Symantec™ Storage Foundation 6.1 Installation Guide - Solaris

October 2015

VERITAS™

Symantec™ 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: 6.1

Document version: 6.1 Rev 3

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, whether delivered by Symantec as on premises or hosted services. 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 15

Chapter 1 Introducing Storage Foundation 16

About Storage Foundation 16
About Symantec Replicator Option 17
About Veritas Operations Manager 17
About Symantec Operations Readiness Tools (SORT) 17

Chapter 2 System requirements 20

Release notes 20
Important preinstallation information for SF 21
Supported operating systems 21
Veritas File System requirements 21
Disk space requirements 22
Checking installed product versions and downloading maintenance releases and hot fixes 23
Obtaining installer hot fixes 24
Disabling external network connection attempts 25
Database requirements 26

Chapter 3 Planning to install Storage Foundation 27

About planning for SF installation 27
About installation and configuration methods 28
About response files 30
Downloading the Storage Foundation software 31

Chapter 4 Licensing Storage Foundation 33

About Symantec product licensing 33
Setting or changing the product level for keyless licensing 34
Installing Symantec product license keys 36

Section 2 Installation of Storage Foundation 38

Chapter 5 Preparing to install Storage Foundation 39

Installation preparation overview 39
About using ssh or rsh with the installer 40
Creating a root user 41
Creating the /opt directory 42
Setting environment variables 42
Mounting the product disc 42
Assessing the system for installation readiness 43
Prechecking your systems using the installer 43
Making the IPS publisher accessible 44

Chapter 6 Installing Storage Foundation using the script-based installer 46

About the script-based installer 46
Installing Storage Foundation using the script-based installer 48
Installing language packages 50

Chapter 7 Installing Storage Foundation using the web-based installer 51

About the web-based installer 51
Before using the web-based installer 52
Starting the web-based installer 52
Obtaining a security exception on Mozilla Firefox 53
Performing a preinstallation check with the web-based installer 54
Installing SF with the web-based installer 54

Chapter 8 Automated installation using response files 57

Installing SF using response files 57
Response file variables to install Storage Foundation 58
Sample response file for SF installation 60
Configuring SF using response files 61
Response file variables to configure Storage Foundation 62

Chapter 9 Installing Storage Foundation using operating system-specific methods 64

About installing SF using operating system-specific methods 64
Installing SF on Solaris 11 using Automated Installer 65

About Automated Installation 65
Using Automated Installer 66
Using AI to install the Solaris 11 operating system and SFHA products 66
Installing SF on Solaris 10 using JumpStart 70
Overview of JumpStart installation tasks 70
Generating the finish scripts 71
Preparing installation resources 72
Adding language pack information to the finish file 73
Using a Flash archive to install SF and the operating system 74
Creating the Symantec post-deployment scripts 75
Manually installing SF using the system command 76
Installing SF on Solaris 10 using the pkgadd command 76
Installing SF on Solaris 11 using the pkg install command 78

Chapter 10 Configuring Storage Foundation 82

Configuring Storage Foundation using the installer 82
Configuring Storage Foundation manually 82
Configuring Veritas Volume Manager 83
Configuring Veritas File System 86
Configuring the SFDB repository database after installation 87

Section 3 Managing your Symantec deployments 88

Chapter 11 Performing centralized installations using the Deployment Server 89

About the Deployment Server 90
How to install the Deployment Script 91
Deployment management overview 92
Setting up a Deployment Server 93
Setting deployment preferences 95
Using the Deployment Server command line option to specify a non-default repository location 96
Using the Deployment Server command line options to load and download the most recent release information 97
Viewing or downloading available release images 98
Viewing or removing repository images stored in your repository 102
Deploying Symantec product updates to your environment 105
Finding out which releases you have, and which upgrades or updates you may need 106

Deploying a specific Symantec release 108
Updating release information on systems without Internet access 109

Section 4 Upgrade of SF 111

Chapter 12 Planning to upgrade SF 112

Upgrade methods for SF 112
Supported upgrade paths for SF 6.1 113
About using the installer to upgrade when the root disk is encapsulated 116
Preparing to upgrade SF 117
Getting ready for the upgrade 117
Creating backups 120
Tasks for upgrading the Storage Foundation for Databases (SFDB) 121
Determining if the root disk is encapsulated 121
Pre-upgrade tasks for migrating the SFDB repository database 121
Pre-upgrade planning for Volume Replicator 122
Verifying that the file systems are clean 125
Upgrading the array support 126
Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes 127

Chapter 13 Upgrading Storage Foundation 130

Upgrading Storage Foundation with the product installer when OS upgrade is not required 130
Upgrading Storage Foundation to 6.1 using the product installer or manual steps 133
Upgrading Storage Foundation with the product installer 134
Upgrading Symantec Storage Foundation to 6.1 using upgrade scripts (OS upgrade) 136
Upgrading SF using the web-based installer 139
Upgrading the Solaris operating system 141
Upgrading Volume Replicator 143
Upgrading VVR without disrupting replication 143
Upgrading language packages 145

Chapter 14 Performing an automated SF upgrade using response files 146

Upgrading SF using response files 146
Response file variables to upgrade Storage Foundation 147
Sample response file for SF upgrade 149

Chapter 15 Upgrading SF using Live Upgrade and Boot Environment upgrade 151

About Live Upgrade 151
About ZFS Boot Environment (BE) upgrade 152
Supported upgrade paths for Live Upgrade and Boot Environment upgrade 153
Performing Live Upgrade in a Solaris zone environment on Solaris 10 155
Performing Live Upgrade on Solaris 10 systems 155
Before you upgrade SF using Solaris Live Upgrade 156
Creating a new Solaris 10 boot environment on the alternate boot disk 157
Upgrading SF using the installer for Solaris 10 Live Upgrade 161
Upgrading SF manually for Solaris 10 Live Upgrade 161
Completing the Solaris 10 Live Upgrade 163
Verifying the Solaris 10 Live Upgrade of SF 164
Administering boot environments in Solaris 10 Live Upgrade 165
Performing Boot Environment upgrade on Solaris 11 systems 167
Creating a new Solaris 11 BE on the primary boot disk 168
Upgrading SF using the installer for upgrading BE on Solaris 11 168
Upgrading SF manually for Solaris 11 Live Upgrade 169
Completing the SF upgrade on BE on Solaris 11 170
Verifying Solaris 11 BE upgrade 171
Administering BEs in Solaris 11 BE upgrade 172
About Live Upgrade in a Volume Replicator (VVR) environment 173

Chapter 16 Migrating from Storage Foundation Basic to Storage Foundation Standard 174

Migrating from Storage Foundation Basic to Storage Foundation Standard 174

Chapter 17 Performing post-upgrade tasks 177

Optional configuration steps 177
Re-joining the backup boot disk group into the current disk group 178
Reverting to the backup boot disk group after an unsuccessful upgrade 178
Post upgrade tasks for migrating the SFDB repository database 179
Migrating from a 5.0 repository database to 6.1 180
Migrating from a 5.1 or higher repository database to 6.1 182
After upgrading from 5.0.x and before migrating SFDB 184
Recovering VVR if automatic upgrade fails 184
Upgrading disk layout versions 185
Upgrading VxVM disk group versions 186
Updating variables 186
Setting the default disk group 186
Upgrading the Array Support Library 187
Adding JBOD support for storage arrays for which there is not an ASL available 187
Unsuppressing DMP for EMC PowerPath disks 188
Converting from QuickLog to Multi-Volume support 197
Verifying the Storage Foundation upgrade 198

Section 5 Post-installation tasks 199

Chapter 18 Performing post-installation tasks 200

Changing root user into root role 200
Switching on Quotas 201
Enabling DMP support for native devices 201

Chapter 19 Verifying the SF installation 202

Verifying that the products were installed 202
Installation log files 203
Using the installation log file 203
Using the summary file 203
Starting and stopping processes for the Symantec products 203
Checking Veritas Volume Manager processes 204
Checking Veritas File System installation 205
Verifying Veritas File System kernel installation 205
Verifying command installation 205

Section 6 Uninstallation of SF 206

Chapter 20 Uninstalling Storage Foundation 207

About removing Storage Foundation 207
Preparing to uninstall 208
Preparing to remove Veritas Volume Manager 208
Preparing to remove Veritas File System 216
Removing the Replicated Data Set 217
Uninstalling SF packages using the script-based installer 218
Uninstalling SF with the web-based installer 220
Uninstalling Storage Foundation using the `pkgrm` or `pkg uninstall` command 221
Uninstalling the language packages using the `pkgrm` command 222
Manually uninstalling Storage Foundation packages on non-global zones on Solaris 11 222
Removing the Storage Foundation for Databases (SFDB) repository after removing the product 223

Chapter 21 Uninstalling SF using response files 225

Uninstalling SF using response files 225
Response file variables to uninstall Storage Foundation 226
Sample response file for SF uninstallation 226

Section 7 Installation reference 228

Appendix A Installation scripts 229

Installation script options 229

Appendix B Tunable files for installation 235

About setting tunable parameters using the installer or a response file 235
Setting tunables for an installation, configuration, or upgrade 236
Setting tunables with no other installer-related operations 237
Setting tunables with an un-integrated response file 238
Preparing the tunables file 239
Setting parameters for the tunables file 239
Tunables value parameter definitions 240

Appendix C Configuring the secure shell or the remote shell for communications 248

About configuring secure shell or remote shell communication modes before installing products 248
Manually configuring and passwordless ssh 249
Restarting the ssh session 253
Enabling and disabling rsh for Solaris 254

Appendix D Storage Foundation components 256

Storage Foundation installation packages 256
Chinese language packages 259
Japanese language packages 259
Symantec Storage Foundation obsolete and reorganized installation packages 260

Appendix E Troubleshooting installation issues 263

Restarting the installer after a failed connection 263
What to do if you see a licensing reminder 263
About the VRTSspt package troubleshooting tools 264
Incorrect permissions for root on remote system 265
Inaccessible system 266
Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.1 (2184482) 266
Workaround 266
Troubleshooting the webinstaller 266

Appendix F Compatability issues when installing Storage Foundation with other products 268

Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present 268
Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present 269
Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present 269

Index 270

Section 1

Installation overview and planning

- Chapter 1. Introducing Storage Foundation
- Chapter 2. System requirements
- Chapter 3. Planning to install Storage Foundation
- Chapter 4. Licensing Storage Foundation

Chapter 1

Introducing Storage Foundation

This chapter includes the following topics:

- About Storage Foundation
- About Veritas Operations Manager
- About Symantec Operations Readiness Tools (SORT)

About Storage Foundation

Symantec Storage Foundation 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.

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

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

About Symantec Replicator Option

Symantec Replicator Option is an optional, separately-licensable feature.

Volume Replicator replicates data to remote locations over any standard IP network to provide continuous data availability.

About Veritas Operations Manager

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

Symantec recommends using Veritas Operations Manager (VOM) to manage Storage Foundation and Cluster Server environments.

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. Symantec Storage Foundation Management Server is deprecated.

About Symantec Operations Readiness Tools (SORT)

Symantec Operations Readiness Tools (SORT) is a website that automates and simplifies some of the most time-consuming administrative tasks. It helps you identify risks in your datacenters and improve operational efficiency, enabling you to manage the complexity that is associated with datacenter architectures and scale

Table 1-1 lists three major datacenter tasks and the SORT tools that can help you accomplish them.

Table 1-1 Datacenter tasks and the SORT tools

Task	SORT tools
Prepare for installations and upgrades	■ Installation and Upgrade checklists Display system requirements including memory, disk space, and architecture. ■ Installation and Upgrade custom reports Create reports that determine if you're ready to install or upgrade a Symantec enterprise product. ■ Array-specific Module Finder List the latest Array Support Libraries (ASLs) and Array Policy Modules (APMs) for UNIX servers, and Device Driver Installers (DDIs) and Device Discovery Layers (DDLs) for Windows servers. ■ High Availability Agents table Find and download the agents for applications, databases, replication, and Symantec partners.
Identify risks and get server-specific recommendations	■ Patch notifications Receive automatic email notifications about patch updates. (Sign in required.) ■ Risk Assessment check lists Display configuration recommendations based on your Symantec product and platform. ■ Risk Assessment custom reports Create reports that analyze your system and give you recommendations about system availability, storage use, performance, and best practices. ■ Error code descriptions and solutions Display detailed information on thousands of Symantec error codes.

Table 1-1 Datacenter tasks and the SORT tools *(continued)*

Task	SORT tools
Improve efficiency	■ Patch Finder List and download patches for your Symantec enterprise products. ■ License/Deployment custom reports Create custom reports that list your installed Symantec products and license keys. Display licenses by product, platform, server tier, and system. ■ Symantec Performance Value Unit (SPVU) Calculator Use the calculator to assist you with the pricing meter transition. ■ Documentation List and download Symantec product documentation, including manual pages, product guides, and support articles. ■ Related links Display links to Symantec product support, forums, customer care, and vendor information on a single page.

SORT is available at no additional charge.

To access SORT, go to:

https://sort.symantec.com

Chapter 2

System requirements

This chapter includes the following topics:

- Release notes
- Important preinstallation information for SF
- Supported operating systems
- Veritas File System requirements
- Disk space requirements
- Checking installed product versions and downloading maintenance releases and hot fixes
- Obtaining installer hot fixes
- Disabling external network connection attempts
- Database requirements

Release notes

The *Release Notes* for each Symantec 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

Important preinstallation information for SF

Before you install SF, make sure that you have reviewed the following information:

- Preinstallation checklist for your configuration. Go to the SORT installation checklist tool. From the drop-down lists, select the information for the Symantec product you want to install, and click **Generate Checklist**.
- Hardware compatibility list for information about supported hardware: http://www.symantec.com/docs/TECH211575
- For important updates regarding this release, review the Late-Breaking News Technote on the Symantec Technical Support website: http://www.symantec.com/docs/TECH211540

Supported operating systems

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

Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x6000 (for Solaris 10) and 0x8000 (for Solaris 11). When you install the Veritas File System package, `VRTSvxfs`, the VRTSvxfs packaging scripts check the values of these variables in the kernel. If the values are less than the required values, VRTSvxfs increases the values and modifies the `/etc/system` file with the required values. If the VRTSvxfs scripts increase the values, the installation proceeds as usual except that you must reboot and restart the installation program. A message displays if a reboot is required.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System. Use the following commands to check the values of the variables:

For Solaris 10:

```
# echo "lwp_default_stksize/X" | mdb -k
lwp_default_stksize:
lwp_default_stksize:            6000

# echo "svc_default_stksize/X" | mdb -k
svc_default_stksize:
svc_default_stksize:            6000
```

For Solaris 11:
```
# echo "lwp_default_stksize/X" | mdb -k
lwp_default_stksize:
lwp_default_stksize:            8000

# echo "svc_default_stksize/X" | mdb -k
svc_default_stksize:
svc_default_stksize:            8000
```

If the values shown are less than 6000 (for Solaris 10) and less than 8000 (for Solaris 11), you can expect a reboot after installation.

Note: The default value of the `svc_default_stksize` variable is 0 (zero), which indicates that the value is set to the value of the `lwp_default_stksize` variable. In this case, no reboot is required, unless the value of the `lwp_default_stksize` variable is too small.

To avoid a reboot after installation, you can modify the `/etc/system` file with the appropriate values. Reboot the system prior to installing the packages. Add the following lines to the `/etc/system` file:

For Solaris 10:
```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
```

For Solaris 11:
```
set lwp_default_stksize=0x8000
set rpcmod:svc_default_stksize=0x8000
```

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 to determine whether there is sufficient space.

Or, go to the installation directory and run the installer with the `-precheck` option.

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

See “About the script-based installer” on page 46.

Checking installed product versions and downloading maintenance releases and hot fixes

Symantec provides a means to check the Symantec packages you have installed, and download any needed maintenance releases and hot fixes.

Use the `installer` command with the `-version` option to determine what is installed on your system, and download any needed maintenance releases or hot fixes. After you have installed the current version of the product, you can use the `showversion` script in the `/opt/VRTS/install` directory to find product information.

The `version` option or the `showversion` script checks the specified systems and discovers the following:

- Storage Foundation and High Availability product versions that are installed on the system
- All the required packages and the optional Symantec packages installed on the system
- Any required or optional packages (if applicable) that are not present
- Installed hot fixes
- Available base releases (major or minor)
- Available maintenance releases
- Available hot fix releases

To check your systems and download maintenance releases and hot fixes

1. Mount the media, or navigate to the installation directory.
2. Start the installer with the `-version` option.

   ```
   # ./installer -version sys1 sys2
   ```

   For each system, the installer lists all of the installed base releases, maintenance releases, and hot fixes, followed by the lists of available downloads.

3 If you have Internet access, follow the prompts to download the available maintenance releases and hot fixes to the local system.

4 If you do not have Internet access, you can download any needed maintenance releases and hot fixes from the Symantec Operations Readiness Tools (SORT) Patch Finder page at:

https://sort.symantec.com/patch/finder

You can obtain installer hot fixes automatically or manually.

See "Obtaining installer hot fixes" on page 24.

Downloading maintenance releases and hot fixes requires the installer to make outbound networking calls. You can also disable external network connection attempts.

See "Disabling external network connection attempts" on page 25.

Obtaining installer hot fixes

Symantec occasionally finds issues with the Storage Foundation installer, and posts public installer hot fixes on the Symantec Operations Readiness Tools (SORT) website's Patch Finder page at:

https://sort.symantec.com/patch/finder

You can access installer hot fixes automatically or manually.

To download installer hot fixes automatically

- Starting with Storage Foundation version 6.1, installer hot fixes are downloaded automatically. No action is needed on your part.

  If you are running Storage Foundation version 6.1 or later, and your system has Internet access, the installer automatically imports any needed installer hot fix, and begins using it.

Automatically downloading installer hot fixes requires the installer to make outbound networking calls. You can also disable external network connection attempts.

See "Disabling external network connection attempts" on page 25.

If your system does not have Internet access, you can download installer hot fixes manually.

To download installer hot fixes manually

1. Go to the Symantec Operations Readiness Tools (SORT) website's Patch Finder page, and save the most current Symantec patch on your local system.
2. Navigate to the directory where you want to unzip the file you downloaded in step 1.
3. Unzip the patch tar file. For example, run the following command:

```
# gunzip cpi-6.1P2-patches.tar.gz
```

4. Untar the file. For example, enter the following:

```
# tar -xvf cpi-6.1P2-patches.tar
patches/
patches/CPI61P2.pl
README
```

5. Navigate to the installation media or to the installation directory.
6. To start using the patch, run the `installer` command with the `-require` option. For example, enter the following:

```
# ./installer -require /target_directory/patches/CPI61P2.pl
```

Disabling external network connection attempts

When you execute the `installer` command, the installer attempts to make an outbound networking call to get information about release updates and installer hot fixes. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network connection attempts by the installer.

To disable external network connection attempts

- Disable inter-process communication (IPC).

  To do this, run the installer with the `-noipc` option.

  For example, to disable IPC for system1 (sys1) and system2 (sys2) enter the following:

```
# ./installer -noipc sys1 sys2
```

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 and Sybase on VxFS and VxVM.

SF does not support running SFDB tools with Sybase.

Chapter 3

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 have the current version of this guide. The latest documentation is available on the Symantec Operations Readiness Tools (SORT) website.

https://sort.symantec.com/documents

Document version: 6.1 Rev 3.

This installation guide is designed for system administrators who already have basic knowledge of UNIX system and network administration. Basic knowledge includes commands such as `tar`, `mkdir`, and simple shell scripting. What is also required is familiarity with the specific platform and operating system where SF is to be installed.

Follow the preinstallation instructions if you want to install Storage Foundation.

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

- Symantec Storage Foundation Basic
- Symantec 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 Symantec installation programs or using native operating system methods.

Table 3-1 shows the installation and configuration methods that SF supports.

Table 3-1 Installation and configuration methods

Method	Description

The script-based installer	Using the script-based installer, you can install Symantec products (version 6.1 and later) from a driver system running a supported platform to target computers running any supported platform.
	
	To install your Symantec product using the installer, choose one of the following:
	
	■ The general product installer:
	`installer`
	The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
	
	■ Product-specific installation scripts:
	`installsf`
	The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installsf` script is identical to running the general product installer and specifying SF from the list of products to install.
	Use the product-specific installation scripts to install or configure individual products you download electronically.
	
	See "About the script-based installer" on page 46.

Table 3-1 Installation and configuration methods *(continued)*

Method	Description
The web-based installer	Using the web-based installer, you can install Symantec products (version 6.1 and later) from a driver system running a supported platform to target computers running any supported platform The web-based installer provides an interface to manage the installation and configuration from a remote site using a standard web browser. `webinstaller` See "About the web-based installer" on page 51.
Deployment Server	Using the Deployment Server, you can store multiple release images in one central location and deploy them to systems of any supported platform. See "About the Deployment Server" on page 90.
Silent installation using response files	Response files automate installation and configuration by using the information that is stored in a specified file instead of prompting you for information. 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 option to install silently on one or more systems. See "About response files" on page 30.
Install Bundles	Beginning with version 6.1, you can easily install or upgrade your systems directly to a base, maintenance, or hot fix level in one step using Install Bundles,. The installer installs both releases as if they were combined in the same release image. The various scripts, packages, and patch components are merged, and multiple releases are installed together as if they are one combined release. See "Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes" on page 127.

