detects the product that is installed on the specified system. Click **Next**.
- 4 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.
- 5 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.
- 6 Click **Next** to complete the upgrade.

After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.
- 7 When prompted, reboot the systems.


```
# /usr/sbin/shutdown -r now
```
- 8 After the upgrade, if the product is not configured, the Web-based installer asks: "Do you want to configure this product?" If the product is already configured, it will not ask any questions.
- 9 Click **Finish**. The installer prompts you for another task.

Upgrading the HP-UX operating system

If you are on an unsupported version of the operating system, you need to upgrade it to HP-UX B.11.31.1103, HP-UX 11i Version 3 March 2011 or later.

If you are upgrading the operating system from HP-UX 11i v2, make sure that you choose the following depots along with the HP-UX 11i v3 March 2011 OEUR release depots:

- `Base-VxFS-version`
where *version* is the base VxFS version bundled with the operating system.
- `Base-VxTools-version`
where *version* is the base VxTools version bundled with the operating system.
- `Base-VxVM-version`
where *version* is the base VxVM version bundled with the operating system.

To upgrade the operating system from HP-UX 11i v2, run the `update-ux` command specifying the Veritas depots along with the HP-UX operating system depots:

```
# swinstall -s os_path Update-UX
# update-ux -s os_path HPUX11i-DC-OE \
Base-VxFS-version Base-VxTools-version \
Base-VxVM-version
```

where *os_path* is the full path of the directory containing the operating system depots.

where *version* is the the base version of Veritas depots bundled with the operating system.

To upgrade the operating system from HP-UX 11i v3, run the `update-ux` command as follows:

```
# update-ux -s os_path HPUX11i-DC-OE
```

where *os_path* is the full path of the directory containing the operating system depots.

For detailed instructions on upgrading the operating system, see the operating system documentation.

Upgrading Veritas Volume Replicator

If a previous version of Veritas Volume Replicator (VVR) is configured, the product installer upgrades VVR automatically when you upgrade the Storage Foundation products.

See [“Upgrading VVR without disrupting replication”](#) on page 81.

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 69.

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.

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_rvgname
```

- 2 Upgrade from VVR 4.1 or later to VVR 6.0 on the Secondary.
- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname sec_hostname
```

Upgrading VVR on the Primary

After you upgrade the Secondary, use the Veritas product installer to upgrade the Primary.

To upgrade the Primary

- 1 Stop replication to the Primary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvgname
```

- 2 Upgrade from VVR 4.1 or later to VVR 6.0 on the Secondary.
- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname  
sec_hostname
```

See [“Planning an upgrade from the previous VVR version”](#) on page 69.

Performing post-upgrade tasks

This chapter includes the following topics:

- [Optional configuration steps](#)
- [Post upgrade tasks for migrating the SFDB repository database](#)
- [Recovering VVR if automatic upgrade fails](#)
- [Upgrading disk layout versions](#)
- [About upgrading disk layout versions](#)
- [Upgrading VxVM disk group versions](#)
- [Updating variables](#)
- [Setting the default disk group](#)
- [Configuring Powerfail Timeout after upgrade](#)
- [Converting from QuickLog to Multi-Volume support](#)
- [Verifying the Storage Foundation upgrade](#)

Optional configuration steps

After the upgrade is complete, additional tasks may need to be performed.

You can perform the following optional configuration steps:

- If Veritas Volume Replicator (VVR) is configured, do the following steps in the order shown:

- Reattach the RLINKs.
- Associate the SRL.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
See [“Upgrading VxVM disk group versions”](#) on page 90.

Post upgrade tasks for migrating the SFDB repository database

Database Storage Checkpoints that have been created by using the SFDB tools before upgrade are visible using the `vxsfadm` CLI, and you can mount these Database Storage Checkpoints and roll back to them, if required. However, creating clones by using migrated Database Storage Checkpoints is not supported.

If you want to continue using previously created FlashSnap snapplans to take snapshots, you must validate them by using the `-o validate` option of the `vxsfadm` command.

To continue using the Database Storage Checkpoints or SmartTier for Oracle policies you created with a 5.0x or earlier version of Storage Foundation for Oracle, you must perform one of the following procedures after upgrading SF to 6.0:

- Rename startup script after upgrading from 5.0x and before migrating the SFDB repository
See [“After upgrading from 5.0.x and before migrating SFDB”](#) on page 87.
- Migrate from a 5.0x SFDB repository database to 6.0
See [“Migrating from a 5.0 repository database to 6.0”](#) on page 84.

Migrating from a 5.0 repository database to 6.0

To migrate from a 5.0 repository database to 6.0

- 1 When upgrading SFDB tools from the previous release in an HP Service Guard environment, first verify that the `cmviewc1` command can be executed by a non-root user.
- 2 Rename the startup script `NO_S*vxdbms3` to `S*vxdbms3`.
See [“After upgrading from 5.0.x and before migrating SFDB”](#) on page 87.
- 3 As root, dump out the old Sybase ASA repository. If you are using SFHA or SF Oracle RAC, you only need to do this on one node.

```
# /opt/VRTSdbed/migrate/sfua_rept_migrate
```

- 4 On the same node that you ran `sfua_rept_migrate` run the following command as Oracle user. For each Oracle instance, migrate the old repository data to the SQLite repository.

```
$ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

- 5 By default, the repository is created on the file system which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:

- Repository path has to be a directory writable by Oracle user.
- The update commands will not be able to verify accessibility of the repository path and will fail if you have not set up the path correctly.

Create an alternate repository path.