Table 3-1 Installation and configuration methods *(continued)*

Method	Description
JumpStart (For Solaris 10 systems)	You can use the product installer of the product-specific installation script to generate a JumpStart script file. Use the generated script to install Symantec packages from your JumpStart server. See "Installing SF on Solaris 10 using JumpStart" on page 70.
Flash Archive (For Solaris 10 systems)	You can use the product installer to clone the system and install the Symantec products on the master system. See "Using a Flash archive to install SF and the operating system" on page 74.
Manual installation and configuration	Manual installation uses the Solaris commands to install SF. To retrieve a list of all packages and patches required for all products in the correct installation order, enter: `# installer -allpkgs` Use the Solaris commands to install SF. Then manually or interactively configure SF. See "Manually installing SF using the system command" on page 76.
Automated Installer (For Solaris 11 systems)	You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Symantec packages on multiple client systems in a network. AI performs a hands-free installation (automated installation without manual interactions) of SPARC systems. See "Installing SF on Solaris 11 using Automated Installer" on page 65.

About response files

The installer script 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 229.

Syntax in the response file

The syntax of the Perl statements that is 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 1 ", "value 2 ", "value 3 "];
```

Downloading the Storage Foundation software

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

For a Trialware download, perform the following. Contact your Symantec 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. In Products and Solutions section, click the **Trialware** link.
3. On the next page near the bottom of the page, click **Business Continuity**.
4. Under Cluster Server, click **Download**.
5. In the new window, click **Download Now**.
6. Review the terms and conditions, and click **I agree**.
7. You can use existing credentials to log in or create new credentials.
8. Find the product that you want to download and select it. Continue with the installation.

If you download a standalone Symantec 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 the script-based installer" on page 46.

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 2 GB for SPARC.

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

   See "Disk space requirements" on page 22.

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.

   ```
   # /usr/bin/df -l filesystem
   ```

   Caution: When you select a location to download files, do not select a directory that contains Symantec 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.

Chapter 4

Licensing Storage Foundation

This chapter includes the following topics:

- About Symantec product licensing
- Setting or changing the product level for keyless licensing
- Installing Symantec product license keys

About Symantec product licensing

You have the option to install Symantec 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.

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

http://www.symantec.com/products-solutions/licensing/activating-software/detail.jsp?detail_id=licensing_portal

The product installer prompts you to select one of the following licensing methods:

- 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. 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 upgrade to this release from a previous release of the Symantec software, the installer asks whether you want to upgrade the key to the new version. 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:

- 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 product level for keyless licensing" on page 34.
  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 Symantec product license keys" on page 36.
  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: To change from one product group to another, you may need to perform additional steps.

Setting or changing the product level for keyless licensing

The keyless licensing method uses product levels to determine the Symantec products and functionality that are licensed.

For more information to use keyless licensing 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.

When you upgrade from a previous release, the product installer prompts you to update the `vxkeyless` license product level to the current release level. If you update the `vxkeyless` license product level during the upgrade process, no further action is required. If you do not update the `vxkeyless` license product level, the output you see when you run the `vxkeyless display` command includes the previous release's `vxkeyless` license product level. Each `vxkeyless` license product level name includes the suffix *_previous_release_version*. For example, DMP_6.0, or SFENT_VR_5.1SP1, or VCS_GCO_5.1. If there is no suffix, it is the current release version.

You would see the suffix *_previous_release_version* if you did not update the `vxkeyless` product level when prompted by the product installer. Symantec highly recommends that you always use the current release version of the product levels. To do so, use the `vxkeyless set` command with the desired product levels. If you see SFENT_60, VCS_60, use the `vxkeyless set SFENT,VCS` command to update the product levels to the current release.

After you install or upgrade, 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. Change your current working directory:

   ```
   # export PATH=$PATH:/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 Symantec products until you install a new key or set a new product level.

See "Installing Symantec product license keys" on page 36.

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 Symantec product license keys

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

`vxlicinst`	Installs a license key for a Symantec product
`vxlicrep`	Displays the currently installed licenses
`vxlictest`	Retrieves the features and their descriptions that are 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 or change a license

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

   ```
   # cd /opt/VRTS/bin

   # ./vxlicinst -k license key
   ```

2. Run the following Veritas Volume Manager (VxVM) command to recognize the new license:

   ```
   # vxdctl license init
   ```

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

If you have `vxkeyless` licensing, you can view or update the keyless product licensing levels.

See "Setting or changing the product level for keyless licensing" on page 34.

Section 2

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. Automated installation using response files
- Chapter 9. Installing Storage Foundation using operating system-specific methods
- Chapter 10. Configuring Storage Foundation

Chapter 5

Preparing to install Storage Foundation

This chapter includes the following topics:

- Installation preparation overview
- About using ssh or rsh with the installer
- Creating a root user
- Creating the /opt directory
- Setting environment variables
- Mounting the product disc
- Assessing the system for installation readiness
- Making the IPS publisher accessible

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 Symantec product licensing" on page 33.

Table 5-1 Installation overview *(continued)*

Installation task	Section
Download the software, or insert the product DVD.	See "Downloading the Storage Foundation software" on page 31. See "Mounting the product disc" on page 42.
Set environment variables.	See "Setting environment variables" on page 42.
Create the `/opt` directory, if it does not exist.	See "Creating the /opt directory" on page 42.
Configure the Secure Shell (ssh) or Remote Shell (rsh) on all nodes.	See "About configuring secure shell or remote shell communication modes before installing products" on page 248.
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 script-based installer" on page 46.

About using ssh or rsh with the installer

The installer uses passwordless Secure Shell (ssh) or Remote Shell (rsh) communications among systems. The installer uses the ssh daemon or rsh daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. Or, you can run the `installer -comsetup` command to set up ssh or rsh explicitly. You then provide the installer with the superuser passwords for the systems where you plan to install. The ssh communication or rsh 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 configuration or rsh configuration from the systems.

See "Installation script options" on page 229.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer can configure ssh or rsh on the target systems. In the following scenarios, you need to set up ssh or rsh manually, or use the `installer -comsetup` option to set up an ssh or rsh configuration from the systems.

- When you perform installer sessions using a response file.

See “About configuring secure shell or remote shell communication modes before installing products” on page 248.

Creating a root user

On Oracle Solaris 11, you need to change the root role into a user as you cannot directly log in as root user.

To change root role into a user

1. Log in as local user and assume the root role.

```
% su  - root
```

2. Remove the root role from local users who have been assigned the role.

```
# roles admin

root

# usermod -R " " admin
```

3. Change the root role into a user.

```
# rolemod -K type=normal root
```

4. Verify the change.

   - `# getent user_attr root`

     ```
     root::::auths=solaris.*;profiles=All;audit_flags=lo\
     :no;lock_after_retries=no;min_label=admin_low;clearance=admin_high
     ```

     If the `type` keyword is not present in the output or is equal to normal, the account is not a role.

   - `# userattr type root`

     If the output is empty or lists normal, the account is not a role.

Note: For more information, see the Oracle documentation on Oracle Solaris 11 operating system.

Note: After installation, you may want to change root user into root role to allow local users to assume the root role.

See "Changing root user into root role" on page 200.

Creating the /opt directory

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

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

Setting environment variables

Most of the commands which are used in the installation are present 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`.

Some VCS custom scripts reside in `/opt/VRTSvcs/bin`. If you want to install a high availability product, add `/opt/VRTSvcs/bin` to the `PATH` also.

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

- If you want to use `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 want to use 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

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 into a DVD drive that is connected to your system.

3. If Solaris volume management software is running on your system, the software disc automatically mounts as /cdrom/cdrom0.

4. If Solaris volume management software is not available to mount the DVD, you must mount it manually. After you insert the software disc, enter:

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

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

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 6.1.

Symantec Operations Readiness Tools	Symantec Operations Readiness Tools (SORT) is a web-based application that is designed to support Symantec enterprise products. See "About Symantec Operations Readiness Tools (SORT)" on page 17.
Prechecking your systems using the installer	Performs a preinstallation check on the specified systems. The product installer reports whether the specified systems meet the minimum requirements for installing Storage Foundation 6.1. See "Prechecking your systems using the installer" on page 43.

Prechecking your systems using the 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 Symantec programs for best performance
- Required operating system versions

To use the precheck option

1. Start the script-based or web-based installer.

   See "Installing Storage Foundation using the script-based installer" on page 48.

   See "Installing SF with the web-based installer" on page 54.

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. Enter the system name or the IP address of the system that you want to check.
4. Review the output and make the changes that the installer recommends.

Making the IPS publisher accessible

The installation of SF 6.1 fails on Solaris 11 if the Image Packaging System (IPS) publisher is inaccessible. The following error message is displayed:

CPI ERROR V-9-20-1273 Unable to contact configured publishers on <node_name>.

Solaris 11 introduces the new Image Packaging System (IPS) and sets a default publisher (solaris) during Solaris installation. When additional packages are being installed, the set publisher must be accessible for the installation to succeed. If the publisher is inaccessible, as in the case of a private network, then package installation will fail. The following commands can be used to display the set publishers:

```
# pkg publisher
```

Example:

```
root@sol11-03:~# pkg publisher
PUBLISHER       TYPE     STATUS   URI
solaris         origin   online   http://pkg.oracle.com/solaris/release/
root@sol11-03:~# pkg publisher solaris
Publisher: solaris
               Alias:
          Origin URI: http://pkg.oracle.com/solaris/release/
```

```
          SSL Key: None
         SSL Cert: None
      Client UUID: 00000000-3f24-fe2e-0000-000068120608
  Catalog Updated: October 09:53:00 PM
          Enabled: Yes
Signature Policy: verify
```

To make the IPS publisher accessible

1. Enter the following to disable the publisher (in this case, solaris):

   ```
   # pkg set-publisher --disable solaris
   ```

2. Repeat the installation of SF 6.1.

3. Re-enable the original publisher. If the publisher is still inaccessible (private network), then the `no-refresh` option can be used to re-enable it.

   ```
   # pkg set-publisher --enable solaris
   ```

   or

   ```
   # pkg set-publisher --enable --no-refresh solaris
   ```

Note: Unsetting the publisher will have a similar effect, except that the publisher can only be re-set if it is accessible. See pkg(1) for further information on the pkg utility.

Chapter 6

Installing Storage Foundation using the script-based installer

This chapter includes the following topics:

- About the script-based installer
- Installing Storage Foundation using the script-based installer
- Installing language packages

About the script-based installer

You can use the script-based installer to install Symantec products (version 6.1 and later) from a driver system that runs any supported platform to a target system that runs different supported platforms.

To install your Symantec product, use one of the following methods:

- The general product installer (`installer`). The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
  See "Installing Storage Foundation using the script-based installer" on page 48.
- Product-specific installation scripts (`installsf`). The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installsf` script is identical to running the general product installer and specifying SF from the list of products to install. Use the product-specific installation scripts to install or configure individual products you download

electronically. You can find these scripts at the root of the product media. These scripts are also installed with the product.

Table 6-1 lists all the SFHA Solutions product installation scripts. The list of product-specific installation scripts that you find on your system depends on the product that you install on your system.

Table 6-1 Product installation scripts

Symantec product name	Script name in the media	Script name after an installation
For all SFHA Solutions products	`installer`	N/A
Symantec ApplicationHA	`installapplicationha`	`installapplicationha` *<version>*
Symantec Cluster Server (VCS)	`installvcs`	`installvcs`*<version>*
Symantec Storage Foundation (SF)	`installsf`	`installsf`*<version>*
Symantec Storage Foundation and High Availability (SFHA)	`installsfha`	`installsfha`*<version>*
Symantec Storage Foundation Cluster File System High Availability (SFCFSHA)	`installsfcfsha`	`installsfcfsha`*<version>*
Symantec Storage Foundation for Oracle RAC (SF Oracle RAC)	`installsfrac`	`installsfrac`*<version>*
Symantec Dynamic Multi-pathing (DMP)	`installdmp`	`installdmp`*<version>*

When you install from the installation media, the script name does not include a product version.

When you configure the product after an installation, the installation scripts include the product version in the script name.

For example, for the 6.1 version:

```
# /opt/VRTS/install/installsf61 -configure
```

Note: The general product installer (`installer`) script does not include the product version.

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 229.

Installing Storage Foundation using the script-based installer

The 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 248.

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 42.

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

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

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_SFHA_Ux_version.pdf file present on the media? [y,n,q,?] y
```

8 Select from one of the following installation options:

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

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,?] sys1
```

10 After the system checks complete, the installer displays a list of the packages to be installed. Press Enter to continue with the installation.

11 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.

12 The installer may prompt to restore previous Veritas Volume Manager configurations.

13 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 Symantec products. Keyless licensing requires that you manage the systems with a Management Server.

See "About Symantec product licensing" on page 33.

14 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
```

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

```
Installation procedures and diagnostic information were saved in the
log files under directory /var/tmp/installer-<platform>-<uuid>.
Analyzing this information helps Symantec discover and
fix failed operations performed by the installer.
Would you like to send the information about this installation to
Symantec to help improve installation in the future? [y,n,q,?] (y) y
```

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

Installing language packages

To install SF in a language other than English, install the required language packages after installing the English packages.

To install the language packages on the server

1 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

2 Install the language packages using the `install_lp` command.

```
# cd /cdrom/cdrom0
# ./install_lp
```

Chapter 7

Installing Storage Foundation using the web-based installer

This chapter includes the following topics:

- About the web-based installer
- Before using the web-based installer
- Starting the web-based installer
- Obtaining a security exception on Mozilla Firefox
- Performing a preinstallation check with the web-based installer
- Installing SF with the web-based installer

About the web-based installer

Use the web-based installer interface to install Symantec 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 `xprtlwid` process. The `webinstaller` script can also be used to check the status of the XPortal Server.

When the `webinstaller` script starts the `xprtlwid` 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 operates. While the installation processes operate, 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. Symantec recommends 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 web-based installer" on page 52.

See "Starting the web-based installer" on page 52.

Before using the web-based installer

The 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 Symantec products.	Must be a supported platform for Storage Foundation 6.1. See "Supported upgrade paths for SF 6.1" on page 113.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	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 web-based installer

This section describes starting the web-based installer.

To start the web-based installer

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

```
# ./webinstaller start
```

The `webinstaller` script displays a URL. Note this URL.

Note: If you do not see the URL, please check your firewall and iptables settings. If you have configured a firewall, ensure that the firewall settings allow access to the port 14172. You can alternatively use the `-port` option to use a free port instead.

You can use the following command to display the details about ports used by webinstaller and its status:

```
# ./webinstaller status
```

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 you are 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 **I Understand the Risks**, or **You can add an exception**.

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 preinstallation check with the web-based installer

This section describes performing a preinstallation check with the web-based installer.

To perform a preinstallation check

1 Start the web-based installer.

See "Starting the web-based installer" on page 52.

2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.

3 Select **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 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.

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

Installing SF with the web-based installer

This section describes installing SF with the Symantec web-based installer.

To install SF using the web-based installer

1 Perform preliminary steps.

See "Performing a preinstallation check with the web-based installer" on page 54.

2 Start the web-based installer.

See "Starting the web-based installer" on page 52.

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 packages. 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 rsh. 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 radio buttons:

- Enable keyless licensing and complete system licensing later

  Note: The keyless license option enables you to install without entering a key. However, 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

  Click **Next**

  Complete the following information:

  - Choose whether you want to install Standard or Enterprise mode.
  - Choose whether you want to enable the Symantec Volume Replicator.
  - Click **Next**.

- Enter a valid license key
  If you have a valid license key, input the license key and click **Next**.

11 For Storage Foundation, click **Next**. If the installer prompts you to restart the system, then restart the system and invoke the web-based installer again for configuration. If the installer does not prompt you to restart the system, then it completes the configuration and starts 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 Symantec Storage Foundation, start the web-based installer and select **Configure a product.** Click the **OK** button. The installers checks for updates. Click the **Next** button.

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

Click the **Finish** button.

13 If you are prompted, enter the option to specify whether you want to send your installation information to Symantec.

```
Installation procedures and diagnostic information were saved in
the log files under directory
/var/tmp/installer-<platform>-<uuid>. Analyzing this information
helps Symantec discover and fix failed operations performed by
the installer. Would you like to send the information about this
installation to Symantec to help improve installation in the
future? [y,n,q,?]
```

Click **Finish**. The installer asks if you want to read the summary file. Select **Yes** to read the summary file. If you select **No**, the installer prompts you for another task.

Chapter 8

Automated installation using response files

This chapter includes the following topics:

- Installing SF using response files
- Response file variables to install Storage Foundation
- Sample response file for SF installation
- Configuring SF using response files
- Response file variables to configure Storage Foundation

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 that 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.

See "About the script-based installer" on page 46.

7 Complete the SF post-installation tasks.

For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.

Response file variables to install Storage Foundation

Table 8-1 lists the response file variables that you can define to install SF.

Table 8-1 Response file variables for installing SF

Variable	Description
CFG{opt}{install}	Installs SF packages. Configuration can be performed at a later time using the `-configure` option. List or scalar: scalar Optional or required: optional
CFG{opt}{installallpkgs} or CFG{opt}{installrecpkgs} or CFG{opt}{installminpkgs}	Instructs the installer to install SF packages based on the variable that has the value set to 1: ■ installallpkgs: Installs all packages ■ installrecpkgs: Installs recommended packages ■ installminpkgs: Installs minimum packages **Note:** Set only one of these variable values to 1. In addition to setting the value of one of these variables, you must set the variable `CFG{opt}{install}` to 1. List or scalar: scalar Optional or required: required

Table 8-1 Response file variables for installing SF *(continued)*

Variable	Description
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{opt}{license}	Installs the product with permanent license. List or scalar: scalar Optional or required: optional
CFG{keys}{hostname}	List of keys to be registered on the system if the variable `CFG{opt}{vxkeyless}` is set to 0 or if the variable `CFG{opt}{licence}` is set to 1. 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

Table 8-1 Response file variables for installing SF *(continued)*

Variable	Description
CFG{opt}{pkgpath}	Defines a location, typically an NFS mount, from which all remote systems can install product packages. The location must be accessible from all target systems. List or scalar: scalar Optional or required: optional
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the packages that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{updatekeys}	Updates the keyless license to the current version. List or scalar: scalar Optional or required: optional
CFG{opt}{rsh}	Defines that *rsh* must be used instead of ssh as the communication method between systems. List or scalar: scalar 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

Sample response file for SF installation

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

```
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{configure}=1;
```

```
$CFG{opt}{install}=1;
$CFG{opt}{installrecpkgs}=1;
$CFG{opt}{prodmode}="SF Enterprise";
$CFG{opt}{updatekeys}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{prod}="SF61";
$CFG{systems}=[ qw(cdgv240a) ];

1;
```

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 packages 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. Copy the response file to one of the cluster systems where you want to configure SF.
4. 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 62.
5. Start the configuration from the system to which you copied the response file. For example:

   ```
   # /opt/VRTS/install/installsf<version>
   -responsefile /tmp/response_file
   ```

   Where *<version>* is the specific release version, and */tmp/response_file* is the response file's full path name.

   See "About the script-based installer" on page 46.

Response file variables to configure Storage Foundation

Table 8-2 lists the response file variables that you can define to configure SF.

Table 8-2 Response file variables specific to configuring Storage Foundation

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the packages 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)
CFG{prod}	Scalar	Defines the product to be configured. The value is SF61 for SF. (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 *rsh* must be used instead of ssh as the communication method between systems. (Optional)

Table 8-2 Response file variables specific to configuring Storage Foundation *(continued)*

Variable	List or Scalar	Description
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 *logpath* 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 website. The value 0 indicates that the installation logs are not uploaded to the Symantec website. (Optional)

Chapter 9

Installing Storage Foundation using operating system-specific methods

This chapter includes the following topics:

- About installing SF using operating system-specific methods
- Installing SF on Solaris 11 using Automated Installer
- Installing SF on Solaris 10 using JumpStart
- Manually installing SF using the system command

About installing SF using operating system-specific methods

On Solaris, you can install SF using the following methods:

- You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Storage Foundation product on multiple client systems in a network.
  See "Installing SF on Solaris 11 using Automated Installer" on page 65.
- The procedure to manually install SF differs depending on the Solaris version.
  See "Manually installing SF using the system command" on page 76.
- You can install SF on Solaris 10 systems using Solaris JumpStart.

See "Installing SF on Solaris 10 using JumpStart" on page 70.

- You can install SF using Flash archive on the Solaris 10 operating system.
  See "Using a Flash archive to install SF and the operating system" on page 74.

Installing SF on Solaris 11 using Automated Installer

You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Storage Foundation product on multiple client systems in a network. AI performs a hands-free installation (automated installation without manual interactions) of SPARC systems. You can also use AI media to install the Oracle Solaris OS on a single SPARC platform. Oracle provides the AI bootable image and it can be downloaded from the Oracle website. All cases require access to a package repository on the network to complete the installation.

About Automated Installation

AI automates the installation of the Oracle Solaris 11 OS on one or more SPARC clients in a network. Automated Installation applies to Solaris 11 only. You can install the Oracle Solaris OS on many different types of clients. The clients can differ in:

- Architecture
- Memory characteristics
- MAC address
- IP address
- CPU

The installations can differ depending on specifications including network configuration and packages installed.

An automated installation of a client in a local network consists of the following high-level steps:

1. A client system boots and gets IP information from the DHCP server
2. Characteristics of the client determine which AI service and which installation instructions are used to install the client.
3. The installer uses the AI service instructions to pull the correct packages from the package repositories and install the Oracle Solaris OS on the client.

Using Automated Installer

To use Automated Installer to install systems over the network, set up DHCP and set up an AI service on an AI server. The DHCP server and AI server can be the same system or two different systems.

Make sure that the systems can access an Oracle Solaris Image Packaging System (IPS) package repository. The IPS package repository can reside on the AI server, on another server on the local network, or on the Internet.

An AI service is associated with a SPARC AI install image and one or more sets of installation instructions. The installation instructions specify one or more IPS package repositories from where the system retrieves the packages that are needed to complete the installation. The installation instructions also include the names of additional packages to install and information such as target device and partition information. You can also specify instructions for post-installation configuration of the system.

Consider the operating systems and packages you want to install on the systems. Depending on your configuration and needs, you may want to do one of the following:

- If two systems have different architectures or need to be installed with different versions of the Oracle Solaris OS, create two AI services. Then, associate each AI service with a different AI image
- If two systems need to be installed with the same version of the Oracle Solaris OS but need to be installed differently in other ways, create two sets of installation instructions for the AI service. The different installation instructions can specify different packages to install or a different slice as the install target.

The installation begins when you boot the system. DHCP directs the system to the AI install server, and the system accesses the install service and the installation instructions within that service.

For more information, see the *Oracle® Solaris 11 Express Automated Installer Guide*.

Using AI to install the Solaris 11 operating system and SFHA products

Use the following procedure to install the Solaris 11 operating system and SFHA products using AI.

To use AI to install the Solaris 11 operating system and SFHA products

1 Follow the Oracle documentation to set up a Solaris AI server and DHCP server.

You can find the documentation at http://docs.oracle.com.

2 Set up the Symantec package repository.

Run the following commands to startup necessary SMF services and create directories:

```
# svcadm enable svc:/network/dns/multicast:default
# mkdir /ai
# zfs create -o compression=on -o mountpoint=/ai rpool/ai
```

3 Run the following commands to set up IPS repository for Symantec SPARC packages:

```
# mkdir -p /ai/repo_symc_sparc
# pkgrepo create /ai/repo_symc_sparc
# pkgrepo add-publisher -s /ai/repo_symc_sparc Symantec
# pkgrecv -s <media_sparc>/pkgs/VRTSpkgs.p5p -d
/ai/repo_symc_sparc '*'
# svccfg -s pkg/server list
# svcs -a | grep pkg/server
# svccfg -s pkg/server add symcsparc
# svccfg -s pkg/server:symcsparc addpg pkg application
# svccfg -s pkg/server:symcsparc setprop pkg/port=10003
# svccfg -s pkg/server:symcsparc setprop pkg/inst_root=
/ai/repo_symc_sparc
# svccfg -s pkg/server:symcsparc addpg general framework
# svccfg -s pkg/server:symcsparc addpropvalue general/complete
astring: symcsparc
# svccfg -s pkg/server:symcsparc addpropvalue general/enable
boolean: true
# svcs -a | grep pkg/server
# svcadm refresh application/pkg/server:symcsparc
# svcadm enable application/pkg/server:symcsparc
```

Or run the following commands to set up the private depot server for testing purposes:

```
# /usr/lib/pkg.depotd -d /ai/repo_symc_sparc -p 10003 > /dev/null &
```

Check the following URL on IE or Firefox browser:

http://<host>:10003

4 Set up the install service on the AI server.

Run the following command:

```
# mkdir /ai/iso
```

Download the AI image from the Oracle website and place the `iso` in the `/ai/iso` directory.

Create an install service.

For example:

To set up the AI install service for SPARC platform::

```
# # installadm create-service -n sol11sparc -s\
 /ai/iso/sol-11-1111-ai-sparc.iso -d /ai/aiboot/
```

5 Run the installer to generate manifest XML files for all the SFHA products that you plan to install.

```
# mkdir /ai/manifests
# <media>/installer -ai /ai/manifests
```

6 For each system, generate the system configuration and include the host name, user accounts, and IP addresses. For example, enter one of the following:

```
# mkdir /ai/profiles
# sysconfig create-profile -o /ai/profiles/profile_client.xml
```

or

```
# cp /ai/aiboot/auto-install/sc_profiles/sc_sample.xml
/ai/profiles/profile_client.xml
```

7 Add a system and match it to the specified product manifest and system configuration.

Run the following command to add a SPARC system, for example:

```
# installadm create-client -e "<client_MAC>" -n sol11sparc
# installadm add-manifest -n sol11sparc -f \
/ai/manifests/vrts_manifest_sfha.xml
# installadm create-profile -n sol11sparc -f \
/ai/profiles/profile_client.xml -p profile_sc
# installadm set-criteria -n sol11sparc -m \
vrts_sfha -p profile_sc -c mac="<client_MAC>"
# installadm list -m -c -p -n sol11sparc
```

8 For SPARC system, run the following command to restart the system and install the operating system and Storage Foundation products:

```
# boot net:dhcp - install
```

Installing SF on Solaris 10 using JumpStart

This installation method applies only to Solaris 10. These JumpStart instructions assume a working knowledge of JumpStart. See the JumpStart documentation that came with your operating system for details on using JumpStart.

Upgrading is not supported. The following procedure assumes a standalone configuration.

For the language pack, you can use JumpStart to install packages. You add the language packages in the script, and put those files in the JumpStart server directory.

You can use a Flash archive to install SF and the operating system with JumpStart.

See "Using a Flash archive to install SF and the operating system" on page 74.

Overview of JumpStart installation tasks

Review the summary of tasks before you perform the JumpStart installation.

Summary of tasks

1 Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.

2 Read the JumpStart installation instructions.

3 Generate the finish scripts.

See “Generating the finish scripts” on page 71.

4 Prepare shared storage installation resources.

See “Preparing installation resources” on page 72.

5 Modify the rules file for JumpStart.

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

6 Install the operating system using the JumpStart server.

7 When the system is up and running, run the installer command from the installation media to configure the Symantec software.

```
# /opt/VRTS/install/installer -configure
```

See “About the script-based installer” on page 46.

Generating the finish scripts

Perform these steps to generate the finish scripts to install SF.

To generate the script

1 Run the product installer program to generate the scripts for all products.

```
./installer -jumpstart directory_to_generate_scripts
```

Or

```
./install<productname> -jumpstart directory_to_generate_script
```

where **<*productname*>** is the product's installation command, and *directory_to_generate_scripts* is where you want to put the product's script.

For example:

```
# ./installsf -jumpstart /js_scripts
```

2 When you are prompted to encapsulate the root disk automatically, choose **yes** to do so. If you do not want to encapsulate it automatically, choose **no** and go to step 6.

3 Specify a disk group name for the root disk.

```
Specify the disk group name of the root disk to be encapsulated:
rootdg
```

4 Specify private region length.

```
Specify the private region length of the root disk to be
encapsulated: (65536)
```

5 Specify the disk's media name of the root disk to encapsulate.

```
Specify the disk media name of the root disk to be encapsulated:
(rootdg_01)
```

6 JumpStart finish scripts and encapsulation scripts are generated in the directory you specified in step 1.

Output resembles:

```
The finish scripts for SF is generated at /js_scripts/
jumpstart_sf.fin
The encapsulation boot disk script for VM is generated at
/js_scripts/encap_bootdisk_vm.fin
```

List the js_scripts directory.

```
# ls /js_scripts
```

Output resembles:

```
encap_bootdisk_vm.fin jumpstart_sf.fin
```

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1 Copy the `pkgs` directory of the installation media to the shared storage.

```
# cd /path_to_installation_media
# cp -r pkgs BUILDSRC
```

2 Generate the response file with the list of packages.

```
# cd BUILDSRC/pkgs/
# pkgask -r package_name.response -d /
BUILDSRC/pkgs/packages_name.pkg
```

3 Create the `adminfile` file under *`BUILDSRC`*`/pkgs/` directory.

```
mail=
instance=overwrite
partial=nocheck
runlevel=quit
idepend=quit
rdepend=nocheck
space=quit
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=default
```

4 If you want to encapsulate the root disk automatically when you perform the JumpStart installation, copy the scripts `encap_bootdisk_vm.fin` generated previously to *ENCAPSRC*.

See "Generating the finish scripts" on page 71.

Adding language pack information to the finish file

To add the language pack information to the finish file, perform the following procedure.

To add the language pack information to the finish file

1. For the language pack, copy the language packages from the language pack installation disc to the shared storage.

   ```
   # cd /cdrom/cdrom0/pkgs
   # cp -r * BUILDSRC/pkgs
   ```

   If you downloaded the language pack:

   ```
   # cd /path_to_language_pack_installation_media/pkgs
   # cp -r * BUILDSRC/pkgs
   ```

2. In the finish script, copy the product package information and replace the product packages with language packages.
3. The finish script resembles:

   ```
   . . .
   for PKG in product_packages
   do
   ...
   done. . .
   for PKG in language_packages
   do
   ...
   done. . .
   ```

Using a Flash archive to install SF and the operating system

You can only use Flash archive on the Solaris 10 operating system. In the following outline, refer to Solaris documentation for Solaris-specific tasks.

Note: Symantec does not support Flash Archive installation if the root disk of the master system is encapsulated.

The following is an overview of the creation and installation of a Flash archive with Symantec software.

- If you plan to start flar (flash archive) creation from bare metal, perform step 1 through step 10.
- If you plan to start flar creation from a system where you have installed, but not configured the product, perform step 1 through step 4. Skip step 5 and finish step 6 through step 10.

- If you plan to start flar creation from a system where you have installed and configured the product, perform step 5 through step 10.

Flash archive creation overview

1. Ensure that you have installed Solaris 10 on the master system.
2. Use JumpStart to create a clone of a system.
3. Restart the cloned system.
4. Install the Symantec products on the master system.

   Perform one of the installation procedures from this guide.
5. If you have configured the product on the master system, create the `vrts_deployment.sh` file and the `vrts_deployment.cf` file and copy them to the master system.

   See "Creating the Symantec post-deployment scripts" on page 75.
6. Use the `flarcreate` command to create the Flash archive on the master system.
7. Copy the archive back to the JumpStart server.
8. Use JumpStart to install the Flash archive to the selected systems.
9. Configure the Symantec product on all nodes in the cluster.

   The scripts that are installed on the system include the product version in the script name. For example, to install the SF script from the install media, run the `installsf` command. However, to run the script from the installed binaries, run the `installsf<version>` command. For example, for the 6.1 version:

   ```
   # /opt/VRTS/install/installsf61 -configure
   ```

   See "About the script-based installer" on page 46.
10. Perform post-installation and configuration tasks.

    See the product installation guide for the post-installation and configuration tasks.

Creating the Symantec post-deployment scripts

The generated files vrts_deployment.sh and vrts_post-deployment.cf are customized Flash archive post-deployment scripts. These files clean up Symantec product settings on a cloned system before you reboot it for the first time. Include these files in your Flash archives.

To create the post-deployment scripts

1. Mount the product disc.
2. From the prompt, run the `-flash_archive` option for the installer. Specify a directory where you want to create the files.

   ```
   # ./installer -flash_archive /tmp
   ```

3. Copy the vrts_postedeployment.sh file and the vrts_postedeployment.cf file to the golden system.
4. On the golden system perform the following:
   - Put the vrts_postdeployment.sh file in the /etc/flash/postdeployment directory.
   - Put the vrts_postdeployment.cf file in the /etc/vx directory.
5. Make sure that the two files have the following ownership and permissions:

   ```
   # chown root:root /etc/flash/postdeployment/vrts_postdeployment.sh
   # chmod 755 /etc/flash/postdeployment/vrts_postdeployment.sh
   # chown root:root /etc/vx/vrts_postdeployment.cf
   # chmod 644 /etc/vx/vrts_postdeployment.cf
   ```

   Note that you only need these files in a Flash archive where you have installed Symantec products.

Manually installing SF using the system command

The procedure to manually install SF differs depending on the Solaris version.

See "Installing SF on Solaris 10 using the pkgadd command" on page 76.

See "Installing SF on Solaris 11 using the pkg install command" on page 78.

Installing SF on Solaris 10 using the pkgadd command

On Solaris 10, the packages must be installed while in the global zone.

To install SF on Solaris 10 using the pkgadd command

1. Mount the software disc.

   See "Mounting the product disc" on page 42.

2. Copy the supplied VRTS* files from the installation media to a temporary location. Modify them if needed.

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

3. Create the admin file in the current directory. Specify the `-a adminfile` option when you use the `pkgadd` command:

   ```
   mail=
   instance=overwrite
   partial=nocheck
    runlevel=quit
   idepend=quit
   rdepend=nocheck
   space=quit
   setuid=nocheck
   conflict=nocheck
   action=nocheck
   basedir=default
   ```

4. Use the product-specific install command with one of the following options to get a list of packages in the order to be installed:

   - minpkgs
   - recpkgs
   - allpkgs

   See "About the script-based installer" on page 46.

   See "Installation script options" on page 229.

5. Install the packages listed in step 4.

   ```
   # pkgadd -a adminfile -d /tmp/pkgs pkgname.pkg
   ```

   On Solaris 10, these packages must be installed while in the global zone. If a package's `pkginfo` file contains the variable SUNW_PKG_ALLZONES set not equal to true, the `-G` option should additionally be specified to the `pkgadd` command.

6 Verify that the packages are installed:

```
# pkginfo -l
     packagename
```

7 Start the processes.

See "Starting and stopping processes for the Symantec products " on page 203.

Installing SF on Solaris 11 using the pkg install command

To install SF on Solaris 11 using the pkg install command

1 Copy the VRTSpkgs.p5p package from the pkgs directory from the installation media to the system at `/tmp/install` directory.

2 Disable the publishers that are not reachable as package install may fail if any of the already added repositories are unreachable.

```
# pkg set-publisher --disable <publisher name>
```

3 Add a file-based repository in the system.

```
# pkg set-publisher -p /tmp/install/VRTSpkgs.p5p Symantec
```

4 Install the required packages.

5 Remove the publisher from the system.

```
# pkg unset-publisher Symantec
```

6 Clear the state of the SMF service if non-global zones are present in the system. In presence of non-global zones, setting the file-based repository causes SMF service `svc:/application/pkg/system-repository:default` to go into maintenance state.

```
# svcadm clear svc:/application/pkg/system-repository:default
```

7 Enable the publishers that were disabled earlier.

```
# pkg set-publisher --enable <publisher name>
```

Manually installing packages on Oracle Solaris 11 systems

To install packages on Solaris 11 system

1. Copy the `VRTSpkgs.p5p` package from the `pkgs` directory from the installation media to the the system at `/tmp/install directory`..
2. Disable the publishers that are not reachable, as package install may fail if any of the already added repositories are unreachable.

   ```
   # pkg set-publisher --disable <publisher name>
   ```

3. Add a file-based repository in the non-global zone.

   ```
   # pkg set-publisher -p/tmp/install/VRTSpkgs.p5p Symantec
   ```

4. Install the required packages.
5. To configure an OracleVMServer logical domain for disaster recovery, install the following required packages inside the logical domain:

   ```
   # pkg install --accept VRTSvcsnr
   ```

6. Remove the publisher on the non-global zone.

   ```
   # pkg unset-publisher Symantec
   ```

7. Clear the state of the SMF service if non-global zones are present in the system. In presence of non-global zones, setting the file-based repository causes SMF service `svc:/application/pkg/system-repository:default` to go into maintenance state. .

   ```
   # svcadm clear svc:/application/pkg/system-repository:default
   ```

8. Enable the publishers that were disabled earlier.

   ```
   # pkg set-publisher --enable <publisher>
   ```

Manually installing packages on Solaris brand non-global zones

With Oracle Solaris 11, you must manually install SF packages inside non-global zones. The native non-global zones are called Solaris brand zones.

To install packages manually on Solaris brand non-global zones

1. Ensure that the SMF service `svc:/application/pkg/system-repository:default` and `svc:/application/pkg/zones-proxyd:default` is online on the global zone.

   ```
   # svcs svc:/application/pkg/system-repository:default
   # svcs  svc:/application/pkg/zones-proxyd:default
   ```

2. Log on to the non-global zone as a superuser.

3. Ensure that the SMF service `svc:/application/pkg/zones-proxy-client:default` is online inside non-global zone

   ```
   # svcs svc:/application/pkg/zones-proxy-client:default
   ```

4. Copy the `VRTSpkgs.p5p` package from the `pkgs` directory from the installation media to the non-global zone (for example at `/tmp/install` directory).

5. Disable the publishers that are not reachable, as package install may fail if any of the already added repositories are unreachable.

   ```
   # pkg set-publisher --disable <publisher name>
   ```

6. Add a file-based repository in the non-global zone.

   ```
   # pkg set-publisher -g/tmp/install/VRTSpkgs.p5p Symantec
   ```

7. Install the required packages.

   ```
   # pkg install --accept VRTSperl VRTSvlic VRTSvxfs VRTSvcs
   VRTSvcsag VRTSvcsea
   ```

8. Remove the publisher on the non-global zone.

   ```
   # pkg unset-publisher Symantec
   ```

9 Clear the state of the SMF service, as setting the file-based repository causes SMF service `svc:/application/pkg/system-repository:default` to go into maintenance state.

```
# svcadm clear svc:/application/pkg/system-repository:default
```

10 Enable the publishers that were disabled earlier.

```
# pkg set-publisher --enable <publisher>
```

Note: Note: Perform steps 2 through 10 on each non-global zone.

Chapter 10

Configuring Storage Foundation

This chapter includes the following topics:

- Configuring Storage Foundation using the installer
- Configuring Storage Foundation manually
- 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

1. Go to the installation directory.
2. Run the installer command with the configure option.

```
# ./installer -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 *Symantec Storage Foundation Administrator's Guide*.

In releases of VxVM (Volume Manager) before 4.0, a system that was installed with VxVM was configured with a default disk group, `rootdg`. The `rootdg` disk group had to contain at least one disk. By default, operations were directed to the `rootdg` disk group. From release 4.0 onward, VxVM can function without any disk group having been configured.

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:

```
# vxdctl 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 `vxconfigd` daemon, enter the following command:

```
# vxdctl enable
```

Once started, `vxconfigd` automatically becomes a background process.

By default, `vxconfigd` writes error messages to the console. However, you can configure it to write errors to a log file. For more information, see the `vxconfigd`(1M) and `vxdctl`(1M) manual pages.

Starting the volume I/O daemon

The volume I/O daemon (`vxiod`) provides extended I/O operations without blocking calling processes. Several `vxiod` daemons are usually started at system boot time after initial installation, and they should be running at all times. The procedure below describes how to verify that the `vxiod` daemons are running, and how to start them if necessary.

To verify that `vxiod` daemons are running, enter the following command:

```
# vxiod
```

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 to16. 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.

Using vxinstall to configure Veritas Volume Manager

If you used the Veritas Installation Menu or the `installvm` script, you do not need to carry out the instructions in this section. Licensing, configuration of enclosure based naming and creation of a default disk group are managed by the menu installer and the `installvm` script.

Because you are no longer required to configure VxVM disks immediately, the `vxinstall` command no longer invokes the `vxdiskadm` program, making it much simpler than in previous releases.

The utility provides the following functions:

- Licensing VxVM.
- Setting up a system-wide default disk group.
- Starting VxVM daemons in case installation of SF has been done manually.

To run the command, enter

```
# vxinstall
```

which will prompt you to enter a license key:

```
Are you prepared to enter a license key [y,n,q,?] (default: y) y
```

The `vxinstall` program then asks if you want to set up a system-wide default disk group, which is optional:

```
Do you want to setup a system wide default disk group ?
[y,n,q,?] (default: y)
```

VxVM will continue with the question:

```
Which disk group [<group>,list,q,?] ?
```

If you know the name of the disk group that you want to use as the default disk group, enter it at the prompt, or use the `list` option and make a selection.

In releases prior to VxVM 4.0, the default disk group was `rootdg` (the root disk group). For VxVM to function, the `rootdg` disk group had to exist and it had to contain at least one disk. This requirement no longer exists, however you may find it convenient to create a system-wide default disk group. For operations that require a disk group, the system-wide default disk group will be used if the VxVM command is not specified with the `-g` option. The main benefit of creating a default disk group is that VxVM commands default to the default disk group and you will not need to use the `-g` option. To verify the default disk group after it has been created, enter the command:

```
# vxdg defaultdg
```

VxVM does not allow you to use the following names for the default disk group because they are reserved words: `bootdg`, `defaultdg` and `nodg`.

At this stage, the installation of VxVM is complete. To carry out further tasks such as disk encapsulation or initialization, see the *Symantec Storage Foundation Administrator's Guide*.

Configuring Veritas File System

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

`/etc/vfstab`

The Symantec-specific commands are described in the Symantec Storage Foundation guides and online manual pages.

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

Loading and unloading the file system module

The `vxfs` file system module automatically loads on the first reference to a VxFS file system. This occurs when a user tries to mount a VxFS disk layout. In some instances, you may want to load the file system module manually. To do this, first load `vxfs`, then `vxportal`. `vxportal` is a pseudo device driver that enables VxFS commands to issue ioctls to the VxFS modules even when there are no file systems mounted on the system.

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

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

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

To determine if the modules successfully loaded, enter:

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

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

You can unload the module by entering:

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

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