```
$ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME -R \
Alternate_path
```

- 6 If you are using Database Flashsnap for off-host processing, and if you have a repository on the secondary host that you need to migrate: perform the previous steps on the secondary host.

- 7 On the primary host, edit your snapplans to remove the "SNAPSHOT_DG=SNAP_*" parameter and add "SNAPSHOT_DG_PREFIX=SNAP_*. The parameter can be any PREFIX value and not necessarily "SNAP_".

For example:

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=system1_data
SNAPSHOT_DG=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=judge_data
SNAPSHOT_DG_PREFIX=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

- 8 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/vxsfadm -s flashsnap \  
-a oracle -c SNAPPLAN -o validate
```

This completes the migration of the repository for Database Storage Checkpoints and Database Tiered Storage parameters.

To begin using the Storage Foundation for Databases (SFDB) tools:

See *Storage Foundation: Storage and Availability Management for Oracle Databases*

After upgrading from 5.0.x and before migrating SFDB

When upgrading from SF version 5.0 or 5.0.1 to SF 6.0 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by sfua_rept_migrate. Thus when sfua_rept_migrate is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found  
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.  
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

To prevent S*vxdbms3 startup script error

- ◆ Rename the startup script NO_S*vxdbms3 to S*vxdbms3.

When upgrading SFDB tools from the previous release in an HP Service Guard environment, first verify that the `cmviewcl` command can be executed by a non-root user. This permission change must be done before executing SFDB upgrade commands.

Recovering VVR if automatic 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  
# addddcm  
# srlprot  
# attrlink  
# start.rvg
```

After the configuration is restored, the current step can be retried.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 5, 7, 8, and 9. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Disk layout Version 6 has been deprecated and you cannot cluster mount an existing file system that has disk layout Version 6. To upgrade a cluster file system with disk layout Version 6, you must local mount the file system and then upgrade the file system using the `vxupgrade` utility to a later version.

See the `vxupgrade(1M)` manual page.

Support for disk layout Version 4 has been removed. You must upgrade any existing file systems with disk layout Version 4 to disk layout Version 7 or later using the `vxfsconvert` command.

See the `vxfsconvert(1M)` manual page.

Note: Symantec recommends that you upgrade existing file systems to the highest supported disk layout version prior to upgrading to this release.

You can check which disk layout version your file system has by using the following command:

```
# fstyp -v /dev/vx/dsk/dg1/vol1 | grep -i version
```

For more information about disk layout versions, see the *Veritas Storage Foundation Administrator's Guide*.

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 6.0 release.

See the *Veritas Storage Foundation Release Notes 6.0* for information on new features.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 7 or later.

See the `vxfsconvert` or `vxupgrade` man pages.

Warning: Never upgrade the `/` and `/stand` file systems to disk layout Version 7 or later. The HP-UX bootloader does not support disk layout Version 7 or later.

Upgrading VxFS disk layout versions

Veritas File System 6.0 allows Version 5, 7, 8, and 9 for locally mounted file systems and disk layout Versions 7, 8 and 9 for cluster mounted file systems. If you have cluster-mounted file systems with disk layout Versions lower than 7, then after upgrading to VxFS 6.0, upgrade the disk layout. Perform the following additional steps to prepare the file system for being mounted on all nodes of the cluster.

Disk layout Versions 4 and earlier are not supported by VxFS 6.0. All file systems created on VxFS 6.0 use disk layout Version 9 by default.

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

To upgrade VxFS disk layout versions

- 1 Select one of the nodes of the cluster and mount the file system locally on this node. Use the `mount`, but without the `-o cluster` option. For example:

```
# mount -F vxfs /dev/vx/dsk/sharedg/voll /mnt1
```

- 2 To find the current disk layout version on a file system:

```
# fstyp -v <char_device_path> | grep version | \
  awk '{print $2}'
```

- 3 On the node selected in step 1, incrementally upgrade the disk layout of this file system to layout Version 7 or later.

For example, if you had a cluster mounted file system of disk layout Version 4 running with previous version of VxFS, after upgrading to VxFS 6.0, you would need to upgrade the disk layout to Version 7 or later. The incremental upgrade is as follows:

```
# vxupgrade -n 5 /mnt1
# vxupgrade -n 6 /mnt1
# vxupgrade -n 7 /mnt1
```

Disk layout Version 4 is not supported for mount purposes. To mount disk layout Version 4, use the offline utility `vxfsconvert` to upgrade to disk layout Version 9.

You can use disk layout Version 5 or later as a local mount and upgrade using the online utility `vxupgrade`.

- 4 On the node selected in step 1, after the disk layout has been successfully upgraded, unmount the file system:

```
# umount /mnt1
```

- 5 This file system can be mounted on all nodes of the cluster.

When to use vxfsconvert

You can use the `vxfsconvert` command to convert an unmounted HFS file system to a Veritas file system with disk layout Version 9.

```
# vxfsconvert /device_name
```

See the `vxfsconvert(1M)` and `fsadm_vxfs(1M)` manual pages.

When to use vxupgrade

You can use the `vxupgrade` command to upgrade older VxFS disk layouts to disk layout Version 7 or later while the file system remains mounted.

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

See the `vxupgrade(1M)` and `fsadm_vxfs(1M)` manual pages.

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

Requirements for upgrading to a later disk layout version

Converting a previous disk layout version to a later disk layout version 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.

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. VxVM 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 or use the features, upgrade the existing disk groups.

For 6.0, the Veritas Volume Manager disk group version is different than in previous VxVM releases. Symantec recommends that you upgrade the disk group version if you upgraded from a previous VxVM release.

After upgrading to SF 6.0, you must upgrade any existing disk groups that are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

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 Storage Foundation Administrator's Guide*.

Updating variables

In `/etc/profile`, update the `PATH` and `MANPATH` variables as needed.

`MANPATH` could include `/opt/VRTS/man` and `PATH` `/opt/VRTS/bin`.

Setting the default disk group

In releases prior to Volume Manager 4.0, the default disk group was `rootdg` (the root disk group). For Volume Manager to function, the `rootdg` disk group had to exist and it had to contain at least one disk.

This requirement no longer exists; however, you may find it convenient to create a system-wide default disk group. The main benefit of creating a default disk group is that VxVM commands default to the default disk group. You do not need to use the `-g` option.

You can set the name of the default disk group after installation by running the following command on a system:

```
# vxdctl defaultdg diskgroup
```

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

Configuring Powerfail Timeout after upgrade

When you install SF, SF configures Powerfail Timeout (PFTO) using tunable parameters. Starting with SF 5.0.1, the Powerfail Timeout (PFTO) has the following default values:

- disabled for devices using the HP-UX native multi-pathing
- enabled for devices using DMP

After installation, you can override the defaults, if required. You can explicitly enable or disable PFTO for native multi-pathing devices and DMP devices.

When you upgrade from SF release 5.0 or earlier, SF does not preserve any user-defined PFTO values. After the upgrade, the PFTO default values apply to all devices. If you want to use the device settings from the previous release, you must set the desired value explicitly. For example, in an SF 5.0 installation, you have set the PFTO state to enabled for a native multi-pathing device. After you upgrade from SF 5.0 to SF 6.0, the native device is set to the default value, which is disabled. In order to use PFTO, you must explicitly enable the PFTO on that device.

When you upgrade from SF release 5.0.1 or higher to 6.0, SF preserves the PFTO state for the devices. After the upgrade, the PFTO values are set to the same values that the device had before the upgrade. For example, in an SF 5.0.1 installation, you have set the PFTO state to enabled for a native multi-pathing device. After you upgrade from SF 5.0.1 to SF 6.0, the native device will have the PFTO state as enabled.

For more information about controlling Powerfail Timeout, see the *Veritas Volume Manager Administrator's Guide*.

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 and later disk layouts do 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 the Version 6 or later 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.

```
# qlogdetach -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 the Version 7 or later disk layout.

See [“About upgrading disk layout versions”](#) on page 88.

For example:

```
# vxupgrade -n 9 /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 Storage Foundation upgrade

Refer to the section about verifying the installation to verify the upgrade.

See [“Verifying that the products were installed”](#) on page 95.

Verification of the installation or the upgrade

- [Chapter 13. Verifying the installation](#)

Verifying the installation

This chapter includes the following topics:

- [Verifying that the products were installed](#)
- [Installation log files](#)
- [Starting and stopping processes for the Veritas products](#)
- [Checking Veritas Volume Manager processes](#)
- [Checking Veritas File System installation](#)

Verifying that the products were installed

Verify that the SF products are installed.

Use the `swlist` command to check which depots have been installed:

```
# swlist -l product | grep VRTS
```

You can verify the version of the installed product. Use the following command:

```
# /opt/VRTS/install/installsf -version
```

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 summary file

The summary file contains the results of the installation by the installer or product installation scripts. The summary includes the list of the depots, and the status (success or failure) of each depot. 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.

Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

To stop the processes

- ◆ Use the `-stop` option to stop the product installation script.

For example, to stop the product's processes, enter the following command:

```
# ./installer -stop
```

To start the processes

- ◆ Use the `-start` option to start the product installation script.

For example, to start the product's processes, enter the following command:

```
# ./installer -start
```

Checking Veritas 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 -ef | grep vx
```

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

If you installed Storage Foundation and High Availability, the `vxodmd` and `vxdbd_11.31` processes are also displayed.

Checking Veritas File System installation

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

Command installation verification

The Veritas File System commands are installed in the `/opt/VRTS/bin` directory. To verify, determine that the subdirectory is present:

```
# ls /opt/VRTS/bin
```

Make sure you have adjusted your environment variables accordingly.

Uninstallation of Storage Foundation

- [Chapter 14. Uninstalling Storage Foundation](#)

Uninstalling Storage Foundation

This chapter includes the following topics:

- [Removing the Replicated Data Set](#)
- [Uninstalling SF depots using the script-based installer](#)
- [Uninstalling SF with the Veritas Web-based installer](#)
- [Removing license files \(Optional\)](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product](#)

Removing the Replicated Data Set

If you use VVR, you need to perform the following steps. This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

Note: If you are upgrading Veritas Volume Replicator, do not remove the Replicated Data Set.

To remove the Replicated Data Set

- 1 Verify that all RLINKs are up-to-date:

```
# vxlink -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.

Uninstalling SF depots using the script-based installer

Use the following procedure to remove SF products.

Not all depots may be installed on your system depending on the choices that you made when you installed the software.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in SF 6.0 with a previous version of SF.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 133.

To shut down and remove the installed SF depots

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/fstab`. 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 depot (`VRTSvxxvm`) is installed, read and follow the uninstallation procedures for VxVM.
- 4 Make sure you have performed all of the prerequisite steps.
- 5 Move to the `/opt/VRTS/install` directory and run the uninstall script.

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

For Veritas Storage Foundation

```
# ./uninstallsf
```

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

```
Enter the system names separated by spaces: [q?] host1 host2
```

- 7 The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the depots are uninstalled.

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

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

Uninstalling SF with the Veritas Web-based installer

This section describes how to uninstall using the Veritas Web-based installer.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout Version in SF 6.0 with a previous version of SF.

To uninstall SF

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 40.
- 3 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 4 Select **Storage Foundation** from the Product drop-down list, and click **Next**.
- 5 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Next**.
- 6 After the validation completes successfully, click **Next** to uninstall SF on the selected system.
- 7 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 8 After the installer stops the processes, the installer removes the products from the specified system.

Click **Next**.

- 9 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.
- 10 Click **Finish**.

The Web-based installer prompts you to reboot the system.

You see a prompt recommending that you reboot the system, and then return to the Web page to complete additional tasks.

Removing license files (Optional)

Optionally, you can remove the license files.

To remove the VERITAS license files

- 1 To see what license key files you have installed on a system, enter:

```
# /sbin/vxlicrep
```

The output lists the license keys and information about their respective products.

- 2 Go to the directory containing the license key files and list them:

```
# cd /etc/vx/licenses/lic  
# ls -a
```

- 3 Using the output from step 1, identify and delete unwanted key files listed in step 2. Unwanted keys may be deleted by removing the license key file.

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file disables the SFDB tools.

To remove the SFDB repository

- 1 Identify the SFDB repositories created on the host.