```
# df -F vxfs
```

vxtunefs command permissions and Cached Quick I/O

By default, you must have superuser (`root`) privileges to use the `/opt/VRTS/bin/vxtunefs` command. The `vxtunefs` command is a tool that lets you change caching policies to enable Cached Quick I/O and change other file system options. Database administrators can be granted permission to change default file system behavior in order to enable and disable Cached Quick I/O. The system administrator must change the `vxtunefs` executable permissions as follows:

```
# chown root /opt/VRTS/bin/vxtunefs
# chgrp dba /opt/VRTS/bin/vxtunefs
# chmod 4550 /opt/VRTS/bin/vxtunefs
```

Setting the permissions for `/opt/VRTS/bin/vxtunefs` to 4550 allows all users in the dba group to use the `vxtunefs` command to modify caching behavior for Quick I/O files.

For more information, see the *Symantec Storage Foundation Administrator's Guide*.

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

Section 3

Managing your Symantec deployments

- Chapter 11. Performing centralized installations using the Deployment Server

Chapter 11

Performing centralized installations using the Deployment Server

This chapter includes the following topics:

- About the Deployment Server
- How to install the Deployment Script
- Deployment management overview
- Setting up a Deployment Server
- Setting deployment preferences
- Using the Deployment Server command line option to specify a non-default repository location
- Using the Deployment Server command line options to load and download the most recent release information
- Viewing or downloading available release images
- Viewing or removing repository images stored in your repository
- Deploying Symantec product updates to your environment
- Finding out which releases you have, and which upgrades or updates you may need
- Deploying a specific Symantec release
- Updating release information on systems without Internet access

About the Deployment Server

The Deployment Server is a script that makes it easier to install or upgrade SFHA releases. The Deployment Server lets you store multiple release images in one central location and deploy them to systems of any supported UNIX or Linux operating system (6.1 or later). Prior to 6.1, releases still require the same platform, architecture, distribution, and version of the operating system. You can use the Deployment Server if you want to install or upgrade multiple releases and or multiple platforms.

The Deployment Server lets you do the following as described in Table 11-1.

Table 11-1 Deployment Server functionality

Feature	Description
Manage release images	■ View available SFHA releases. ■ Download maintenance and hot fix release images from the Symantec Operations Readiness Tools (SORT) website into a repository. ■ Load the downloaded release image files from FileConnect and SORT into the repository. ■ View and remove release image files stored in the repository.
Check versions	■ Discover packages and patches installed on designated systems and informs you of the product and version installed, including installed hot fixes. ■ Identify base, maintenance, and hot fix level upgrades to your system and download maintenance and hot fix releases. ■ Query SORT for the most recent updates.

Install or upgrade systems	■ Install or upgrade a release stored in the repository on selected systems.
	■ In release 6.1 and later:
	  ■ Install hot fix level releases.
	  ■ Install SFHA from any supported UNIX or Linux operating system to any other supported UNIX or Linux operating system.
	  ■ Automatically load the script-based installer hot fixes that apply to that release.

Update metadata and preferences	■ Download, load the release matrix updates, and script-based installer updates for systems behind a firewall. ■ Define or reset program settings.

Note: The Deployment Server is available only for the script-based installer, not the web-based installer.

How to install the Deployment Script

The Deployment Script is the utility for managing your Deployment Server.

You can obtain the Deployment Script by either:

- Installing the Deployment Script manually.
- Running the Deployment Script after installing at least one Symantec 6.1 product.

Note: The `VRTSperl` and the `VRTSsfcpi<version>`packages are included in all Storage Foundation (SF) products, so installing any Symantec 6.1 product lets you access the Deployment Script.

To install the Deployment Script manually

1. Log in as superuser.
2. Mount the installation media.

   See "Mounting the product disc" on page 42.
3. For Solaris 10, move to the top-level directory on the disc.

   ```
   # cd /cdrom/cdrom0
   ```
4. For Solaris 10, navigate to the following directory:

   ```
   # cd pkgs
   ```
5. For Solaris 10, run the following commands to install the `VRTSperl` and the `VRTSsfcpi<version>` packages:

   ```
   # pkgadd -d ./VRTSperl.pkg VRTSperl
   # pkgadd -d ./VRTSsfcpi<version>.pkg VRTSsfcpi<version>
   ```
6. For Solaris 11, move to the top-level directory on the disc.

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

7 For Solaris 11, navigate to the following directory:

```
# cd pkgs
```

8 For Solaris 11, run the following commands to install the `VRTSperl` and the `VRTSsfcpi<version>` packages:

```
# pkg install --accept -g ./VRTSpkgs.p5p VRTSperl VRTSsfcpi<version>
```

To run the Deployment Script

1 Log in as superuser.

2 Navigate to the following directory:

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

3 Run the Deployment Script.

```
# ./deploy_sfha
```

Deployment management overview

After obtaining and installing the Deployment Server and defining a central repository, you can begin managing your deployments from that repository. You can load and store product images for Symantec products up to version 5.1 in your Deployment Server. The Deployment Server is a central installation server for storing and managing your product updates.

See "How to install the Deployment Script" on page 91.

Setting up and managing your repository involves the following tasks:

- Setting up a Deployment Server.
  See "Setting up a Deployment Server" on page 93.
- Finding out which products you have installed, and which upgrades or updates you may need.
  See "Viewing or downloading available release images" on page 98.
- Adding release images to your Deployment Server.
  See "Viewing or downloading available release images" on page 98.
- Removing release images from your Deployment Server.
  See "Viewing or removing repository images stored in your repository" on page 102.

Later, when your repository is set up, you can use it to deploy Symantec products to other systems in your environment.

See "Deploying Symantec product updates to your environment" on page 105.

Setting up a Deployment Server

For large deployments, Symantec recommends that you create a dedicated Deployment Server to manage your product updates.

A Deployment Server is useful for doing the following tasks:

- Downloading and storing release images for the latest upgrades and updates from Symantec in a central repository directory.
- Installing and updating systems directly by accessing the release images that are stored within a central repository (direct installation).
- Performing heterogeneous push installations (installing Symantec products from the Deployment Server to systems running any supported platform).

  Note: The script-based installer for version 6.1 and higher supports installations from one operating system node onto a different operating system. Therefore, heterogeneous push installations are supported for 6.1 and higher releases only. To perform push installations for product versions 5.1, 6.0, or 6.0.1 releases, you must have a separate Deployment Server for each operating system.

- Creating a file share on the repository directory provides a convenient, central location from which systems running any supported platform can install the latest Symantec products and updates.

Create a central repository on the Deployment Server to store and manage the following types of Symantec releases:

- Base releases. These major releases and minor releases are available for all Symantec products. They contain new features, and you can download them from FileConnect.
- Maintenance releases. These releases are available for all Symantec products. They contain bug fixes and a limited number of new features, and you can download them from the Symantec Operations Readiness Tools (SORT) website.
- Hot fixes. These releases contain fixes for specific products, and you can download them from the SORT website.

Note: All SF base releases and maintenance releases can be deployed using the install scripts that are included in the release. Hot fixes are typically installed manually, however, from the 6.1 release and onwards, install scripts are included with hot fix releases.

You can set up a Deployment Server with or without Internet access.

- If you set up a Deployment Server that has Internet access, you can download SF maintenance releases and hot fixes from Symantec directly. Then, you can deploy them to your systems.
  Setting up a Deployment Server that has Internet access
- If you set up a Deployment Server that does not have Internet access, you can download SF maintenance releases and hot fixes from Symantec on another system that has Internet access. Then, you can load the images onto the Deployment Server separately.
  Setting up a Deployment Server that does not have Internet access

Setting up a Deployment Server that has Internet access

Figure 11-1 shows a Deployment Server that can download product images directly from Symantec using the Deployment Server.

Figure 11-1 Example Deployment Server that has Internet access

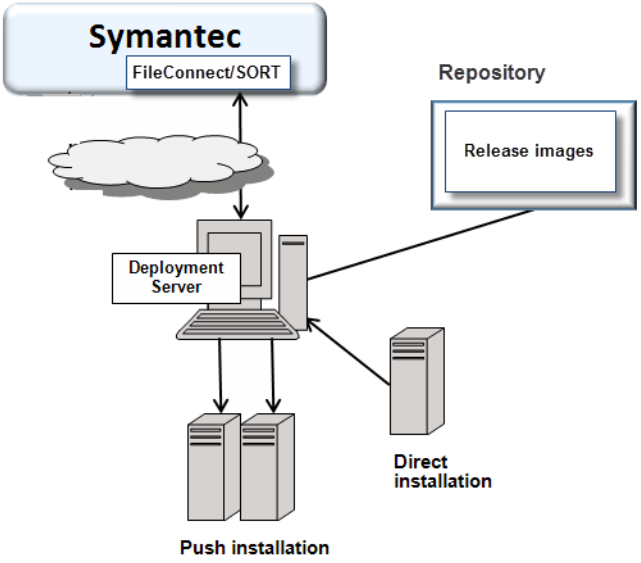

Setting up a Deployment Server that does not have Internet access

Figure 11-2 shows a Deployment Server that does not have Internet access. In this scenario, release images and metadata updates are downloaded from another system. Then, they are copied to a file location available to the Deployment Server, and loaded.

Figure 11-2 Example Deployment Server that does not have Internet access

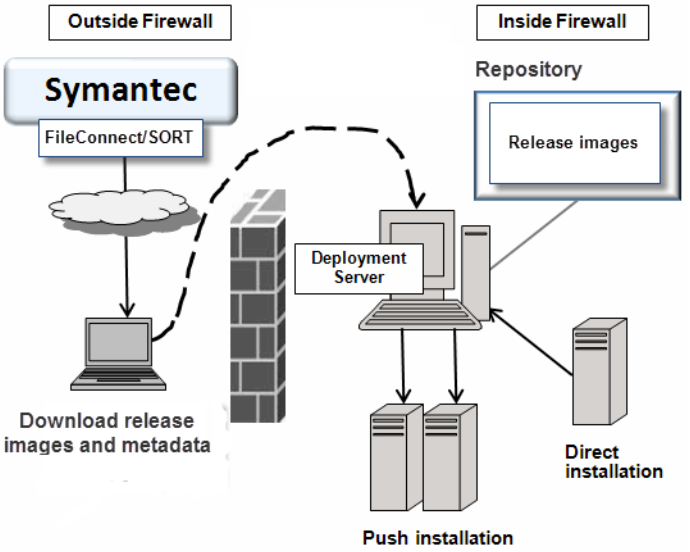

Setting deployment preferences

You can set preferences for managing the deployment of products dating back to version 5.1.

Note: You can select option **T (Terminology and Usage)** to obtain more information about Deployment Server terminology and usage.

To set deployment preferences

1. Launch the Deployment Server.

   ```
   # /opt/VRTS/install/deploy_sfha
   ```

   You see the following output:

   ```
   Task Menu:

   R) Manage Repository Images              M) Update Metadata
   V) Version Check Systems                 S) Set Preferences
   U) Upgrade/Install Systems               T) Terminology and Usage
   ?) Help                                  Q) Quit

   Enter a Task: [R,M,V,S,U,T,?,Q]
   ```

2. Select option **S**, **Set Preferences**.
3. In the current preferences page, select option **S**, **Set Preferences**.
4. Do one of the following:
   - To set the default repository, enter **1**. Then enter the name of the repository in which you want to store your downloads. For example, enter the following:

     ```
     # /opt/VRTS/install/ProductDownloads
     ```

     If the specified repository replaces a previous repository, the installer asks if you want to move all your files into the new repository. To move your files to the new repository, enter **y**.
   - To set the option for saving or removing tar files, enter **2**. At the prompt, if you want to save the tar files after untarring them, enter **y**. Or, if you want to remove tar files after untarring them, enter **n**.
     By default, the installer does not remove tar files after the releases have been untarred.

Using the Deployment Server command line option to specify a non-default repository location

You can specify a repository location other than the default that has been set within the system preferences by using the command line option. The command line option is mainly used to install a release image from a different repository location. When you use the command line option, the designated repository folder is used instead

of the default for the execution time of the script. Using the command line option does not override the repository preference set by the **Set Preference** menu item.

To use the command line option to specify a non-default repository location

- At the command line, to specify a non-default repository location, enter the following:

  ```
  # ./deploy_sfha -repository repository_path
  ```

  where *repository_path* is the preferred location of the repository.

Using the Deployment Server command line options to load and download the most recent release information

You can use the Deployment Server command line options to perform the following operations:

- Load the most recent SF release information and installer hot fixes on systems without Internet access using the `deploy_sfha` command.
- Download a `.tar` file containing the most recent SF release information and installer hot fixes from the SORT site. This `.tar` file is used to update release information on systems without Internet access using the `deploy_sfha -load_metadata` command.

To use the Deployment Server command line option to load the most recent SF release information and installer hot fixes without Internet access

- At the command line, enter the following:

  ```
  # ./deploy_sfha -load_metadata metadata_tar_file
  ```

  where *metadata_tar_file* is the name of the metadata tar file where you want to load the most recent SF release information and installer hot fixes.

  You can download the `.tar` file from the SORT site at:

  https://sort.symantec.com/support/related_links/offline-release-updates

  Or, you can create it by running the `deploy_sfha -download_metadata` command from a system that does have Internet access and can connect to the SORT site.

To use the Deployment Server command line option to download the most recent SF release information and installer hot fixes from the SORT site

- At the command line, enter the following:

```
# ./deploy_sfha -download_metadata
```

Viewing or downloading available release images

You can use the Deployment Server to conveniently view or download available SF release images to be deployed on other systems in your environment.

Note: If you have Internet access, communication with the Symantec Operations Readiness Tools (SORT) provides the latest release information. If you do not have Internet access, static release matrix files are referenced, and the most recent updates may not be included.

See "Updating release information on systems without Internet access" on page 109.

To view or download available release images

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images      M) Update Metadata
V) Version Check Systems         S) Set Preferences
U) Upgrade/Install Systems       T) Terminology and Usage
?) Help                          Q) Quit

Enter a Task: [R,M,V,S,U,T,?,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

Select the option you would like to perform [1-3,b,q,?]
```

3 Select option **1**, **View/Download Available Releases**, to view or download what is currently installed on your system.

You see a list of platforms and release levels.

```
To view or download available releases, the platform type and release
level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  SLES10 x86_64             8)  SLES11 x86_64
9)  Solaris 9 Sparc          10)  Solaris 10 Sparc
11)  Solaris 10 x64          12)  Solaris 11 Sparc
13)  Solaris 11 x64           b)  Back to previous menu

Select the platform of the release to view/download [1-13,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels.

```
1) Base
2) Maintenance
3) Hot Fix
b) Back to previous menu

Select the level of the <platform> releases to view/download
[1-3,b,q,?]
```

5 Select the number corresponding to the type of release you want to view (Base, Maintenance, or Hot Fix).

You see a list of releases available for download.

```
Available Maintenance releases for sol10_sparc:

release_version SORT_release_name           DL OBS AI rel_date   size_KB
================================================================================
5.1SP1PR2RP2  sfha-sol10_sparc-5.1SP1PR2RP2 -  Y   Y  2011-09-28 145611
5.1SP1PR2RP3  sfha-sol10_sparc-5.1SP1PR2RP3 -  Y   Y  2012-10-02 153924
5.1SP1PR2RP4  sfha-sol10_sparc-5.1SP1PR2RP4 -  -   -  2013-08-21 186859
5.1SP1PR3RP2  sfha-sol10_sparc-5.1SP1PR3RP2 -  Y   Y  2011-09-28 145611
5.1SP1PR3RP3  sfha-sol10_sparc-5.1SP1PR3RP3 -  Y   Y  2012-10-02 153924
5.1SP1PR3RP4  sfha-sol10_sparc-5.1SP1PR3RP4 -  -   -  2013-08-21 186859
6.0RP1        sfha-sol10_sparc-6.0RP1       Y  -   -  2012-03-22 245917
6.0.3         sfha-sol10_sparc-6.0.3        Y  -   -  2013-02-01 212507

Enter the release_version to view details about a release or press 'Enter' to continue [b,q,?]
```

The following are the descriptions for the column headers:

- release_version: The version of the release.
- SORT_release_name: The name of the release, used when accessing SORT (https://sort.symantec.com).
- DL: An indicator that the release is present in your repository.
- OBS: An indicator that the release has been obsoleted by another higher release.
- AI: An indicator that the release has scripted install capabilities. All base and maintenance releases have auto-install capabilities. Hot Fix releases

with auto-install capabilities are available beginning with version 6.1. Otherwise the hot fix will require a manual installation.

- rel_date: The date the release is available.
- size_KB: The file size of the release in kilobytes.

6 If you are interested in viewing more details about any release, type the release version. For example, enter the following:

```
6.0.3
```

You see the following output:

```
release_version: 6.0.3
release_name: sfha-sol10_sparc-6.0.3
release_type: MR
release_date: 2013-02-01
downloaded: Y
install_path: sol10_sparc/installmr
upload_location: ftp://ftp.veritas.com/pub/support/patchcentral
/Solaris/6.0.3/sfha/sfha-sol10_sparc-6.0.3-patches.tar.gz
obsoletes: 6.0.1.200-fs,6.0.1.200-vm,6.0.1.300-fs
obsoleted_by: None
Would you like to download this patch? [y,n,q] (y) n

Enter the release_version to view the details about a release or press
'Enter' to continue [b,q,?]
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to download a sol10_sparc Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all releases that are not currently in the repository.

```
1)  5.1SP1PR1RP2
2)  5.1SP1PR1RP4
3)  6.0RP1
4)  All non-obsoleted releases
5)  All releases
b)  Back to previous menu

Select the patch release to download, 'All non-obsoleted releases' to
download all non-obsoleted releases, or 'All releases' to download
all releases [1-5,b,q] 3
```

8 Select the number corresponding to the release that you want to download. It is possible to download a single release, all non-obsoleted releases, or all releases.

The selected release images are downloaded to the Deployment Server.

```
Downloading sfha-sol10_sparc-6.0RP1 from SORT - https://sort.symantec.com
Downloading 215118373 bytes (Total 215118373 bytes [205.15 MB]): 100%
Untarring sfha-sol10_sparc-6.0RP1 ..................................... Done

sfha-sol10_sparc-6.0RP1 has been downloaded successfully.
```

9 From the menu, select option **2**, **View/Remove Repository Images**, and follow the prompts to check that the release images are loaded.

See "Viewing or downloading available release images" on page 98.

Viewing or removing repository images stored in your repository

You can use the Deployment Server to conveniently view or remove the release images that are stored in your repository.

To view or remove release images stored in your repository

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images      M) Update Metadata
V) Version Check Systems         S) Set Preferences
U) Upgrade/Install Systems       T) Terminology and Usage
?) Help                          Q) Quit

Enter a Task: [R,M,V,S,U,T,?,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

Select the option you would like to perform [1-3,b,q,?]
```

3 Select option **2**, **View/Remove Repository Images**, to view or remove the release images currently installed on your system.

You see a list of platforms and release levels if you have downloaded the corresponding Base, Maintenance, or Hot Fix release on that platform.

```
To view or remove repository images, the platform type and
release level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  SLES10 x86_64             8)  SLES11 x86_64
9)  Solaris 9 Sparc          10)  Solaris 10 Sparc
11)  Solaris 10 x64          12)  Solaris 11 Sparc
13)  Solaris 11 x64           b)  Back to previous menu

Select the platform of the release to view/remove [1-13,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels if you have downloaded the corresponding Base, Maintenance, or Hot Fix release.

```
1) Base
2) Maintenance
3) Hot Fix
b) Back to previous menu

Select the level of the <platform> releases to view/remove
[1-3,b,q]
```

5 Select the number corresponding to the type of release you want to view or remove (Base, Maintenance, or Hot Fix).

You see a list of releases that are stored in your repository.

```
Stored Repository Releases:

release_version SORT_release_name        OBS AI
================================================
6.0RP1          sfha-sol10_sparc-6.0RP1   -   Y
6.0.3           sfha-sol10_sparc-6.0.3    -   Y
```

6 If you are interested in viewing more details about a release image stored in your repository, type the release version. For example, enter the following:

```
6.0.3
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to remove a sol10_sparc Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all releases stored in your repository that match the selected platform and release level.

```
1)  6.0RP1
2)  6.0.3
b)  Back to previous menu

Select the patch release to remove [1-2,b,q] 1
```

8 Type the number corresponding to the release version you want to remove.

The release images are removed from the Deployment Server.

```
Removing sfha-sol10_sparc-6.0RP1-patches ..................... Done
sfha-sol10_sparc-6.0RP1-patches has been removed successfully.
```

9 From the menu, select option **2**, **View/Remove Repository Images**, and follow the prompts to check that the release images are removed.

See "Viewing or downloading available release images" on page 98.

Deploying Symantec product updates to your environment

After you install at least one Symantec 6.1 product on a server, you can use the Deployment Server to deploy release images to the systems in your environment as follows:

- If you are not sure what to deploy, perform a version check. A version check tells you if there are any Symantec products installed on your systems. It suggests patches and maintenance releases, and gives you the option to install updates.
  See "Finding out which releases you have, and which upgrades or updates you may need" on page 106.

- If you know which update you want to deploy on your systems, use the Upgrade/Install Systems script to deploy a specific Symantec release.
  See "Deploying a specific Symantec release" on page 108.

Finding out which releases you have, and which upgrades or updates you may need

Use the Version Check script to determine which Symantec product you need to deploy. The Version Check script is useful if you are not sure which releases you already have installed, or want to know about available releases.

The Version Check script gives the following information:

- Installed products and their versions (base, maintenance releases, and hot fixes)
- Installed packages (required and optional)
- Available releases (base, maintenance releases, and hot fixes) relative to the version which is installed on the system

To determine which Symantec product updates to deploy

1. Launch the Deployment Server. For example, enter the following:

   ```
   # /opt/VRTS/install/deploy_sfha
   ```

   You see the following output:

   ```
   Task Menu:

   R) Manage Repository Images              M) Update Metadata
   V) Version Check Systems                 S) Set Preferences
   U) Upgrade/Install Systems               T) Terminology and Usage
   ?) Help                                  Q) Quit

   Enter a Task: [R,M,V,S,U,T,?,Q]
   ```

2. Select option **V**, **Version Check Systems**.

3 At the prompt, enter the system names for the systems you want to check. For example, enter the following:

```
sys1
```

You see output for the installed packages (required, optional, or missing).

You see a list of releases available for download.