```
# cat /var/vx/vxdba/rep_loc

{
  "sfae_rept_version" : 1,
  "oracle" : {
    "SFAEDB" : {
      "location" : "/data/sfaedb/.sfae",
      "old_location" : "",
      "alias" : [
        "sfaedb"
      ]
    }
  }
}
```

- 2 Remove the directory identified by the `location` key.

```
# rm -rf /data/sfaedb/.sfae
```

- 3 Remove the repository location file.

```
# rm -rf /var/vx/vxdba/rep_loc
```

This completes the removal of the SFDB repository.

Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Response files](#)
- [Appendix C. Tunable files for installation](#)
- [Appendix D. Configuring the secure shell or the remote shell for communications](#)
- [Appendix E. Storage Foundation components](#)
- [Appendix F. Troubleshooting installation issues](#)
- [Appendix G. Compatability issues when installing Storage Foundation with other products](#)

Installation scripts

This appendix includes the following topics:

- [About Veritas Storage Foundation and High Availability Solutions installation scripts](#)
- [Installation script options](#)

About Veritas Storage Foundation and High Availability Solutions installation scripts

Veritas Storage Foundation and High Availability Solutions products 6.0 provides several installation scripts. You can find these scripts at the root of the product media in the scripts directory.

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

See [“Downloading the Storage Foundation software”](#) on page 18.

The following product installation scripts are available:

Table A-1 Product installation scripts

Product installation script	Veritas product name
<code>installvcs</code>	Veritas Cluster Server (VCS)
<code>installsf</code>	Veritas Storage Foundation (SF)
<code>installsfha</code>	Veritas Storage Foundation and High Availability (SFHA)

Table A-1 Product installation scripts (*continued*)

Product installation script	Veritas product name
<code>installsfcfsha</code>	Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)
<code>installsffrac</code>	Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
<code>installvm</code>	Veritas Volume Manager
<code>installfs</code>	Veritas File System
<code>installdmp</code>	Veritas Dynamic Multi-Pathing

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-2](#) shows command line options for the 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 Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 106.

Table A-2 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>-allpkgs</code>	Displays all depots and patches required for the specified product. The depots and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-comcleanup</code>	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
<code>-configure</code>	Configures the product after installation.
<code>-ignite</code>	The <code>-ignite</code> option allows you to create SD bundles for a product. When you create the SD bundle for the product, the Veritas product disc must be mounted on the Ignite-UX Server.
<code>-hostfile full_path_to_file</code>	Specifies the location of a file that contains a list of hostnames on which to install.
<code>-install</code>	The <code>-install</code> option is used to install products on systems.
<code>-installallpkgs</code>	Specifies that all depots are installed.
<code>-installminpkgs</code>	Specifies that the minimum depot set is installed.
<code>-installrecpkgs</code>	Specifies that the required depot set is installed.
<code>-keyfile ssh_key_file</code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i ssh_key_file</code> to every SSH invocation.
<code>-license</code>	Registers or updates product licenses on the specified systems.
<code>-logpath log_path</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>-makeresponsefile</code>	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-minpkgs</code>	Displays the minimal depots and patches required for the specified product. The depots and patches are listed in correct installation order. Optional depots are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.
<code>-nolic</code>	Allows installation of product depots without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
<code>-pkginfo</code>	Displays a list of depots and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the <code>-pkginfo</code> option with the <code>installvcs</code> script to display VCS depots.
<code>-pkgpath <i>package_path</i></code>	Designates the path of a directory that contains all depots to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
<code>-pkgset</code>	Discovers and displays the depot group (minimum, recommended, all) and depots that are installed on the specified systems.
<code>-pkgtable</code>	Displays product's depots in correct installation order by group.
<code>-postcheck</code>	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
<code>-precheck</code>	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-recpkgs</code>	Displays the recommended depots and patches required for the specified product. The depots and patches are listed in correct installation order. Optional depots are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.
<code>-redirect</code>	Displays progress details without showing the progress bar.
<code>-requirements</code>	The <code>-requirements</code> option displays required OS version, required depots and patches, file system space, and other system requirements in order to install the product.
<code>-responsefile <i>response_file</i></code>	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.
<code>-rootpath <i>root_path</i></code>	Specifies an alternative root directory on which to install depots. On HP-UX operating systems, <code>-rootpath</code> passes <code>-I path</code> to <code>swinstall</code> .
<code>-rsh</code>	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See “About configuring secure shell or remote shell communication modes before installing products” on page 133.
<code>-serial</code>	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.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
-set tunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the <code>-tunablesfile</code> option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-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 depots are copied on remote systems before installation.
-uninstall	The <code>-uninstall</code> option is used to uninstall products from systems.
-tunablesfile	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing depots and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing depots and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.

Response files

This appendix includes the following topics:

- [About response files](#)
- [Installing SF using response files](#)
- [Configuring SF using response files](#)
- [Upgrading SF using response files](#)
- [Uninstalling SF using response files](#)
- [Syntax in the response file](#)
- [Response file variables to install, upgrade, or uninstall Storage Foundation](#)
- [Response file variables to configure Storage Foundation](#)
- [Sample response file for SF upgrade](#)

About response files

The installer or product installation script generates a response file during any installation, configuration, upgrade, or uninstall procedure. The response file contains the configuration information that you entered during the procedure. When the procedure completes, the installation script displays the location of the response files.

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.

You can generate a response file using the `-makeresponsefile` option.

See [“Installation script options”](#) on page 107.

Installing SF using response files

Typically, you can use the response file that the installer generates after you perform SF installation on a system to install SF on other systems. You can also create a response file using the `-makeresponsefile` option of the installer.

To install SF using response files

- 1 Make sure the systems where you want to install SF meet the installation requirements.
- 2 Make sure the preinstallation tasks are completed.
- 3 Copy the response file to the system where you want to install SF.
- 4 Edit the values of the response file variables as necessary.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installation from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file  
  
# ./installsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Configuring SF using response files

Typically, you can use the response file that the installer generates after you perform SF configuration on one system to configure SF on other systems. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure SF using response files

- 1 Make sure the SF depots are installed on the systems where you want to configure SF.
- 2 Copy the response file to the system where you want to configure SF.

- 3 Edit the values of the response file variables as necessary.

To configure optional features, you must define appropriate values for all the response file variables that are related to the optional feature.

See [“Response file variables to configure Storage Foundation”](#) on page 118.

- 4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Upgrading SF using response files

Typically, you can use the response file that the installer generates after you perform SF upgrade on one system to upgrade SF on other systems. You can also create a response file using the `makeresponsefile` option of the installer.

To perform automated SF upgrade

- 1 Make sure the systems where you want to upgrade SF meet the upgrade requirements.
- 2 Make sure the pre-upgrade tasks are completed.
- 3 Copy the response file to one of the systems where you want to upgrade SF.
- 4 Edit the values of the response file variables as necessary.
- 5 Mount the product disc and navigate to the folder that contains the installation program.
- 6 Start the upgrade from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file
```

```
# ./installsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Uninstalling SF using response files

Typically, you can use the response file that the installer generates after you perform SF uninstallation on one system to uninstall SF on other systems.

To perform an automated uninstallation

- 1 Make sure that you meet the prerequisites to uninstall SF.
- 2 Copy the response file to one of the cluster systems where you want to uninstall SF.
- 3 Edit the values of the response file variables as necessary.
- 4 Start the uninstallation from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/uninstallsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Syntax in the response file

The syntax of the Perl statements that are included in the response file variables varies. It can depend on whether the variables require scalar or list values.

For example, in the case of a string value:

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value", "value", "value"];
```