```
Available Base Releases for Veritas Storage Foundation HA 6.0.1:
None

Available Maintenance Releases for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name      DL OBS AI rel_date   size_KB
=====================================================================
6.0.3           sfha-sol10_sparc-6.0.3 Y  -   -  2013-02-01 212507

Available Public Hot Fixes for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name        DL OBS AI rel_date   size_KB
=========================================================================
6.0.1.200-fs    fs-sol10_sparc-6.0.1.200 -  Y   -  2012-09-20 14346
6.0.1.200-vm    vm-sol10_sparc-6.0.1.200 -  Y   -  2012-10-10 47880

Would you like to download the available Maintenance or Public Hot
Fix releases which cannot be found in the repository? [y,n,q] (n) y
```

4 If you want to download any of the available maintenance releases or hot fixes, enter **y**.

5 If you have not set a default repository for releases you download, the installer prompts you for a directory. (You can also set the default repository in **Set Preferences**).

See "Setting deployment preferences" on page 95.

You can also specify a non-default repository location using the command line.

See "Using the Deployment Server command line option to specify a non-default repository location" on page 96.

6 Select an option for downloading products.

The installer downloads the releases you specified and stores them in the repository.

Deploying a specific Symantec release

After you install at least one Symantec 6.1 product on a server, you can use the Deployment Server to deploy your licensed Symantec products dating back to version 5.1. If you know which product version you want to install, follow the steps in this section to install it.

To deploy a specific Symantec release

1. From the directory in which you installed your Symantec product (version 6.1 or later), launch the Deployment Server with the upgrade and install systems option. For example, enter the following:

   ```
   # /opt/VRTS/install/deploy_sfha
   ```

   You see the following output:

   ```
   Task Menu:

   R) Manage Repository Images             M) Update Metadata
   V) Version Check Systems                S) Set Preferences
   U) Upgrade/Install Systems              T) Terminology and Usage
   ?) Help                                 Q) Quit

   Enter a Task: [R,M,V,S,U,T,?,Q]
   ```

2. Select option **U**, **Upgrade/Install Systems**.

   You see the following output:

   ```
       1)  AIX 5.3
       2)  AIX 6.1
       3)  AIX 7.1
       4)  RHEL5 x86_64
       b)  Back to previous menu

   Select the platform of the available release(s) to be
   upgraded/installed [1-4,b,q,?]
   ```

3 Select the number corresponding to the platform for the release you want to deploy. For example, select the number for the **RHEL5 x86_64** release or the **AIX 6.1** release.

4 Select a Symantec product release.

The installation script is executed and the release is deployed on the specified server.

Updating release information on systems without Internet access

When you install the Deployment Server, the installation includes product metadata that includes information about Symantec, all prior base, maintenance, and hot fix releases across all the products and the platforms. If your system has Internet access, release matrix information is automatically updated from the Symantec Operations Readiness Tools (SORT) site with each use. If your system is behind a firewall, these updates are not possible and the release matrices eventually get out of date.

To update release information on systems without Internet access, you can download a `.tar` file (`deploy_sfha.tar`) containing all the latest release matrices. Then, load it on to your Deployment Server. The `deploy_sfha.tar` file available from the SORT site is updated on a daily basis, and there are typically several release updates every week captured in the updates.

Downloading a .tar file from the SORT site

Loading releases and hot fixes onto your Deployment Server

Downloading a .tar file from the SORT site

To obtain a `.tar` file with release updates, the easiest method is to download a copy from the SORT website.

To download a .tar file from the SORT site

1 Navigate to the following link:

https://sort.symantec.com/support/related_links/offline-release-updates

2 Click on **deploy_sfha.tar [Download]**.

3 Save the file to your desktop.

Loading releases and hot fixes onto your Deployment Server

In this procedure, the Internet-enabled system is the system to which you downloaded the `deploy_sfha.tar` file.

To load releases and hot fixes onto your Deployment Server

1. On the Internet-enabled system, copy the `deploy_sfha.tar` file you downloaded to a location accessible by the Deployment Server.
2. On the Deployment Server, change to the installation directory. For example, enter the following:

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

3. Run the Deployment Script. Enter the following:

   ```
   # ./deploy_sfha
   ```

4. Select option **M**, **Update Metadata**, and enter the location of the `deploy_sfha.tar` file (the installer calls it a "meta-data tar file").

   ```
   Enter the location of the meta-data tar file: [b]
   (/opt/VRTS/install/deploy_sfha.tar)
   ```

   For example, enter the location of the meta-data tar file:

   ```
   /tmp/deploy_sfha.tar
   ```

Section 4

Upgrade of SF

- Chapter 12. Planning to upgrade SF
- Chapter 13. Upgrading Storage Foundation
- Chapter 14. Performing an automated SF upgrade using response files
- Chapter 15. Upgrading SF using Live Upgrade and Boot Environment upgrade
- Chapter 16. Migrating from Storage Foundation Basic to Storage Foundation Standard
- Chapter 17. Performing post-upgrade tasks

Chapter 12

Planning to upgrade SF

This chapter includes the following topics:

- Upgrade methods for SF
- Supported upgrade paths for SF 6.1
- About using the installer to upgrade when the root disk is encapsulated
- Preparing to upgrade SF
- Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes

Upgrade methods for SF

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 12-1 Review this table to determine how you want to perform the upgrade

Upgrade types and considerations	Methods available for upgrade
Typical upgrades—use a Symantec provided tool or you can perform the upgrade manually. Requires some server downtime.	Script-based—you can use this method to upgrade for the supported upgrade paths Web-based—you can use this method to upgrade for the supported upgrade paths Manual—you can use this method to upgrade from the previous release Response file—you can use this method to upgrade from the supported upgrade paths

Table 12-1 Review this table to determine how you want to perform the upgrade *(continued)*

Upgrade types and considerations	Methods available for upgrade
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
Upgrade from any supported UNIX or Linux platform to any other supported UNIX or Linux platform.	Deployment Server See "About the Deployment Server" on page 90.
Simultaneously upgrade base releases, maintenance patches, and hot fixes.	Install Bundles See "Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes" on page 127.

Supported upgrade paths for SF 6.1

The following tables describe upgrading to 6.1.

Table 12-2 Solaris SPARC upgrades using the script- or web-based installer

Symantec product versions	Solaris 8 or older	Solaris 9	Solaris 10	Solaris 11
3.5 3.5 MP4 4.0 4.0 MP1 4.0 MP2	Upgrade the operating system to at least Solaris 10 Update 9, 10, or 11, then upgrade product to 5.0 MP3. Upgrade to 6.1 using the installer script. **Note:** For Symantec product version 3.5, upgrade to 3.5 MP4 prior to upgrading to 5.0 MP3. For Symantec product versions 4.0 and 4.0 MP1, upgrade to 4.0 MP2 prior to upgrading to 5.0 MP3.	Upgrade the operating system to at least Solaris 10 Update 9, 10, or 11, then upgrade product to 5.0 MP3. Upgrade to 6.1 using the installer script. **Note:** For Symantec product version 3.5, upgrade to 3.5 MP4 prior to upgrading to 5.0 MP3. For Symantec product versions 4.0 and 4.0 MP1, upgrade to 4.0 MP2 prior to upgrading to 5.0 MP3.	N/A	N/A
4.1 4.1 MP1 4.1 MP2 5.0 5.0 MP1	Upgrade the operating system to at least Solaris 10 Update 9, 10, or 11, then upgrade product to 5.0 MP3. Upgrade to 6.1 using the installer script.	Upgrade the operating system to at least Solaris 10 Update 9, 10, or 11, then upgrade product to 5.0MP3. Upgrade to 6.1 using the installer script.	Upgrade the product to 5.0 MP3. Upgrade to 6.1 using the installer script.	N/A

Table 12-2 Solaris SPARC upgrades using the script- or web-based installer *(continued)*

Symantec product versions	Solaris 8 or older	Solaris 9	Solaris 10	Solaris 11
5.0 MP3 5.0 MP3 RPx	Upgrade the operating system to at least Solaris 10 Update 9, 10, or 11. Upgrade to 6.1 using the installer script.	Upgrade operating system to at least Solaris 10 Update 9, 10, or 11. Upgrade to 6.1 using the installer script.	Upgrade directly to 6.1 using the installer script.	N/A
5.1 5.1 RPx 5.1 SP1 5.1 SP1 RPx	N/A	Upgrade the operating system to at least Solaris 10 Update 9, 10, or 11. Upgrade to 6.1 using the installer script.	Upgrade directly to 6.1 using the installer script.	N/A
6.0 6.0 RP1	N/A	N/A	Upgrade directly to 6.1 using the installer script.	N/A
6.0 PR1	N/A	N/A	N/A	Upgrade operating system to one of the supported Solaris versions, and then upgrade to 6.1 using the installer script. See the *Storage Foundation Release Notes* for the supported Solaris versions.

Table 12-2 Solaris SPARC upgrades using the script- or web-based installer *(continued)*

Symantec product versions	Solaris 8 or older	Solaris 9	Solaris 10	Solaris 11
6.0.1 6.0.3	N/A	N/A	Upgrade directly to 6.1 using the installer script.	Upgrade operating system to one of the supported Solaris versions, and then upgrade to 6.1 using the installer script. See the *Storage Foundation Release Notes* for the supported Solaris versions.

Note: Starting with Solaris version 11.1, DMP native support provides support for ZFS root devices. On Solaris 11.1 or later, if DMP native support is enabled, then upgrading SF enables ZFS root support automatically. However, if you upgrade from a previous Solaris release to Solaris 11.1, DMP support for ZFS root devices is not automatically enabled. You must enable support explicitly.

See "Enabling DMP support for native devices" on page 201.

About using the installer to upgrade when the root disk is encapsulated

In prior versions of SF, when upgrading a system with an encapsulated root disk, you first had to unencapsulate. When upgrading to SF 6.1, that is no longer necessary, as shown in the table below.

Table 12-3 Upgrading using the installer when the root disk is encapsulated

Starting version	Ending version	Action required
5.0 MP3 RPx	6.1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Table 12-3 Upgrading using the installer when the root disk is encapsulated *(continued)*

Starting version	Ending version	Action required
5.1 or 5.1 RPx	6.1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 SP1 or 5.1 SP1 RPx	6.1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0 or 6.0 RPx	6.1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0.1 or 6.0.3	6.1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Preparing to upgrade SF

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 *Symantec 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/
- Perform the following system-level settings:
  - Set `diag-level` to `min` to perform the minimum number of diagnostics when the system boots. Depending on the configuration of your systems you may want to turn it on after you perform the upgrade.

```
{1} ok setenv diag-level min

 diag-level=min
```

- Set **auto-boot?** to `false`. For tight control when systems restart, set this variable to false. Re-enable this variable after the upgrade.

  ```
  {1} ok setenv auto-boot? false

  auto-boot?=false
  ```

- Deactivate cron to make sure that extraneous jobs are not performed while you upgrade the systems. Do one of the following:
  Solaris 9:

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

  Solaris 10:

  ```
  # svcadm disable -t svc:system/cron:default
  ```

  Solaris 11:

  ```
  # ps -ef | grep cron
  # kill cron pid
  # svcadm disable svc:/system/cron:default
  ```

- If zones are present, make sure that all non-global zones are booted and are in the running state before you use the Symantec product installer to upgrade the Storage Foundation products in the global zone so that any packages present inside non-global zones also gets updated automatically.
  For Oracle Solaris 10, if the non-global zones are not mounted and running at the time of the upgrade, you have to attach the zone with `-U` option to upgrade the SFHA packages inside non-global zone.
  For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you restart the alternative root, you can install `VRTSodm`.
- 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.
  See "Creating backups" on page 120.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the packages, 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 under `/tmp`, which is erased during a system restart.
Do not put the files on a file system that is inaccessible before running the upgrade script.
You can use a Symantec-supplied disc for the upgrade as long as modifications to the upgrade script are not required.
If `/usr/local` was originally created as a slice, modifications are required.

- Unmount all the file systems not on the `root disk`. Comment out their entries in `/etc/vfstab`. Stop the associated volumes and deport the associated disk groups. Any file systems that the Solaris operating system or Storage Foundation assumes should be in `rootdg` but are not, must be unmounted, and the associated entry in `/etc/vfstab` commented out.
- For any startup scripts in `/usr/sbin/svcadm disable`, 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 6.1 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 Symantec products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/vfstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/vfstab` and not mounted during the upgrade. The active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure that the file systems are clean before upgrading.
See "Verifying that the file systems are clean" on page 125.
- Symantec recommends that you upgrade VxFS disk layouts to a supported version before installing VxFS 6.1. Unsupported disk layout versions 4, 5, and 6 can be mounted for the purpose of online upgrading in VxFS 6.1. You can upgrade unsupported layout versions online before installing VxFS 6.1.
- Upgrade arrays (if required).
See "Upgrading the array support" on page 126.
- 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 failback point.
- Determine if the root disk is encapsulated.
See "Determining if the root disk is encapsulated" on page 121.

- If CP server-based coordination points are used in your current fencing configuration, then check that your CP servers are upgraded to 6.1 before starting the upgrade process.
- Make sure that DMP support for native stack is disabled (`dmp_native_support=off`). If DMP support for native stack is enabled (`dmp_native_support=on`), the installer may detect it and ask you to restart the system.

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.

   Back up the `/etc/system` file.
3. Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

   If not, a warning message is displayed.

   Warning: Backup `/etc/vx/cbr/bk` directory.
4. Copy the `vfstab` file to `vfstab.orig`:

   ```
   # cp /etc/vfstab /etc/vfstab.orig
   ```
5. Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
6. If you install the high availability version of the Symantec Storage Foundation 6.1 software, follow the guidelines that are given in the *Symantec Cluster Server Installation Guide* and *Symantec Cluster Server Release Notes* for information on preserving your VCS configuration across the installation procedure.
7. Back up the external `quotas` and `quotas.grp` files.

   If you are upgrading from 6.0.3, you must also back up the `quotas.grp.64` and `quotas.64` files.

8 If you are planning on performing a Phased or Rolling upgrade from 6.0.3 and use quotas, you need to disable them:

```
# vxquotaoff -av
```

9 Verify that quotas are turned off on all the mounted file systems.

Tasks for upgrading the Storage Foundation for Databases (SFDB)

Tasks for upgrading SFDB tools to version 6.1:

- Preparing to migrate the repository database before upgrading from 5.0x or earlier to 6.1
  See "Pre-upgrade tasks for migrating the SFDB repository database" on page 121.
- Migrating the repository database after upgrading from 5.0.x or earlier to 6.1
  See "Post upgrade tasks for migrating the SFDB repository database" on page 179.

Caution: If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 6.1: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 6.1.

Determining if the root disk is encapsulated

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

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

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

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

See "About using the installer to upgrade when the root disk is encapsulated" on page 116.

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 6.1 before upgrading to Storage Foundation or Storage Foundation for Oracle RAC 6.1.

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 6.1.

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 6.1.

Pre-upgrade planning for Volume Replicator

Before installing or upgrading 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.
- If replication using VVR is configured, Symantec recommends that the disk group version is at least 110 prior to upgrading.
  You can check the Disk Group version using the following command:

```
# vxdg list diskgroup
```

- If replication using VVR is configured, make sure the size of the SRL volume is greater than 110 MB.
  Refer to the *Symantec 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.

```
# /usr/sbin/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.
- Make sure that you have worked out all terminal emulation issues. Make sure that the terminal you use is fully functional for OpenBoot prompts and single-user and multi-user run levels.

See the *Symantec 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 *Symantec 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 12-4, if either the Primary or Secondary are running a version of VVR prior to 6.1, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 6.1, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 12-4 VVR versions and checksum calculations

VVR prior to 6.1 (DG version <= 140)	VVR 6.1 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes

Table 12-4 VVR versions and checksum calculations *(continued)*

VVR prior to 6.1 (DG version <= 140)	VVR 6.1 (DG version >= 150)	VVR calculates checksum TCP connections?
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

Additional settings for using VVR in a localized environment

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

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

- Install the required client locale from the Operating System disc.
- Install the required Volume Manager and VVR localized packages.
- Set the client locale, before using any of the VVR interfaces:
  - For the VVR command line, set the locale using the appropriate method for your operating system.
  - For VRW, select the locale from the VRW login page.

Verifying that the file systems are clean

Verify that all file systems have been cleanly unmounted.

To make sure the file systems are clean

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

```
# echo "8192B.p S" | /opt/VRTS/bin/fsdb filesystem | \
    grep clean
    flags 0 mod 0 clean clean_value
```

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

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

```
# /opt/VRTS/bin/fsck -F vxfs filesystem
# /opt/VRTS/bin/mount -F vxfs Block_Device
   mountpoint
# /opt/VRTS/bin/umount mountpoint
```

   These commands should complete any extended operations on the file system and unmount the file system cleanly.

   A pending large package clone removal extended operation might be in progress if the `umount` command fails with the following error:

```
file system device busy
```

   An extended operation is in progress if the following message is generated on the console:

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

3. If an extended operation is in progress, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large package clone can take several hours.
4. Repeat step 1 to verify that the unclean file system is now clean.

Upgrading the array support

The Storage Foundation 6.1 release includes all array support in a single package, `VRTSaslapm`. The array support package includes the array support previously included in the `VRTSvxvm` package. The array support package also includes support previously packaged as external Array Support Libraries (ASLs) and array policy modules (APMs).

See the 6.1 Hardware Compatibility List for information about supported arrays.

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` package exits with an error if external ASLs or APMs are detected.

After you have installed Storage Foundation 6.1, Symantec provides support for new disk arrays through updates to the `VRTSaslapm` package.

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

Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes

Beginning with version 6.1, Symantec offers you a method to easily install or upgrade your systems directly to a base, maintenance, or hot fix level in one step using Install Bundles. With Install Bundles, the installers have the ability to merge so that customers can install or upgrade directly to maintenance or hot fix levels in one execution. The various scripts, packages, and patch components are merged, and multiple releases are installed together as if they are one combined release. You do not have to perform two or more install actions to install or upgrade systems to maintenance levels or hot fix levels.

Releases are divided into the following categories:

Table 12-5 Release Levels

Level	Content	Form factor	Applies to	Release types	Download location
Base	Features	packages	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect
Maintenance	Fixes, new features	packages, patches	All products	Maintenance Release (MR)	Symantec Operations Readiness Tools (SORT)
Hot fix	Fixes	packages	Single product	P-Patch, Public hot fix, Private hot fix	SORT, Support site

When you install or upgrade using Install Bundles:

- SFHA products are discovered and assigned as a single version to the maintenance level. Each system can also have one or more hot fixes applied.
- Base releases are accessible from FileConnect that requires customer serial numbers. Maintenance and hot fix releases can be automatically downloaded from SORT. You can download them from the SORT website manually or use the `deploy_sfha` script.
- Public hot fix releases can be installed using automated installers from the 6.1 version or later.
- Private hot fixes can now be detected to prevent upgrade conflict. Private hot fix releases are not offered as a combined release. They are only available from Symantec Technical Support on a need basis.

You can use the `-base_path` and `-hotfix_path` options to import installation code from multiple releases. You can find packages and patches from different media paths, and merge package and patch definitions for multiple releases. You can use these options to use new task and phase functionality to correctly perform required operations for each release component. You can install the packages and patches in defined phases using these options, which helps you when you want to perform a single start or stop process and perform pre and post operations for all level in a single operation.

Four possible methods of integration exist. All commands must be executed from the highest base or maintenance level install script.

For example:

1. Base + maintenance:

   This integration method can be used when you install or upgrade from a lower version to 6.1.1.

   Enter the following command:

   ```
   # installmr -base_path <path_to_base>
   ```

2. Base + hot fix:

   This integration method can be used when you install or upgrade from a lower version to 6.1.0.100.

   Enter the following command:

   ```
   # installer -hotfix_path <path_to_hotfix>
   ```

3. Maintenance + hot fix:

This integration method can be used when you upgrade from version 6.1 to 6.1.1.100.

Enter the following command:

```
# installmr -hotfix_path <path_to_hotfix>
```

4. Base + maintenance + hot fix:

   This integration method can be used when you install or upgrade from a lower version to 6.1.1.100.

   Enter the following command:

   ```
   # installmr -base_path <path_to_base>
   -hotfix_path <path_to_hotfix>
   ```

   Note: For the 6.1 release, you can add a maximum of five hot fixes using *-hotfix_path <path_to_hotfix> -hotfix2_path <path_to_hotfix> ... -hotfix5_path <path_to_hotfix>*

Chapter 13

Upgrading Storage Foundation

This chapter includes the following topics:

- Upgrading Storage Foundation with the product installer when OS upgrade is not required
- Upgrading Storage Foundation to 6.1 using the product installer or manual steps
- Upgrading SF using the web-based installer
- Upgrading the Solaris operating system
- Upgrading Volume Replicator
- Upgrading language packages

Upgrading Storage Foundation with the product installer when OS upgrade is not required

This section describes upgrading to the current Storage Foundation if the root disk is unencapsulated, and you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 6.1.

To upgrade Storage Foundation

1. Log in as superuser.
2. If the root disk is encapsulated under VxVM, unmirror and unencapsulate the root disk as described in the following steps, to be performed in the order listed:

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt`, and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

  ```
  # vxplex -o rm dis mirrootvol-01 mirswapvol-01
  ```

  Warning: Do not remove the plexes on the root disk that corresponds to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

  ```
  # /etc/vx/bin/vxunroot
  ```

  Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

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

  - Partition table entries for the private regions and public regions
  - GRUB or LILO configuration entries for VxVM

3. Unmount any mounted VxFS file systems.

   The installer supports the upgrade of multiple hosts, if each host is running the same version of VxVM and VxFS. Hosts must be upgraded separately if they are running different versions.

4. If your system has separate `/opt` and `/var` file systems, make sure that they are mounted before proceeding with installation.

5. If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date:

   ```
   # /usr/sbin/vxrlink -g diskgroup status rlink_name
   ```

   Note: Do not continue until the Primary RLINKs are up-to-date.

6. Load and mount the disc. If you downloaded the software, navigate to the top level of the download directory.

7 From the disc, run the `installer` command. If you downloaded the software, run the `./installer` command.

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

8 Enter `G` to upgrade and select the **Full Upgrade**.

9 You are prompted to enter the system names (in the following example, "sys1") on which the software is to be installed. Enter the system name or names and then press Return.

```
Enter the system names separated by spaces on which to
install SF: sys1 sys2
```

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

10 The installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.

11 The installer lists the packages to install or to update. You are prompted to confirm that you are ready to upgrade.

12 The installer discovers if any of the systems that you want to upgrade have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the system's book disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer **y**.

13 The installer then prompts you to name the backup boot disk group. Enter the name for it or press **Enter** to accept the default.

14 You are prompted to start the split operation. Press **y** to continue.

Note: The split operation can take some time to complete.

---

15 Stop the product's 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 it upgrades.

16 The installer stops, uninstalls, reinstalls, and starts specified packages.

17 If the upgrade was done from 5.0 or if the Storage Foundation was done without `vxkeyless` keys, the installer shows the following warning:

```
CPI WARNING V-9-40-5323 SF license version 5.0 is not
updated to 6.1 on sys1. It's recommended to upgrade to a 6.1 key.
CPI WARNING V-9-40-5323 SF license version 5.0 is not updated
to 6.1 on sys2. It's recommended to upgrade to a 6.1 key.
SF is licensed on the systems.
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

18 The Storage Foundation software is verified and configured.

19 The installer prompts you to provide feedback, and provides the log location for the upgrade.

20 Only perform this step if you have split the mirrored root disk to back it up. After a successful restart, verify the upgrade and re-join the backup disk group into the upgraded boot disk group. If the upgrade fails, revert the upgrade boot disk group to the backup disk group.

See "Re-joining the backup boot disk group into the current disk group" on page 178.

See "Reverting to the backup boot disk group after an unsuccessful upgrade" on page 178.

Upgrading Storage Foundation to 6.1 using the product installer or manual steps

This section describes upgrading SF from a previous release to 6.1. Symantec recommends that you perform this upgrade from single-user mode.

No VxFS file systems can be in use at the time of the upgrade.

Choose the appropriate procedure for your situation.

- If the current Storage Foundation product is installed on an operating system supported by 6.1, you do not need to upgrade the operating system. If you do not plan to upgrade the operating system, use one of the following upgrade procedures:
  - Upgrade SF but not OS with the product installer.
    For the recommended upgrade procedure:
    See "Upgrading Storage Foundation with the product installer" on page 134.
  - Upgrade SF but not OS with manual steps (`pkgadd` command).

- If you plan to upgrade the operating system, you must perform additional steps to upgrade. If the current Storage Foundation product is installed on an operating system which is no longer supported by 6.1, you must upgrade the operating system. If you plan to upgrade the operating system, use the following upgrade procedure:
  See "Upgrading Symantec Storage Foundation to 6.1 using upgrade scripts (OS upgrade)" on page 136.

Upgrading Storage Foundation with the product installer

This section describes upgrading to the current Storage Foundation, and you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 6.1.

To upgrade Storage Foundation

1. Log in as superuser.
2. Unmount any mounted VxFS file systems.

   The installer supports the upgrade of multiple hosts, if each host is running the same version of VxVM and VxFS. Hosts must be upgraded separately if they are running different versions.

   If any VxFS file systems are mounted with the QuickLog feature, QuickLog must be disabled before you upgrade. See the "Veritas QuickLog" chapter of the *Veritas File System Administrator's Guide* for more information.
3. If your system has separate `/opt` and `/var` file systems, make sure that they are mounted before proceeding with installation.
4. If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date:

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

   Note: Do not continue until the Primary RLINKs are up-to-date.
5. Load and mount the disc.

   See "Mounting the product disc" on page 42.
6. To invoke the common installer, run the `installer` command on the disc as shown in this example:

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

7 Enter `G` to upgrade and press Return.

8 You are prompted to enter the system names (in the following example, "host1"). Enter the system name or names and then press Return.

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

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

9 Installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.

10 You can perform this step if you want to upgrade from SF 5.1 SP1 for Solaris.

The installer discovers if any of the systems that you want to upgrade have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the system's book disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer **y**.

Note: Splitting the mirrors for the root disk group backup requires a restart upon completion of the upgrade.

---

11 The installer then prompts you to name the backup boot disk group. Enter the name for it or press **Enter** to accept the default.

Note: The split operation can take some time to complete.

---

12 You are prompted to start the split operation. Press **y** to continue.

13 Stop the product's processes.

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

14 The installer lists the packages to install or upgrade, and performs the installation or upgrade.

15 If the upgrade was done from 5.0 or if the Storage Foundation was done without `vxkeyless` keys, the installer shows the following warning:

```
CPI WARNING V-9-40-5323 SF license version 5.0 is not
updated to 6.1 on sys1. It's recommended to upgrade to a 6.1 key.
CPI WARNING V-9-40-5323 SF license version 5.0 is not updated
to 6.1 on sys2. It's recommended to upgrade to a 6.1 key.
SF is licensed on the systems
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

16 The installer verifies, configures, and starts the Symantec Storage Foundation software.

17 Only perform this step if you have split the boot disk group into a backup disk group. After a successful restart, verify the upgrade and re-join the backup disk group. If the upgrade fails, revert to the backup disk group.

See "Re-joining the backup boot disk group into the current disk group" on page 178.

See "Reverting to the backup boot disk group after an unsuccessful upgrade" on page 178.

Upgrading Symantec Storage Foundation to 6.1 using upgrade scripts (OS upgrade)

This section describes upgrading to the current Symantec Storage Foundation and the need to upgrade the Solaris version. If the operating system is not at a supported Solaris version, you must follow this procedure.

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

These steps have to be followed in the specified order.

To begin the upgrade

1 If VCS agents for VVR are configured, you must perform the pre-upgrade steps before proceeding.

2 Load and mount the disc.

See "Mounting the product disc" on page 42.

3 Verify that an upgrade is possible on the system. Enter the following command:

```
# /dvd_mount/scripts/upgrade_start -check
```

4 Run the `upgrade_start` script to preserve the previous configuration of Volume Manager.

```
# /dvd_mount/scripts/upgrade_start
```

5 If the `upgrade_start` script fails for any reason, run the `upgrade_finish` script to undo any changes already made. Compare `/etc/system`, `/etc/vfstab`, and the output of the `format` command and verify that the system is restored. Then determine and correct the cause of the `upgrade_start` failure. If you cannot correct the problem in a timely manner, restore the `vfstab` file to the version saved, restore any other applications, and perform an `init 6` to completely restore the system.

6 Verify that all the Primary RLINKs are up-to-date on all the hosts.

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

Caution: Do not continue until the Primary RLINKs are up-to-date.

7 If VVR is configured, run the `vvr_upgrade_start` script on all hosts to save the original VVR configuration:

```
# /dvd_mount/scripts/vvr_upgrade_start
```

8 If you have VxFS file systems specified in the `/etc/vfstab` file, comment them out.

9 Remove the existing Storage Foundation packages in one of the following ways:

- Using the uninstallsf script
- Using the `pkgrm` command

For details, refer to the *Storage Foundation Installation Guide* for the existing Storage Foundation version.

After you run the uninstallsf script, verify that all VRTS* packages are removed; otherwise, remove them manually using `pkgrm`.

10 If you want to upgrade the operating system, do so now.

Refer to the Solaris installation documentation.

11 Install the Storage Foundation packages in one of the following ways:

- Using the common installer

See “To upgrade the Symantec Storage Foundation packages with the product installer” on page 138.

- Using manual steps
  See “To upgrade the Symantec Storage Foundation packages with manual steps” on page 138.

To upgrade the Symantec Storage Foundation packages with the product installer

1. Load and mount the disc.

   See “Mounting the product disc” on page 42.

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

   ```
   # cd /dvd_mount
   # ./installer
   ```

3. Select **I** to upgrade the product. The installer asks you if you want to use the previous configuration.
4. Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
5. If you commented out VxFS File System entries in the `/etc/vfstab` file, uncomment them.
6. Complete the upgrade by restoring the configuration.

To upgrade the Symantec Storage Foundation packages with manual steps

1. If you want to upgrade from Symantec Storage Foundation for Oracle, resynchronize all existing snapshots before upgrading.

   For Symantec Storage Foundation for Oracle:

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

2. Load and mount the software disc.
3. Change to the directory containing the packages.

   ```
   # cd /dvd_mount
   ```

4 Run the following command to obtain a list of recommended packages to install:

```
./installsf -recpkgs
```

Run the following command to obtain a list of all packages to install:

```
./installsf -allpkgs
```

5 Add packages with the `pkgadd` command.

6 If you commented out VxFS File System entries in the `/etc/vfstab` file, uncomment them.

7 Complete the upgrade by restoring the configuration.

Restoring the configuration and completing the upgrade

1 Complete the upgrade using the `upgrade_finish` script.

```
# devlinks
# /dvd_mount/scripts/upgrade_finish
```

2 Configure the product using the following command:

```
# /dvd_mount/installer -configure
```

If some Veritas modules fail to unload, perform the following steps:

- Restart the systems.

3 Importing a pre-6.1 Veritas Volume Manager disk group does not automatically upgrade the disk group version to the VxVM 6.1 level. You may need to manually upgrade each of your disk groups following a VxVM upgrade.

See "Upgrading VxVM disk group versions" on page 186.

Upgrading SF using the web-based installer

This section describes upgrading SF with the 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 want to preserve. For example, make configuration file backups.

2 Start the web-based installer.

See "Starting the web-based installer" on page 52.

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

The product is discovered once you specify the system. Click **Next**.

4 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.

5 Installer detects the product that is installed on the specified system. It shows the cluster information and lets you confirm if you want to perform upgrade on the cluster. Select **Yes** and click **Next**.

6 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**.

7 The installer discovers if any of the systems that you want to upgrade have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the book disk group. To create the backup, check the **Split mirrors on all the systems** box. Check the appropriate box to use the same name for the backup disk group on all systems. You can use the default name or choose a new one. Check the systems where you want to create the backup. When you are ready, click the **Next** button.

8 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.

9 If you are prompted to restart the systems, enter the following restart command:

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

10 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 does not ask any questions.

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

12 Only perform this step if you have split the mirrored root disk to back it up. After a successful restart, verify the upgrade and re-join the backup disk group into the upgraded boot disk group. If the upgrade fails, revert the upgrade boot disk group to the backup disk group.

See “Re-joining the backup boot disk group into the current disk group” on page 178.

See “Reverting to the backup boot disk group after an unsuccessful upgrade” on page 178.

Upgrading the Solaris operating system

If you are running Storage Foundation 6.1 with an earlier release of the Solaris operating system, you can upgrade the Solaris operating system using the following procedure.

Warning: You should only use this procedure to upgrade the Solaris operating system if you are running Storage Foundation 6.1.

The directory `/opt` must exist, be writable, and must not be a symbolic link as the volumes that the `upgrade_start` command does not temporarily convert are unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link does not function during the upgrade and items in `/opt` are not installed.

To upgrade the Solaris operating system only

1. Bring the system down to single-user mode using the following command:

   ```
   # init S
   ```

   You must mount `/opt` manually if `/opt` is on its own partition.

2. Load and mount the software disc from the currently installed version of Storage Foundation.

   See "Mounting the product disc" on page 42.

3. Change directory:

   ```
   # cd /mount_point/scripts
   ```

4. Run the `upgrade_start` with the `-check` argument to detect any problems that exist which can prevent a successful upgrade. Use the `upgrade_start` script that was supplied with the currently installed SF release. If this command reports success, you can proceed with running the `upgrade_start` script, but if it reports errors, correct the problem(s) and rerun `upgrade_start -check`.

   ```
   # ./upgrade_start -check
   ```

5. Run the `upgrade_start` script so that the system can come up with partitions. The `upgrade_start` script searches for volumes containing file systems, and if any are found, converts them to partitions:

   ```
   # ./upgrade_start
   ```

6 Bring the system down to run level 0.

```
# init 0
```

7 Upgrade the operating system to a supported version of Solaris.

You should boot up the system from run level 0 depending on the Solaris upgrade procedure that you want to follow. Refer to the Solaris installation documentation for instructions on how to upgrade the Solaris operating system.

8 After installing the Solaris operating system, install any Solaris patches required by Storage Foundation 6.1.

See the *Storage Foundation Release Notes*.

9 After the system is up with the upgraded Solaris operating system, bring the system down to single-user mode by entering:

```
# init S
```

10 Ensure that `/opt` is mounted.

11 Load and mount the software disc from the currently installed version of Storage Foundation.

12 If you upgraded to Solaris 10, you must reinstall certain Storage Foundation packages to support Solaris 10 functionality.

To reinstall the required packages, follow the steps below:

- Remove the existing packages in the reverse order of their installation. For example, if you chose the installation of all packages then uninstall those in the following order.
  For Storage Foundation:

  ```
  # pkgrm VRTSat VRTSodm VRTSdbed
  VRTSfssdk VRTSvxfs VRTSsfmh VRTSob VRTSaslapm
  VRTSvxvm VRTSspt VRTSperl VRTSvlic
  ```

- Run the following commands.
  To obtain a list of recommended packages to install:

  ```
  # ./installsf -recpkgs
  ```

  Or
  To obtain a list of all packages to install:

  ```
  # ./installsf -allpkgs
  ```

- Change to the directory containing the appropriate packages.

  ```
  # cd /mount_point/pkgs
  ```

- Use the `pkgadd` command to install the packages from the list you generated.
- Restart the system.

13 Complete the upgrade from the software disc from the currently installed version of Storage Foundation by entering:

```
# devlinks
# ./upgrade_finish
```

Upgrading Volume Replicator

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

When upgrading from 4.1 MP1 or later, you have the option to upgrade without disrupting replication.

See "Upgrading VVR without disrupting replication" on page 143.

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 123.

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 sec_hostname
   ```

2. Upgrade from VVR 5.1 or later to VVR 6.1 on the Secondary.

3. Do one of the following:

   - Upgrade the disk group now. Enter the following:

     ```
     # vxdg upgrade dgname
     ```

   - Upgrade the disk group later.
     If you upgrade the disk group later, be sure to pause replication before you upgrade the disk group.

4. 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 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 5.1 or later to VVR 6.1 on the Secondary.

3. Do one of the following:

   - Upgrade the disk group now. Enter the following:

     ```
     # vxdg upgrade dgname
     ```

   - Upgrade the disk group later.

If you upgrade the disk group later, be sure to pause replication before you upgrade the disk group.

4 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 123.

Upgrading language packages

If you want to upgrade Symantec products in a language other than English, you must install the required language packages after installing the English packages. Verify that the English installation is correct before you proceed.

Install the language packages as for an initial installation.

See “Installing language packages” on page 50.

Chapter 14

Performing an automated SF upgrade using response files

This chapter includes the following topics:

- Upgrading SF using response files
- Response file variables to upgrade Storage Foundation
- Sample response file for SF upgrade

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.

Response file variables to upgrade Storage Foundation

Table 14-1 lists the response file variables that you can define to configure SF.

Table 14-1 Response file variables for upgrading SF

Variable	Description
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list 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}{tmppath}	Defines the location where a working directory is created to store temporary files and the packages that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional

Table 14-1 Response file variables for upgrading SF *(continued)*

Variable	Description
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}{upgrade}	Upgrades all packages installed, without configuration. List or scalar: list Optional or required: optional
CFG{mirrordgname}{system}	If the root dg is encapsulated and you select split mirror is selected: Splits the target disk group name for a system. List or scalar: scalar Optional or required: optional
CFG{splitmirror}{system}	If the root dg is encapsulated and you select split mirror is selected: Indicates the system where you want a split mirror backup disk group created. List or scalar: scalar Optional or required: optional
CFG{opt}{disable_dmp_native_support}	If it is set to 1, Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools is disabled after upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases package upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system. List or scalar: scalar Optional or required: optional

Table 14-1 Response file variables for upgrading SF *(continued)*

Variable	Description
CFG{opt}{hotfix_path}	Defines the path of a hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed . List or scalar: scalar Optional or required: optional
CFG{opt}{hotfix2_path}	Defines the path of a second hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{opt}{hotfix3_path}	Defines the path of a third hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{opt}{hotfix4_path}	Defines the path of a fourth hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{opt}{hotfix5_path}	Defines the path of a fifth hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional

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(sys1) ];
1;
```

Chapter 15

Upgrading SF using Live Upgrade and Boot Environment upgrade

This chapter includes the following topics:

- About Live Upgrade
- About ZFS Boot Environment (BE) upgrade
- Supported upgrade paths for Live Upgrade and Boot Environment upgrade
- Performing Live Upgrade on Solaris 10 systems
- Performing Boot Environment upgrade on Solaris 11 systems
- About Live Upgrade in a Volume Replicator (VVR) environment

About Live Upgrade

Solaris Live Upgrade provides a method of upgrading a system while the system continues to operate. This is done by creating an alternate boot environment (ABE) from the current boot environment and then upgrading the ABE. Once the ABE is upgraded, you can activate the ABE and then reboot the system.

On Solaris 10 or previous releases, you can use Live Upgrade technology to reduce downtime associated with the OS upgrade and SF product upgrade by creating a boot environment on a alternate boot disk.

- See "Performing Live Upgrade on Solaris 10 systems" on page 155.

Figure 15-1 illustrates an example of an upgrade of Symantec products from 5.1 SP1 to 6.1, and the operating system from Solaris 9 to Solaris 10 using Live Upgrade.

Figure 15-1 Live Upgrade process

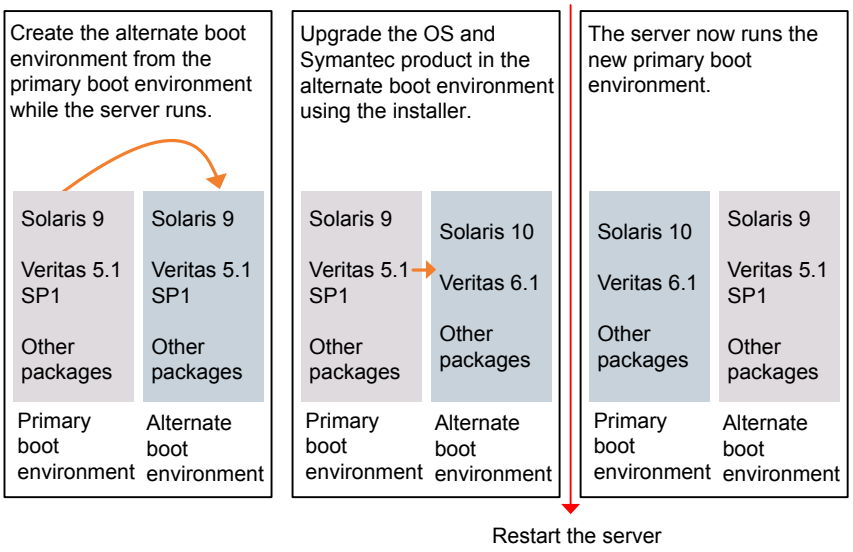

Some service groups (failover and parallel) may be online in this cluster and the Live Upgrade process does not affect them. Downtime is experienced only when the server is restarted to boot into the alternate boot environment.

About ZFS Boot Environment (BE) upgrade

A Boot Environment (BE) is a bootable instance of the Oracle Solaris operating system image along with any other application software packages installed into that image. System administrators can maintain multiple BEs on their systems, and each BE can have different software versions installed. Upon the initial installation of the Oracle Solaris 11 release onto a system, a BE is created.

On Solaris 11, you can use the `beadm` utility to create and administer additional BEs on your system.