Response file variables to install, upgrade, or uninstall Storage Foundation

[Table B-1](#) lists the response file variables that you can define to configure SF.

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SF

Variable	Description
CFG{opt}{install}	Installs SF depots. Configuration can be performed at a later time using the <code>-configure</code> option. List or scalar: scalar Optional or required: optional
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required
\$CFG{opt}{vxkeyless}	Installs the product with keyless license. List or scalar: scalar Optional or required: optional
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
CFG{opt}{pkgpath}	Defines a location, typically an NFS mount, from which all remote systems can install product depots. The location must be accessible from all target systems. List or scalar: scalar Optional or required: optional

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SF (*continued*)

Variable	Description
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the depots that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{rsh}	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. List or scalar: scalar Optional or required: optional
CFG{donotinstall} {depot}	Instructs the installation to not install the optional depots in the list. List or scalar: list Optional or required: optional
CFG{donotremove} {depot}	Instructs the uninstallation to not remove the optional depots in the list. List or scalar: list Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
\$CFG{opt}{prodmode}	List of modes for product List or scalar: list Optional or required: optional
CFG{opt}{upgrade}	Upgrades all depots installed, without configuration. List or scalar: list Optional or required: optional

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SF (*continued*)

Variable	Description
CFG{opt}{uninstall}	Uninstalls SF depots. List or scalar: scalar Optional or required: optional
CFG{mirrordgname}{system}	Splits the target disk group name for a system. List or scalar: scalar Optional or required: optional
CFG{splitmirror}{system}	Indicates the system where you want a split mirror backup disk group created. List or scalar: scalar Optional or required: optional

Response file variables to configure Storage Foundation

[Table B-2](#) lists the response file variables that you can define to configure SF.

Table B-2 Response file variables specific to configuring Storage Foundation

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the depots are already installed. (Required) Set the value to 1 to configure SF.
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{systems}	List	List of systems on which the product is to be configured. (Required)

Table B-2 Response file variables specific to configuring Storage Foundation
(continued)

Variable	List or Scalar	Description
CFG{prod}	Scalar	Defines the product to be configured. The value is VCS60 for VCS. (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>remsh</i> must be used instead of ssh as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtpprec, and smtpsev), the SNMP

trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Sample response file for SF upgrade

The following example shows a response file for upgrading Storage Foundation.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{upgrade}=1;
$CFG{systems}=[ qw(system01) ];
1;
```

Tunable files for installation

This appendix includes the following topics:

- [About setting tunable parameters using the installer or a response file](#)
- [Setting tunables for an installation, configuration, or upgrade](#)
- [Setting tunables with no other installer-related operations](#)
- [Setting tunables with an un-integrated response file](#)
- [Preparing the tunables file](#)
- [Setting parameters for the tunables file](#)
- [Tunables value parameter definitions](#)

About setting tunable parameters using the installer or a response file

You can set non-default product and system tunable parameters using a tunables file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing. The tunables file passes arguments to the installer script to set tunables. With the file, you can set the tunables for the following operations:

- When you install, configure, or upgrade systems.

```
# ./installer -tunablesfile tunables_file_name
```

See [“Setting tunables for an installation, configuration, or upgrade”](#) on page 122.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -settunables [  
system1 system2 ...]
```

See [“Setting tunables with no other installer-related operations”](#) on page 123.