Figure 15-2 Boot Environment upgrade process

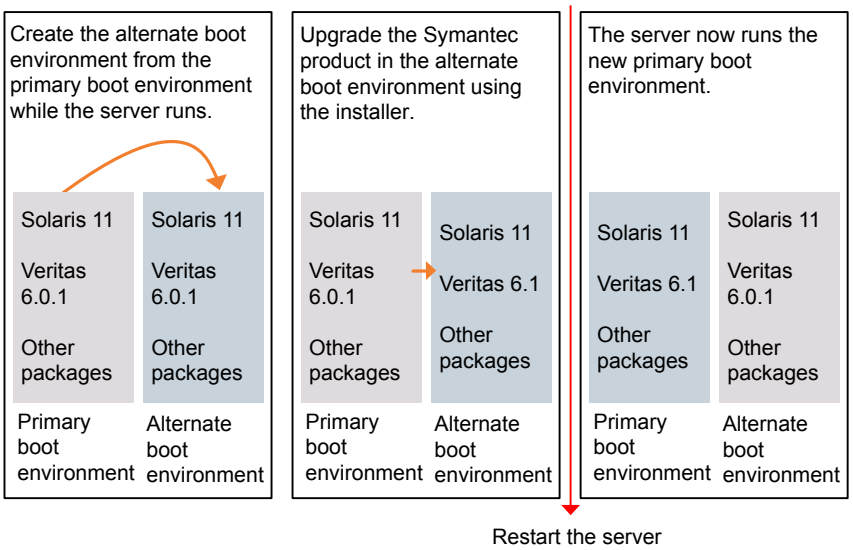

Supported upgrade paths for Live Upgrade and Boot Environment upgrade

The systems where you plan to use Live Upgrade must run Solaris 9 or Solaris 10. Boot environment upgrade can be used on Solaris 11 system only. You can upgrade from those systems that run Solaris 9, but SF 6.1 is not supported on Solaris 9.

For Live Upgrade method, existing SF version must be at least 5.0 MP3. For Boot Environment upgrade method, the SF version you are upgrading to must be at least 6.1.0.

Symantec requires that both global and non-global zones run the same version of Symantec products.

You can use Live Upgrade or Boot environment upgrade in the following virtualized environments:

Table 15-1 Live Upgrade or Boot environment upgrade support in virtualized environments

Environment	Procedure
Solaris native zones	Perform Live Upgrade or Boot Environment upgrade to upgrade both global and non-global zones. If you have a zone root that resides on a VxVM volume, use the following procedure. See "Performing Live Upgrade in a Solaris zone environment on Solaris 10" on page 155. Use the standard procedure for the other standby nodes. See "Performing Live Upgrade on Solaris 10 systems" on page 155. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 167.
Solaris branded zones (BrandZ)	Perform Live Upgrade or Boot Environment upgrade to upgrade the global zone. See "Performing Live Upgrade on Solaris 10 systems" on page 155. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 167. Manually upgrade the branded zone separately. Note that while you can perform a Live Upgrade or Boot Environment upgrade in the presence of branded zones, the branded zones are not upgraded.
Oracle VM Server for SPARC	Use Live upgrade or Boot Environment upgrade procedure for Control domain as well as guest domains. See "Performing Live Upgrade on Solaris 10 systems" on page 155. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 167.

Performing Live Upgrade in a Solaris zone environment on Solaris 10

If you have a zone root that reside on a VxVM volume, for the purpose of Live Upgrade, create another VxVM volume of same or bigger size than that of the existing zone root for copying the file system contents to alternate boot environment. Use VxVM commands for creating the volume.

Use the standard procedure for the other standby nodes.

See "Performing Live Upgrade on Solaris 10 systems" on page 155.

By default, Zone agent `BootState` is set to "multi-user." After you complete the upgrade, you may need to adjust this attribute to the appropriate value before you start your zone through VCS.

Note: Symantec recommends that you set `BootState` to "multi-user-server" to run applications inside non-global zones.

For Solaris 10, make sure that all non-global zones are either in the running or configured state before you use the Symantec product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must attach each non-global zone with update option manually after upgrade.

For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you restart the alternative root, you can install `VRTSodm`.

Performing Live Upgrade on Solaris 10 systems

Perform the Live Upgrade using the installer.

Table 15-2 Upgrading SF using Solaris 10 Live Upgrade

Step	Description
Step 1	Prepare to upgrade using Solaris Live Upgrade. See "Before you upgrade SF using Solaris Live Upgrade" on page 156.
Step 2	Create a new boot environment on the alternate boot disk. See "Creating a new Solaris 10 boot environment on the alternate boot disk" on page 157.

Table 15-2 Upgrading SF using Solaris 10 Live Upgrade *(continued)*

Step	Description
Step 3	Upgrade SF using the installer. See "Upgrading SF using the installer for Solaris 10 Live Upgrade" on page 161.
	To upgrade only Solaris See the Oracle documentation on Solaris 10 operating system **Note:** If you choose to upgrade the operating system, the Solaris operating system on the alternate boot environment is upgraded. A new boot environment is created on the alternate boot disk by cloning the primary boot environment.
Step 4	Switch the alternate boot environment to be the new primary. See "Completing the Solaris 10 Live Upgrade" on page 163.
Step 5	Verify Live Upgrade of SF. See "Verifying the Solaris 10 Live Upgrade of SF" on page 164.

Before you upgrade SF using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

To prepare for the Live Upgrade

1. Make sure that the SF installation media and the operating system installation images are available and on hand.
2. On the nodes to be upgraded, select an alternate boot disk that is at least the same size as the root partition of the primary boot disk
3. Before you perform the Live Upgrade, take offline any services that involve non-root file systems. This prevents file systems from being copied to the alternate boot environment that can potentially cause a root file system to run out of space.
4. On the primary boot disk, patch the operating system for Live Upgrade.

   For upgrade from Solaris 9 to 10:

   - SPARC system: Patch 137447-01 or later is required.

   Verify that the patches are installed.

5 The version of the Live Upgrade packages must match the version of the operating system to which you want to upgrade on the alternate boot disk. If you upgrade the Solaris operating system, do the following steps:

- Remove the installed Live Upgrade packages for the current operating system version:
  All Solaris versions: `SUNWluu`, `SUNWlur` packages.
  Solaris 10 update 7 or later also requires: `SUNWlucfg` package.
  Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.
- From the new Solaris installation image, install the new versions of the following Live Upgrade packages:
  All Solaris versions: `SUNWluu`, `SUNWlur`, and `SUNWlucfg` packages.
  Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.

Solaris installation media comes with a script for this purpose named liveupgrade20. Find the script at /*cdrom*/*solaris_release*/Tools/Installers/liveupgrade20. If scripting, you can use:

```
# /cdrom/solaris_release/Tools/Installers/liveupgrade20 \
-nodisplay -noconsole
```

If the specified image has some missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you have to install any missing patches on the alternate boot disk.

6 If the specified image has some missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you need to install any missing patches on the alternate boot disk.

Creating a new Solaris 10 boot environment on the alternate boot disk

Symantec provides the `vxlustart` script that runs a series of commands to create the alternate boot environment for the upgrade.

To preview the commands, specify the `vxlustart` script with the `-V` option.

Symantec recommends that you preview the commands with `-V` option to ensure there are no problems before beginning the Live Upgrade process. The vxlustart script is located in the scripts directory on the distribution media.

Note: This step can take several hours to complete. Do not interrupt the session as it may leave the boot environment unstable.

```
# cd /cdrom/scripts

# ./vxlustart -V -u targetos_version -s osimage_path -d diskname
```

Table 15-3

vxlustart option	Usage
-V	Lists the commands to be executed during the upgrade process without executing them and pre-checks the validity of the command. If the operating system is upgraded, the user is prompted to compare the patches that are installed on the image with the patches installed on the primary boot disk. This determines if any critical patches are not present from the new operating system image.
-v	Indicates verbose, print commands before executing them.
-f	Forces the vtoc creation on the disk.
-Y	Indicates a default yes with no questions asked.
-m	Uses the already existing vtoc on the disk.
-D	Prints with debug option on, and is for debugging.
-U	Specifies that only the Storage Foundation products are upgraded. The operating system is cloned from the primary boot disk.
-g	Specifies the DG to which the rootdisk belongs. Optional.
-d	Indicates the name of the alternate boot disk `c#t#d#s2` on which you intend to upgrade. The default disk is mirrordisk.

Table 15-3 *(continued)*

vxlustart option	Usage
-u	Specifies the operating system version for the upgrade on the alternate boot disk. For example, use `5.9` for Solaris 9 and `5.10` for Solaris 10. If you want to upgrade only SF products, specify the current OS version.
-F	Specifies the root disk's file system, where the default is `ufs`.
-S	Specifies the path to the Solaris image. It can be a network/directory path. If the installation uses the CD, this option must not be specified. See *Solaris Live Upgrade installation guide* for more information about the path.
-r	Specifies that if the computer crashes or restarts before the `vxlufinish` command is run, the alternate disk is remounted using this option.
-k	Specifies the location of file containing auto-registration information. This file is required by `luupgrade(1M)` for OS upgrade to Solaris 10 9/10 or a later release.
-x	Excludes file from newly created BE. (`lucreate -x` option)
-X	Excludes file list from newly created BE. (`lucreate -f` option
-i	Includes file from newly created BE. (`lucreate -y` option)
-I	Includes file list from newly created BE. (`lucreate -Y` option)
-z	Filters file list from newly created BE. (`lucreate -z` option)
-w	Specifies additional mount points. (`lucreate -m` option)

Table 15-3 *(continued)*

vxlustart option	Usage
-W	Specifies additional mount points in a file (`lucreate -M` option

If the `-U` option is specified, you can omit the `-s` option. The operating system is cloned from the primary boot disk.

For example, to preview the commands to upgrade only the Symantec product:

```
# ./vxlustart -V -u 5.10 -U -d disk_name
```

For example, to preview the commands for an upgrade to Solaris 10 update 6:

```
# ./vxlustart -V -u 5.10 -s /mnt/Solaris_10u6 -d c0t1d0s2
```

In the procedure examples, the primary or current boot environment resides on Disk0 (c0t0d0s2) and the alternate or inactive boot environment resides on Disk1 (c0t1d0s2).

At the end of the process:

- A new boot environment is created on the alternate boot disk by cloning the primary boot environment.
- The Solaris operating system on the alternate boot disk is upgraded, if you have chosen to upgrade the operating system.

To create a new boot environment on the alternate boot disk

1. Navigate to the install media for the Symantec products:

   ```
   # cd /cdrom/scripts
   ```

2. Review the output and note the new mount points. If the system is restarted before completion of the upgrade or if the mounts become unmounted, you may need to remount the disks.

   If you need to remount, run the command:

   ```
   # vxlustart -r -u targetos_version -d disk_name
   ```

3. After the alternate boot disk is created and mounted on */altroot.5.10*, install any operating system patches or packages on the alternate boot disk that are required for the Symantec product installation.

   ```
   # pkgadd -R /altroot.5.10 -d pkg_dir
   ```

Upgrading SF using the installer for Solaris 10 Live Upgrade

You can use the Symantec product installer to upgrade SF as part of the Live Upgrade.

At the end of the process, Storage Foundation 6.1 is installed on the alternate boot disk.

To perform Live Upgrade of SF using the installer

1. Insert the product disc with Storage Foundation 6.1 or access your copy of the software on the network.
2. Run the installer script specifying the root path as the alternate boot disk:

   ```
   # ./installsf -upgrade -rootpath /altroot.5.10
   ```

3. Enter the names of the nodes that you want to upgrade to Storage Foundation 6.1.

   Note: Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the SF installation.

   The installer displays the list of packages to be installed or upgraded on the nodes.
4. Press **Return** to continue with the installation.
5. Verify that the version of the Veritas packages on the alternate boot disk is 6.1.

   ```
   # pkginfo -R /altroot.5.10 -l VRTSpkgname
   ```

   For example:

   ```
   # pkginfo -R /altroot.5.10 -l VRTSvxvm
   ```

   Review the installation logs at `/altroot.5.10/opt/VRTS/install/logs`.

Upgrading SF manually for Solaris 10 Live Upgrade

You can perform a manual upgrade of SF using Live Upgrade.

At the end of the process, Storage Foundation 6.1 is installed on the alternate boot disk.

To perform Live Upgrade of SF manually

1 Remove the SF packages on the alternate boot disk in the reverse order in which they were installed:

- `# pkgrm -R /altroot.5.10 \`
  `VRTSmapro VRTSgapms VRTSvxmsa VRTSfasag VRTSfas VRTSvail \`
  `VRTSfsmnd VRTSfssdk VRTSfsman VRTSvrw VRTSweb VRTSjre15 \`
  `VRTSvcsvr VRTSvrpro VRTSddlpr VRTSvdid VRTSalloc VRTSdcli \`
  `VRTSvmpro VRTSvmman VRTSfspro VRTSdsa VRTSvxvm VRTSvxfs \`
  `VRTSspt VRTSaa VRTSmh VRTSccg VRTSobgui VRTSob VRTSobc33 \`
  `VRTSat VRTSpbx VRTSicsco VRTSvlic VRTSperl`

  Note that this package list is an example. Full package lists vary from release to release and by product option.

The `-R` option removes the packages from the root path `/altroot.5.10` on the alternate boot disk.

2 Install the SF packages from the `/pkgs` directory. You must install the packages in the following order one at a time to the alternate boot disk using the `pkgadd` command:

- For Storage Foundation:

  `VRTSvlic.pkg VRTSperl.pkg VRTSspt.pkg VRTSvxvm.pkg VRTSaslapm.pkg VRTSob.pkg VRTSsfmh.pkg VRTSvxfs.pkg VRTSfssdk.pkg VRTSdbed.pkg VRTSodm.pkg VRTSsfcpi61.pkg`

For example:

`# pkgadd -R /altroot.5.10 -d package_name.pkg`

Where you replace *package_name.pkg* with a package's name, for example `VRTSperl.pkg`.

`# pkgadd -R /altroot.5.10 -d VRTSperl.pkg`

3 Verify that the version of the Veritas packages on the alternate boot disk is 6.1.

`# pkginfo -R /altrootpath -l VRTSpkgname`

For example:

`# pkginfo -R /altroot.5.10 -l VRTSvxvm`

Completing the Solaris 10 Live Upgrade

At the end of the process:

- If the original primary boot disk was encapsulated, the alternate boot disk is encapsulated.
- The alternate boot environment is activated.
- The system is booted from the alternate boot disk.

To complete the Live Upgrade

1. Complete the Live upgrade process using one of the following commands.

   If the primary root disk is not encapsulated, run the following command:

   ```
   # ./vxlufinish -u target_os_version
   Live Upgrade finish on the Solaris release <5.10>
   ```

   If the primary root disk is encapsulated by VxVM, run the following command:

   ```
   # ./vxlufinish -u target_os_version -g diskgroup
   Live Upgrade finish on the Solaris release <5.10>
   ```

   The Live Upgrade process encapsulates the alternate root disk if the primary root disk was encapsulated.

2. After the successful completion of `vxlustart`, if the system crashes or restarts before Live Upgrade completes successfully, you can remount the alternate disk using the following command:

   ```
   # ./vxlustart -r -u target_os_version
   ```

   Then, rerun the `vxlufinish` command from step 1

   ```
   # ./vxlufinish -u target_os_version
   ```

   If you have enabled VVR, See “About Live Upgrade in a Volume Replicator (VVR) environment” on page 173.

3. If you want to upgrade VVR, run the `vvr_upgrade_lu_start` command.

   Note: Only run the `vvr_upgrade_lu_start` command when you are ready to restart the nodes and switch over to the alternate boot environment.

4 Restart the system. The boot environment on the alternate disk is activated when you restart it.

Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

Note: Do not use the `reboot`, `halt`, or `uadmin` commands to restart the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate boot environment.

You can ignore the following error if it appears: Error: boot environment *<dest.13445>* already mounted on *</altroot.5.10>*.

```
# shutdown -g0 -y -i6
```

5 After the alternate boot environment is activated, you can switch boot environments. If the root disk is encapsulated, refer to the procedure to switch the boot environments manually.

See "Administering boot environments in Solaris 10 Live Upgrade" on page 165.

6 After the upgrade, perform any required post-upgrade tasks such as upgrading the disk group.

7 After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

Verifying the Solaris 10 Live Upgrade of SF

To ensure that Live Upgrade has completed successfully, verify that the system have booted from the alternate boot environment.

To verify that Live Upgrade completed successfully

1 Verify that the alternate boot environment is active.

```
# lustatus
```

If the alternate boot environment fails to be active, you can revert to the primary boot environment.

See "Reverting to the primary boot environment on a Solaris 10 system" on page 165.

2 Perform other verification as required to ensure that the new boot environment is configured correctly.

3 In a zone environment, verify the zone configuration.

Administering boot environments in Solaris 10 Live Upgrade

Use the following procedures to perform relevant administrative tasks for boot environments.

Reverting to the primary boot environment on a Solaris 10 system

If the alternate boot environment fails to start, you can revert to the primary boot environment.

Start the system from the primary boot environment in the PROM monitor mode.

```
ok> boot disk0
```

where *disk0* is the primary boot disk.

Switching the boot environment for Solaris SPARC

You do not have to perform the following procedures to switch the boot environment when you use the `vxlufinish` scripts to process Live Upgrade. You must perform the following procedures when you perform a manual Live Upgrade.

Two different procedures exist to switch the boot environment, choose one of the following procedures based on the encapsulation of the root disk:

- See "To switch the boot environment if the root disk is not encapsulated" on page 166.
- See "To switch the boot environment if the root disk is encapsulated" on page 167.

The switching procedures for Solaris SPARC vary, depending on whether VxVM encapsulates the root disk.

To switch the boot environment if the root disk is not encapsulated

1 Display the status of Live Upgrade boot environments.

```
# lustatus

Boot Environment  Is       Active Active    Can    Copy
Name              Complete Now    On Reboot Delete Status
----------------  -------- ------ --------- ------ ------
source.2657       yes      yes    yes       no     -
dest.2657         yes      no     no        yes    -
```

In this example, the primary boot environment is currently (source.2657). You want to activate the alternate boot environment (dest.2657).

2 Unmount any file systems that are mounted on the alternate boot environment (dest.2657).

```
# lufslist dest.2657

               boot environment name: dest.2657

Filesystem        fstype device size  Mounted on Mount Options
----------------- ------ ------------ ---------- -------------
/dev/dsk/c0t0d0s1 swap     4298342400 -          -
/dev/dsk/c0t0d0s0 ufs     15729328128 /          -
/dev/dsk/c0t0d0s5 ufs      8591474688 /var       -
/dev/dsk/c0t0d0s3 ufs      5371625472 /vxfs      -

# luumount dest.2657
```

3 Activate the Live Upgrade boot environment.

```
# luactivate dest.2657
```

4 Restart the system.

```
# shutdown -g0 -i6 -y
```

The system automatically selects the boot environment entry that was activated.

To switch the boot environment if the root disk is encapsulated

1 Display the current boot disk device and device aliases

```
# eeprom
boot-device=vx-rootdg vx-int_disk
use-nvramrc?=true
nvramrc=devalias vx-int_disk /pci@1c,600000/scsi@2/disk@0,0:a
devalias vx-rootdg01 /pci@1c,600000/scsi@2/disk@1,0:a
```

2 Set the device from which to boot using the `eeprom` command. This example shows booting from the primary root disk.

```
# eeprom boot-device=vx-rootdg01
```

3 Restart the system.

```
# shutdown -g0 -i6 -y
```

Performing Boot Environment upgrade on Solaris 11 systems

Perform the Storage Foundation 6.1 BE upgrade manually or use the installer.

Table 15-4 Upgrading SF using BE upgrade

Step	Description
Step 1	Create a new BE on the primary boot disk. See "Creating a new Solaris 11 BE on the primary boot disk" on page 168.
Step 2	Upgrade SF using the installer. See "Upgrading SF using the installer for upgrading BE on Solaris 11" on page 168.
	To upgrade only Solaris See the Oracle documentation on Oracle Solaris 11 operating system.
Step 3	Switch the alternate BE to be the new primary. See "Completing the SF upgrade on BE on Solaris 11" on page 170.

Table 15-4 Upgrading SF using BE upgrade *(continued)*

Step	Description
Step 4	Verify BE upgrade of SF. See "Verifying Solaris 11 BE upgrade" on page 171.

Creating a new Solaris 11 BE on the primary boot disk

Run the `beadm create` command to create a new BE on the primary boot disk.

At the end of the process, a new BE is created on the primary boot disk by cloning the primary BE.

To create a new BE on the primary boot disk

1. View the list of BE in the primary disk.

```
# beadm list
```

2. Create a new BE in the primary boot disk.

```
# beadm create beName
```

```
# beadm mount beName mountpoint
```

If VVR is configured, it is recommended that <*beName*> should have the value *altroot.5.11* and <*mountpoint*> should have the value `/altroot.5.11`.

Upgrading SF using the installer for upgrading BE on Solaris 11

You can use the Symantec product installer to upgrade SF on a BE.

At the end of the process, the Storage Foundation 6.1 is installed on the alternate BE.

To perform BE upgrade of SF using the installer

1. Insert the product disc with Storage Foundation 6.1 or access your copy of the software on the network.
2. Run the installer script specifying the root path as the alternate BE:

```
# ./installsf -upgrade -rootpath /altroot.5.11
```

3 Enter the names of the nodes that you want to upgrade to Storage Foundation 6.1.

Note: Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the SF installation.

The installer displays the list of packages to be installed or upgraded on the nodes.

4 Press **Return** to continue with the installation.

5 Verify that the version of the Veritas packages on the alternate BE is 6.1.

```
# pkg -R /altroot.5.11 list VRTS\*
```

For example:

```
# pkg -R /altroot.5.1list VRTSvxvm
```

Review the installation logs at `/altroot.5.11/opt/VRTS/install/logs`.