- When you apply the tunables file with an un-integrated response file.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name
```

See [“Setting tunables with an un-integrated response file”](#) on page 124.

See [“About response files”](#) on page 112.

You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 126.

Setting tunables for an installation, configuration, or upgrade

You can use a tunables file for installation procedures to set non-default tunables. You invoke the installation script with the `tunablesfile` option. The tunables file passes arguments to the script to set the selected tunables. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 126.

Note: Certain tunables only take effect after a system reboot.

To set the non-default tunables for an installation, configuration, or upgrade

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 125.
- 2 Make sure the systems where you want to install SF meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 126.

Note: Certain tunables only take effect after a system reboot.

To set tunables with no other installer-related operations

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 125.
- 2 Make sure the systems where you want to install SF meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-setttunables` option.

```
# ./installer -tunablesfile tunables_file_name -setttunables [  
  sys123 sys234 ...]
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 126.

Note: Certain tunables only take effect after a system reboot.

To set tunables with an un-integrated response file

- 1 Make sure the systems where you want to install SF meet the installation requirements.
- 2 Complete any preinstallation tasks.
- 3 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 125.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-settunables` option.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name -settunables
```

Where *response_file_name* is the full path name for the response file and *tunables_file_name* is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Preparing the tunables file

A tunables file is a Perl module and consists of an opening and closing statement, with the tunables defined between. Use the hash symbol at the beginning of the line to comment out the line. The tunables file opens with the line "our %TUN;" and ends with the return true "1;" line. The final return true line only needs to appear once at the end of the file. Define each tunable parameter on its own line.

Format the tunable parameter as follows:

```
$TUN{"tunable_name"}{"system_name"|"*"}=value_of_tunable;
```

For the *system_name*, use the name of the system, its IP address, or a wildcard symbol. The *value_of_tunable* depends on the type of tunable you are setting. End the line with a semicolon.

The following is an example of a tunables file.

```
#  
# Tunable Parameter Values:  
#  
our %TUN;  
  
$TUN{"tunable1"}{"*"}=1024;  
$TUN{"tunable3"}{"sys123"}="SHA256";  
  
1;
```

Setting parameters for the tunables file

Each tunables file defines different tunable parameters. The values that you can use are listed in the description of each parameter. Select the tunables that you want to add to the tunables file and then configure each parameter.

See [“Tunables value parameter definitions”](#) on page 126.

Each line for the parameter value starts with \$TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{"node123"}=16;
```

In this example, you are changing the dmp_daemon_count value from its default of 10 to 16. You can use the wildcard symbol "*" for all systems. For example:

```
$TUN{"dmp_daemon_count"}{"*"}=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Veritas Storage Foundation and High Availability Solutions Tuning Guide* for detailed information on product tunable ranges and recommendations .

[Table C-1](#) describes the supported tunable parameters that can be specified in a tunables file.

Table C-1 Supported tunable parameters

Tunable	Description
dmp_cache_open	(Veritas Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_daemon_count	(Veritas Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_delayq_interval	(Veritas Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_evm_handling	(Veritas Dynamic Multi-Pathing) Whether EVM should be handled or not.
dmp_fast_recovery	(Veritas Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_health_time	(Veritas Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_log_level	(Veritas Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_low_impact_probe	(Veritas Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_lun_retry_timeout	(Veritas Dynamic Multi-Pathing) The retry period for handling transient errors. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_fabric	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Veritas Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_support	(Veritas Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_path_age	(Veritas Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_pathswitch_blks_shift	(Veritas Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_idle_lun	(Veritas Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_probe_threshold	(Veritas Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_cycles	(Veritas Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_interval	(Veritas Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_policy	(Veritas Dynamic Multi-Pathing) The policy used by DMP path restoration thread. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_state	(Veritas Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_retry_count	(Veritas Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_scsi_timeout	(Veritas Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_sfg_threshold	(Veritas Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_stat_interval	(Veritas Dynamic Multi-Pathing) The time interval between gathering DMP statistics. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer sets only the system default value of max_diskq. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer sets only the system default value of read_ahead. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer sets only the system default value of read_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer sets only the system default value of read_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
vol_checkpoint_default	(Veritas File System) Size of VxVM checkpoints (sectors). This tunable requires system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for VERITAS Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for VERITAS Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires system reboot to take effect.
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of VERITAS Volume Replicator (bytes).
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (sectors). This tunable requires system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (sectors). This tunable requires system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Veritas Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by VERITAS Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions that can exist on a non-sequential DRL. This tunable requires system reboot to take effect.
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (sectors). This tunable requires system reboot to take effect.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires system reboot to take effect.
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires system reboot to take effect.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
vx_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. This tunable requires system reboot to take effect.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. This tunable requires system reboot to take effect.
vxfs_mbuf	(Veritas File System) Maximum memory used for VxFS buffer cache. This tunable requires system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer sets only the system default value of write_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer sets only the system default value of write_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

- [About configuring secure shell or remote shell communication modes before installing products](#)
- [Manually configuring and passwordless ssh](#)
- [Enabling remsh](#)

About configuring secure shell or remote shell communication modes before installing products

Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a system. The node from which the installer is run must have permissions to run `remsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install Veritas software.

You can install products to remote systems using either secure shell (`ssh`) or remote shell (`remsh`). Symantec recommends that you use `ssh` as it is more secure than `remsh`.

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.

Note: The script- and Web-based installers support establishing password less communication for you.


Manually configuring and passwordless 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 D-1 illustrates this procedure.

Figure D-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 # 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 `/opt/ssh/etc/sshd_config` file must contain the following two lines:

```
PermitRootLogin          yes
Subsystem                 sftp    /opt/ssh/libexec/sftp-server
```

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

```
system1 # /sbin/init.d/secsh 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 `authorized_keys` 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, enter the following command on system2:

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

- 11 To log out of the `ssh` session, enter 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 to 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), enter 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.

Enabling remsh

Remote shell functionality is enabled automatically after installing HP-UX .

Typically, the only requirement to enable remote installations is to modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. You must modify this file for each user who remotely accesses the system using `remsh`. Each line of the `.rhosts` file must contain a fully qualified domain name or IP address for each remote system that has access to the local system. For example, if the root user must remotely access `system1` from `system2`, add an entry for `system2.companyname.com` to the `.rhosts` file on `system1`.

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

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

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

For more information on configuring the remote shell, see the operating system documentation and the `remsh(1M)` manual page.

Storage Foundation components

This appendix includes the following topics:

- [Storage Foundation installation depots](#)
- [Veritas Storage Foundation obsolete and reorganized installation depots](#)

Storage Foundation installation depots

[Table E-1](#) shows the depot name and contents for each English language depot for Storage Foundation. The table also gives you guidelines for which depots to install based on whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and Veritas Cluster Server (VCS) depots, the combined functionality is called Storage Foundation and High Availability.

Table E-1 Storage Foundation depots

depots	Contents	Configuration
VRTSaslapm	Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries Required for the support and compatibility of various storage arrays.	Minimum
VRTSperl	Perl 5.12.2 for Veritas	Minimum

Table E-1 Storage Foundation depots (*continued*)

depots	Contents	Configuration
VRTSvlic	Veritas License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Minimum
VRTSvxfs	Veritas File System binaries Required for VxFS file system support.	Minimum
VRTSvxvm	Veritas Volume Manager binaries, scripts, and utilities. Required for VxVM volume manager support.	Minimum
VRTSdbed	Veritas Storage Foundation for Databases	Recommended
VRTSob	Veritas Enterprise Administrator	Recommended
VRTSodm	Veritas 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 Oracle 9i and 10g to improve performance and manage system bandwidth.	Recommended

Table E-1 Storage Foundation depots (*continued*)

depots	Contents	Configuration
VRTSsfcp60	Veritas Storage Foundation Common Product Installer The Storage Foundation Common Product installer depot contains the installer libraries and product scripts that perform the following: ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ removing nodes ■ etc. You can use these script to simplify the native operating system installations, configurations, and upgrades.	Minimum
VRTSsfmh	Veritas Storage Foundation Managed Host Discovers configuration information on a Storage Foundation managed host. This information is stored on a central database, which is not part of this release. You must download the database separately at:	Recommended
VRTSspt	Veritas Software Support Tools	Recommended
VRTSfsadv	Minimum Veritas File System Advanced Solutions by Symantec (Solaris SPARC only).	Minimum
VRTSfssdk	Veritas File System Software Developer Kit For VxFS APIs, the depot contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.	All

Veritas Storage Foundation obsolete and reorganized installation depots

Table E-2 lists the depots that are obsolete or reorganized for Storage Foundation.

Table E-2 Veritas Storage Foundation obsolete and reorganized depots

depot	Description
Obsolete and reorganized for 6.0	
VRTSat	Obsolete
Obsolete and reorganized for 5.1	
Infrastructure	
SYMCIma	Obsolete
VRTSaa	Included in VRTSsfmh
VRTSccg	Included in VRTSsfmh
VRTSdbms3	Obsolete
VRTSicsco	Obsolete
VRTSjre	Obsolete
VRTSjre15	Obsolete
VRTSmh	Included in VRTSsfmh
VRTSobc33	Obsolete
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product depots	

Table E-2 Veritas Storage Foundation obsolete and reorganized depots
(continued)

depot	Description
VRTSacclib	Obsolete The following information is for installations, upgrades, and uninstallations using the script- or Web-based installer. ■ For fresh installations VRTSacclib is not installed. ■ For upgrades, VRTSacclib is not uninstalled. ■ For uninstallation, VRTSacclib is not uninstalled.
VRTSalloc	Obsolete
VRTScmccc	Obsolete
VRTScmcs	Obsolete
VRTScscm	Obsolete
VRTScscw	Obsolete
VRTScsocw	Obsolete
VRTScssim	Obsolete
VRTScutil	Obsolete
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSdsa	Obsolete
VRTSfsmn	Included in the product's main depot.
VRTSfsmnd	Included in the product's main depot.
VRTSfspro	Included in VRTSsfmh
VRTSvcldb	Included in VRTSvcsea
VRTSvcsor	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvc
VRTSvdid	Obsolete

Table E-2 Veritas Storage Foundation obsolete and reorganized depots
(continued)

depot	Description
VRTSvmman	Included in the product's main depot.
VRTSvmpro	Included in VRTSsfmh
VRTSvrpro	Included in VRTSob
VRTSvrw	Obsolete
VRTSvxmsa	Obsolete
Documentation	All Documentation depots obsolete

Troubleshooting installation issues

This appendix includes the following topics:

- [Restarting the installer after a failed connection](#)
- [What to do if you see a licensing reminder](#)
- [Incorrect permissions for root on remote system](#)
- [Resource temporarily unavailable](#)
- [Inaccessible system](#)
- [Upgrading Veritas Storage Foundation for Databases \(SFDB\) tools from 5.0x to 6.0 \(2184482\)](#)

Restarting the installer after a failed connection

If an installation is killed because of a failed connection, you can restart the installer to resume the installation. The installer detects the existing installation. The installer prompts you whether you want to resume the installation. If you resume the installation, the installation proceeds from the point where the installation failed.

What to do if you see a licensing reminder

In this release, you can install without a license key. In order to comply with the End User License Agreement, you must either install a license key or make the host managed by a Management Server. If you do not comply with these terms within 60 days, the following warning messages result:

```
WARNING V-365-1-1 This host is not entitled to run Veritas Storage
Foundation/Veritas Cluster Server.As set forth in the End User
License Agreement (EULA) you must complete one of the two options
set forth below. To comply with this condition of the EULA and
stop logging of this message, you have <nn> days to either:
- make this host managed by a Management Server (see
  http://go.symantec.com/sfhakeyless for details and free download),
  or
- add a valid license key matching the functionality in use on this host
  using the command 'vxlicinst'
```

To comply with the terms of the EULA, and remove these messages, you must do one of the following within 60 days:

- Install a valid license key corresponding to the functionality in use on the host. See “[Installing Veritas SFHA Solutions product license keys](#)” on page 27. After you install the license key, you must validate the license key using the following command:

```
# /opt/VRTS/bin/vxkeyless
```

- Continue with keyless licensing by managing the server or cluster with a management server.
For more information about keyless licensing, see the following URL:
<http://go.symantec.com/sfhakeyless>

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.
```

```
Failed to setup rsh communication on 10.198.89.241:
'rsh 10.198.89.241 <command>' failed
Trying to setup ssh communication on 10.198.89.241.
Failed to setup ssh communication on 10.198.89.241:
Login denied
```

```
Failed to login to remote system(s) 10.198.89.241.
Please make sure the password(s) are correct and superuser(root)
can login to the remote system(s) with the password(s).
If you want to setup rsh on remote system(s), please make sure
rsh with command argument ('rsh <host> <command>') is not
denied by remote system(s).
```

```
Either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

```
Would you like the installer to setup ssh/rsh communication
automatically between the nodes?
Superuser passwords for the systems will be asked. [y,n,q] (y) n
```

```
System verification did not complete successfully
```

```
The following errors were discovered on the systems:
```

```
The ssh permission denied on 10.198.89.241
rsh exited 1 on 10.198.89.241
either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

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

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 133.

Note: Remove remote shell permissions after completing the SF installation and configuration.

Resource temporarily unavailable

If the installation fails with the following error message on the console:

```
fork() failed: Resource temporarily unavailable
```

The value of `nkthread` tunable parameter may not be large enough. The `nkthread` tunable requires a minimum value of 600 on all systems in the cluster. To determine the current value of `nkthread`, enter:

```
# kctune -q nkthread
```

If necessary, you can change the value of `nkthread` using the SAM (System Administration Manager) interface, or by running the `kctune` command. If you change the value of `nkthread`, the kernel must be rebuilt for the new value to take effect. It is easier to change the value using SAM because there is an option to process the new kernel immediately.

See the `kctune(1M)` and `sam(1M)` manual pages.

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

Verifying systems: 12% .....
Estimated time remaining: 0:10 1 of 8
Checking system communication ..... Done
System verification did not complete successfully
The following errors were discovered on the systems:
cannot resolve hostname host1
Enter the system names separated by spaces: q,? (host1)
```

Suggested solution: Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

If a system cannot access the software source depot, either `swagentd` is not running on the target system or the `swlist` command cannot see the source depot.

```
Correct /etc/{hosts, nsswitch.conf} and continue from here
Continue? [Y/N] :
```

Suggested solutions: check that `swagentd` is running. Check whether there is an entry for the target system in `/etc/hosts`. If there is no entry, then ensure the `hosts` file is not the primary lookup for the "hosts" entry.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 (2184482)

When upgrading from SF version 5.0 or 5.0.1 to SF 6.0 the `S*vxdmbs3` startup script is renamed to `NO_S*vxdmbs3`. The `S*vxdmbs3` startup script is required by

`sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbms3` startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround

Before running `sfua_rept_migrate`, rename the startup script `NO_S*vxdbms3` to `S*vxdbms3`.

Compatibility issues when installing Storage Foundation with other products

This appendix includes the following topics:

- [Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present](#)

Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present

Installing Storage Foundation when other Veritas products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where SFM or VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the SFM Central Server and Managed Host depots as is.
- When uninstalling Storage Foundation products where SFM or VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where SFM or VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspbx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspbx, VRTSicsco, and VRTSat.

Index

C

- cluster functionality
 - enabling 54
 - environment requirements 21
 - shared disks 55
- configuration daemon (vxconfigd)
 - starting 52
- configuring
 - rsh 30
 - shared disks 55
 - ssh 30
- creating SD bundle 45

D

- disk groups
 - rootdg 51

I

- I/O daemon (vxiod)
 - starting 53
- Ignite
 - installing 45
- Ignite-UX 47
 - installing standalone 46
- Installing
 - SF with the Web-based installer 42
- installing
 - Ignite 45
 - Ignite-UX 46–47
 - standalone 46

K

- kctune command 149

M

- mounting
 - software disc 31

P

- planning to upgrade VVR 68
- preinstallation 68

R

- removing
 - the Replicated Data Set 99
- Replicated Data Set
 - removing the 99
- root disk group 51
- rsh
 - configuration 30

S

- sam command 149
- SD bundle 45
- shared disks, configuring 55
- ssh
 - configuration 30
- starting vxconfigd configuration daemon 52
- starting vxiod daemon 53

T

- tunables file
 - about setting parameters 121
 - parameter definitions 126
 - preparing 125
 - setting for configuration 122
 - setting for installation 122
 - setting for upgrade 122
 - setting parameters 125
 - setting with no other operations 123
 - setting with un-integrated response file 124

U

- upgrading VVR
 - from 4.1 69
 - planning 68

V

- verifying installation
 - kernel component 97
- Veritas Operations Manager 15
- vradmin
 - delpri 100
 - stoprep 100
- VVR 4.1
 - planning an upgrade from 69
- vxconfigd configuration daemon
 - starting 52
- vxctl mode command 53
- vxiod I/O daemon
 - starting 53

W

- Web-based installer 42