Upgrading SF manually for Solaris 11 Live Upgrade

You can perform a manual upgrade of SF using Live Upgrade.

At the end of the process, theStorage Foundation 6.1 is installed on the alternate boot environment.

To perform Live Upgrade of SF manually

1 Remove the SF packages on the alternate boot environment in the reverse order in which they were installed:

- For Storage Foundation (SF):

```
# pkg -R /altroot.5.11 uninstall \
VRTSmapro VRTSgapms VRTSvxmsa VRTSfasag VRTSfas VRTSvail \
VRTSfsmnd VRTSfssdk VRTSfsman VRTSvrw VRTSweb VRTSjre15 \
VRTSvcsvr VRTSvrpro VRTSddlpr VRTSvdid VRTSalloc VRTSdcli \
VRTSvmpro VRTSvmman VRTSfspro VRTSdsa VRTSvxvm VRTSvxfs \
VRTSspt VRTSaa VRTSmh VRTSccg VRTSobgui VRTSob VRTSobc33 \
VRTSat VRTSpbx VRTSicsco VRTSvlic VRTSperl
```

Note that this package list is an example. Full package lists vary from release to release and by product option.

The `-R` option removes the packages from the root path `/altroot.5.11` on the alternate boot environment.

2. Set the package repository from the `/pkgs/VRTSpkgs.p5p` directory.

```
# pkg set-publisher -p VRTSpkgs.p5p Symantec
```

3. Install the SF packages from the repository. You must install the packages one at a time to the alternate boot environment using the `pkg install` command in the following order:

```
# pkg install --accept --no-backup-be
VRTSperl VRTSvlic VRTSspt VRTSvxvm VRTSaslapm VRTSvxfs VRTSfsadv
VRTSfssdk VRTSllt VRTSgab VRTSvxfen VRTSamf VRTSvcs VRTScps VRTSvcsag
VRTSvcsea VRTSdbed VRTSodm VRTSsfcpi61 VRTSsfmh VRTSvbs VRTSvcswiz
```

4. Verify that the version of the Veritas packages on the alternate boot environment is 6.1.

```
# pkg -R /altrootpath info VRTSpkgname
```

For example:

```
# pkg -R /altroot.5.11 info VRTSvxvm
```

5. Unset the repository.

```
# pkg unset-publisher Symantec
```

Completing the SF upgrade on BE on Solaris 11

At the end of the process:

- The alternate BE is activated.
- The system is booted from the alternate BE.

To complete the BE upgrade

1 Activate the alternate BE.

```
# beadm activate altroot.5.11
```

2 Restart the system. The BE on the alternate disk is activated when you restart it.

Note: Do not use the `reboot`, `halt`, or `uadmin` commands to restart the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate BE.

```
# shutdown -g0 -y -i6
```

3 After the alternate BE is activated, you can switch BEs. If the root disk is encapsulated, refer to the procedure to switch the BEs manually.

See "Administering boot environments in Solaris 10 Live Upgrade" on page 165.

4 After the upgrade, perform any required post-upgrade tasks such as upgrading the disk group.

5 After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

Verifying Solaris 11 BE upgrade

To ensure that BE upgrade has completed successfully, verify that the system have booted from the alternate BE.

To verify that BE upgrade completed successfully

1 Verify that the alternate BE is active.

```
# beadm list
```

If the alternate BE fails to be active, you can revert to the primary BE.

See "Reverting to the primary BE on a Solaris 11 system" on page 172.

2 Perform other verification as required to ensure that the new BE is configured correctly.

3 In a zone environment, verify the zone configuration.

If you have solaris10 brand zone on your system, you must manually upgrade the packages inside the solaris10 brand zone with packages from Solaris 10 install media.

If you have installed `VRTSvxfs` or `VRTSodm` packages inside the zones, you need to manually upgrade these packages inside the zone.

Administering BEs in Solaris 11 BE upgrade

Use the following procedures to perform relevant administrative tasks for BEs.

Switching the BE for Solaris SPARC

1. Display the status of BEs.

```
# beadm list

Boot Emvironment Active Mountpoint Space  Policy Created
--               ------ ---------  -----  ------ -------
solaris          NR        /       13.08G static 2012-11-14 10:22
altroot.5.11     -         -        3.68G static 2013-01-06 18:41
```

   In this example, the primary boot disk is currently *solaris*. You want to activate the alternate boot disk *altroot.5.11*.

2. Activate the BE.

```
# beadm activate altroot.5.11
```

3. Restart the system to complete the BE activation.

```
# shutdown -g0 -i6 -y
```

   The system automatically selects the BE entry that was activated.

4. You can destroy an existing BE.

```
# beadm destroy altroot.5.11
```

Reverting to the primary BE on a Solaris 11 system

Boot the system to `ok` prompt.

View the available BEs.

To view the BEs, enter the following:

```
ok> boot -L
```

Select the option of the original BE to which you need to boot.

To boot to the BE

```
# boot -Z <path to boot env>
```

For example:

```
{0} ok boot -L
Boot device: /virtual-devices@100/channel-devices@200/disk@0:a
File and args: -L
1 Oracle Solaris 11 11/11 SPARC
2 solaris-backup-1
Select environment to boot: [ 1 - 2 ]: 1
```

To boot the selected entry, enter the following:

```
boot [<root-device>] -Z rpool/ROOT/solaris

Program terminated
{0} ok boot -Z rpool/ROOT/solaris
```

About Live Upgrade in a Volume Replicator (VVR) environment

This section provides an overview of the VVR upgrade process.

In an SF environment that uses Volume Replicator, the following scripts provide the means to upgrade the VVR configuration:

- vvr_upgrade_lu_start
- vvr_upgrade_lu_finish

The scripts are available in the scripts directory in the install media.

- Immediately before restarting the system to switch over to the alternate boot environment, run the `vvr_upgrade_lu_start` script.

  Note: Use the `vvr_upgrade_lu_start` script only when the applications are stopped and the next step is to switch over to the alternate boot environment.

- After the `vvr_upgrade_lu_start` script completes successfully, restart the system. This restart results in the system booting from the alternate boot environment.
- After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

Chapter 16

Migrating from Storage Foundation Basic to Storage Foundation Standard

This chapter includes the following topics:

- Migrating from Storage Foundation Basic to Storage Foundation Standard

Migrating from Storage Foundation Basic to Storage Foundation Standard

Use this procedure to migrate from Storage Foundation (SF) Basic to Storage Foundation Standard.

To migrate from Storage Foundation Basic to Storage Foundation Standard

1. Log in as superuser on a system where you want to install Storage Foundation Standard.
2. Use the following command to confirm that you are currently running Storage Foundation Basic.

   ```
   # /opt/VRTSvlic/bin/vxlicrep | grep  Basic
   ```

   You should see the following output:

   ```
   Product Name = VERITAS Storage Foundation Basic
   ```

3. Mount the installation media for Storage Foundation Standard.
4. Run the installer command.

   ```
   # ./installer
   ```

   The installer will first execute a set of prechecks.
5. Make sure that the prechecks complete successfully. Make any changes that the installer recommends.
6. On the Installer Task menu, select **Install a Product**.
7. On the Product Selection menu, select **Symantec Storage Foundation**.
8. At the prompt, specify whether you accept the terms of the End User License Agreement (EULA). Press **y** to agree and continue.
9. Select the package level to be installed.

   From the menu, select the option that corresponds to **Install Recommended**.
10. You are prompted to enter the system names on which to install Storage Foundation Standard.
11. The installer prompts with a warning that Storage Foundation is already installed, and asks for confirmation to continue. Press **y** to continue the installation.
12. The installer will identify two additional packages to be installed, `VRTSodm` and `VRTSdbed`. Press Enter to continue.
13. After installing the packages in step 12, the installer will prompt if additional licenses need to be installed. Press **y** to continue.

    You will be provided two options:

    - `Enter a license key`
    - `Utilize Keyless licensing`
14. If you chose to enter a license key, you should install the Storage Foundation Standard license key.
15. If you chose to utilize a keyless license, you will be asked to choose the version of Storage Foundation (Standard or Enterprise). Choose **Standard** to install a Storage Foundation Standard license.

    The installer will go through the configuration and startup process.

16 Confirm if you want to send information about this installation to Symantec to help improve the installation in the future.

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

17 If desired, press **y** to view the summary file.

The migration is complete.

Chapter 17

Performing post-upgrade tasks

This chapter includes the following topics:

- Optional configuration steps
- Re-joining the backup boot disk group into the current disk group
- Reverting to the backup boot disk group after an unsuccessful upgrade
- Post upgrade tasks for migrating the SFDB repository database
- Recovering VVR if automatic upgrade fails
- Upgrading disk layout versions
- Upgrading VxVM disk group versions
- Updating variables
- Setting the default disk group
- Upgrading the Array Support Library
- 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 Volume Replicator (VVR) is configured, do the following steps in the order shown:
  - Reattach the RLINKs.
  - Associate the SRL.
- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Symantec Storage Foundation Administrator's Guide*.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
  See "Upgrading VxVM disk group versions" on page 186.

Re-joining the backup boot disk group into the current disk group

Perform this procedure to rejoin the backup boot disk if you split the mirrored boot disk during upgrade. After a successful upgrade and reboot, you no longer need to keep the boot disk group backup.

To re-join the backup boot disk group

- Re-join the *backup_bootdg* disk group to the boot disk group.

  ```
  # /etc/vx/bin/vxrootadm -Y join backup_bootdg
  ```

  where the `-Y` option indicates a silent operation, and *backup_bootdg* is the name of the backup boot disk group that you created during the upgrade.

Reverting to the backup boot disk group after an unsuccessful upgrade

Perform this procedure if your upgrade was unsuccessful and you split the mirrored boot disk to back it up during upgrade. You can revert to the backup that you created when you upgraded.

To revert the backup boot disk group after an unsuccessful upgrade

1. To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

   ```
   # vxprint
   ```

2. Use the `vxdg` command to find the boot disk group where you are currently booted.

   ```
   # vxdg bootdg
   ```

3. Boot the operating system from the backup boot disk group.
4. Join the original boot disk group to the backup disk group.

   ```
   # /etc/vx/bin/vxrootadm -Y join original_bootdg
   ```

   where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

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 6.1:

- 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 184.
- Migrate from a 5.0x SFDB repository database to 6.1
  See “Migrating from a 5.0 repository database to 6.1” on page 180.
- Migrate from a 5.1 or 5.1SP1 repository database to 6.1
  See “Migrating from a 5.1 or higher repository database to 6.1” on page 182.

Migrating from a 5.0 repository database to 6.1

To migrate from a 5.0 repository database to 6.1

1. Rename the startup script NO_S*vxdbms3 to S*vxdbms3.

   See "After upgrading from 5.0.x and before migrating SFDB" on page 184.

2. 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
   ```

3. 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
   ```

4. 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
   ```

5. 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.

6 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
```

7 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.

Note: While you revalidate the snapshot configuration file (`snapplan`) from an older release, use the `vxsfadm -c <configfile>` option to avoid the default values from overriding the old values.

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

see *Storage Foundation: Storage and Availability Management for Oracle Databases*.

Migrating from a 5.1 or higher repository database to 6.1

To migrate from a 5.1 or higher repository database to 6.1

1 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
```

2 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
```

3 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.

4 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
```

5 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.

Note: While you revalidate the snapshot configuration file (`snapplan`) from an older release, use the `vxsfadm -c <configfile>` option to avoid the default values from overriding the old values.

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 to SF 6.1 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:

```
/etc/rc.d/rc2.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.

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
# adddcm
# 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 7, 8, 9, and 10. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Note: If you plan to use 64-bit quotas, you must upgrade to the latest disk layout Version 10. The use of 64-bit quota on earlier disk layout versions is deprecated in this release.

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.

To upgrade the disk layout versions

- To get to disk layout Version 10 from Version 6. You must incrementally upgrade the disk layout of this file system. For example:

```
# vxupgrade -n 7 /mnt
# vxupgrade -n 8 /mnt
# vxupgrade -n 9 /mnt
# vxupgrade -n 10 /mnt
```

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

Support for disk layout Version 4 and 5 has been removed. You must upgrade any existing file systems with disk layout Version 4 or 5 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. Once a disk layout version has been upgraded, it is not possible to downgrade to the previous version.

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 *Symantec Storage Foundation Administrator's Guide*.

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 6.1, 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 6.1, 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.

For more information about ISP disk groups, refer to the *Symantec Storage Foundation Administrator's Guide*.

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 *Symantec Storage Foundation Administrator's Guide*.

Updating variables

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

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

Setting the default disk group

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 *Symantec Storage Foundation Administrator's Guide*.

Upgrading the Array Support Library

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software package.

Adding JBOD support for storage arrays for which there is not an ASL available

If an array is of type A/A-A, A/P or ALUA and a suitable ASL is not available, the array must be claimed as a JBOD of type A/P. This is to prevent path delays and I/O failures arising. As JBODs are assumed to be type A/A by default, you must create appropriate JBOD entries for such arrays.

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

1. Stop all applications, such as databases, from accessing the VxVM volumes that are configured on the array, and unmount all VxFS file systems and Storage Checkpoints that are configured on the array.
2. Add the array as a JBOD of type A/P:

   ```
   # vxddladm addjbod vid=SUN pid=T300 policy=ap
   ```

3. If you have not already done so, upgrade the Storage Foundation or VxVM software to 6.1. Device discovery is performed during the upgrade, and the array is claimed as a JBOD of appropriate type.

   If you have already upgraded your system to 6.1, run the following command to perform device discovery:

   ```
   # vxdctl enable
   ```

4 Verify that the array has been added with the policy set to `APdisk`:

```
# vxddladm listjbod
VID    PID      Opcode Page Code Page Offset SNO length Policy
==========================================================
SUN    T300     18     -1        36          12         APdisk
```

5 Check that the correct devices are listed for the array:

```
# vxdisk list
DEVICE      TYPE           DISK     GROUP    STATUS
APdisk_0    auto:cdsdisk   -        -        online invalid
APdisk_1    auto:cdsdisk   -        -        online invalid
APdisk_2    auto:cdsdisk   -        -        online invalid
...
```

Unsuppressing DMP for EMC PowerPath disks

This section is only applicable if you want to upgrade a system that includes EMC PowerPath disks.

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

VxVM 4.1 and later releases have the ability to discover EMCpower disks, and configure them as autodiscovered disks that DMP recognizes are under the control of a separate multi-pathing driver. This has the benefit of allowing such disks to reconfigured in cluster-shareable disk groups. Before upgrading to VxVM 6.1, you must remove the suppression of the subpaths and controllers so that DMP can determine the association between EMCpower metadevices and `c#t#d#` disk devices.

In the following scenarios, you may need to unsuppress DMP subpaths and controllers:

- Converting a foreign disk
  See "Converting a foreign disk to auto:simple" on page 189.
- Converting a defined disk
  See "Converting a defined disk to auto:simple" on page 191.
- Converting a powervxvm disk
  See "Converting a powervxvm disk to auto:simple" on page 194.

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

```
# powervxvm remove
```

Converting a foreign disk to auto:simple

Release 4.0 of VxVM provides the `vxddladm addforeign` command to configure foreign disks with default disk offsets for the private regions and public regions, and to define them as simple disks. A foreign disk must be manually converted to `auto:simple` format before you upgrade to VxVM 6.1.

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

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

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

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

```
# vxprint
Disk group: fdg
TY NAME    ASSOC        KSTATE  LENGTH    PLOFFS  STATE TUTIL0 PUTIL0
dg fdg     fdg          -       -         -       -     -      -
dm fdisk   emcpower10c  -       17673456  -       -     -      -
...
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5 Upgrade to VxVM 6.1 using the appropriate upgrade procedure.

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

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

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

```
# vxdmpadm getsubpaths dmpnodename=emcpower10s2
```

7 Import the disk group and start the volumes:

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

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

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

Converting a defined disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (`darec`), and identified as simple disks. If an EMCpower disk is defined with a persistent `darec`, it must be manually converted to `auto:simple` format before upgrading to VxVM 6.1.

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

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

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

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

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

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

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

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

```
# vxprint
Disk group: fdg
TY NAME    ASSOC     KSTATE  LENGTH    PLOFFS  STATE TUTIL0 PUTIL0
dg fdg     fdg       -       -         -       -     -      -
dm fdisk   emcdisk1  -       17673456  -       -     -      -
...
```

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

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

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

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

```
# vxdisk rm emcdisk1
```

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

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

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

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

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

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

# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE     TAG  FLAGS    START   SIZE
  4         0x0  0x200    0       0
  5         0x0  0x200    3591000 2100375
  6         0x0  0x200    0       0

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE     TAG  FLAGS    START   SIZE
  4         0x0  0x200    0       0
  5         0xf  0x200    3591000 2100375
  6         0x0  0x200    0       0

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

5 Upgrade to VxVM 6.1 using the appropriate upgrade procedure.

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

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

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

```
# vxdmpadm getsubpaths dmpnodename=emcpower10s2
```

7 Import the disk group and start the volumes:

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

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

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

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

Converting a powervxvm disk to auto:simple

In VxVM 4.0, and particularly in previous releases, EMCpower disks can be defined by a persistent disk access record (darec) using `powervxvm` script, and identified as simple disks. If an EMCpower disk is used using `powervxvm`, it must be manually converted to auto:simple format before you upgrade to VxVM 6.1.

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

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

```
# ls -l /dev/vx/rdmp/
crw-------   1 root     root     260, 76 Feb  7 02:36 emcpower0c

# vxdisk list
DEVICE       TYPE            DISK         GROUP        STATUS
c6t0d12s2    auto:sliced       -            -          online
emcpower0c   simple          ppdsk01      ppdg         online

# vxprint
Disk group: fdg
TY NAME     ASSOC      KSTATE  LENGTH   PLOFFS STATE TUTIL0 PUTIL0
dg ppdg     ppdg       -       -        -      -     -      -
dm ppdsk01  emcpower0c -       2094960  -      -     -      -
```

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

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

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

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

   ```
   # vxdisk rm emcpower0c
   ```

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

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

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

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

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

```
# /etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE     TAG  FLAGS   START  SIZE
  0         0x0  0x201   0      0
  1         0x0  0x200   0      0
  2         0x5  0x201   0      17675520

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

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

5 Upgrade to VxVM 6.1 using the appropriate upgrade procedure.

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

```
# vxdisk list
DEVICE       TYPE            DISK         GROUP        STATUS
c6t0d12s2    auto:sliced       -            -          online
emcpower0s2  auto:simple       -            -          online
```

7 Import the disk group and start the volumes.

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

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

Converting from QuickLog to Multi-Volume support

The 4.1 release of the Veritas File System is the last major release to support QuickLog. The Version 6 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.

   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 202.

Section 5

Post-installation tasks

- Chapter 18. Performing post-installation tasks
- Chapter 19. Verifying the SF installation

Chapter 18

Performing post-installation tasks

This chapter includes the following topics:

- Changing root user into root role
- Switching on Quotas
- Enabling DMP support for native devices

Changing root user into root role

On Oracle Solaris 11, you need to create root user to perform installation. This means that a local user cannot assume the root role. After installation, you may want to turn root user into root role for a local user, who can log in as root.

1. Log in as root user.
2. Change the root account into role.

```
# rolemod -K type=role root

# getent user_attr root

root::::type=role;auths=solaris.*;profiles=All;audit_flags=lo\
:no;lock_after_retries=no;min_label=admin_low;clearance=admin_high
```

3. Assign the root role to a local user who was unassigned the role.

```
# usermod -R root admin
```

For more information, see the Oracle documentation on Oracle Solaris 11 operating system.

Switching on Quotas

This turns on the group and user quotas once all the nodes are upgraded to 6.1, if it was turned off earlier.

To turn on the group and user quotas

- Switch on quotas:

```
# vxquotaon -av
```

Enabling DMP support for native devices

Symantec Dynamic Multi-Pathing (DMP) is a component of SF. DMP supports Veritas Volume Manager (VxVM) volumes on DMP metadevices, and Veritas File System (VxFS) file systems on those volumes.

DMP can also provide multi-pathing functionality for the native operating system volumes and file systems on DMP devices.

For more information on using DMP with native devices, see the *Symantec Dynamic Multi-Pathing Administrator's Guide*.

After you install SF for the first time, use the following procedure to enable DMP support for native devices.

If DMP native support for native devices is enabled on a system before you upgrade SF, DMP native support is maintained when SF is upgraded.

Starting with Solaris 11.1, enabling DMP support for native devices also enables support for ZFS root on DMP devices. If DMP native support is enabled with an earlier Solaris version, ZFS root devices are not supported on DMP. Upgrading the operating system to version 11.1 or later does not enable support for ZFS root devices by default. To enable DMP support for the ZFS root devices, use the following procedure to enable DMP support for native devices again.

To enable DMP support for native devices

1. Turn on the tunable parameter to enable DMP support:

   ```
   # vxdmpadm set dmp_native_support=on
   ```

   The `dmp_native_support` parameter is persistent.

2. If the system has Solaris version 11.1 or later installed, turning on DMP support also enables support for the ZFS root device. Reboot the system for the changes to take effect.

Chapter 19

Verifying the SF installation

This chapter includes the following topics:

- Verifying that the products were installed
- Installation log files
- Starting and stopping processes for the Symantec 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 `pkginfo` (Solaris 10) or `pkg info` (Solaris 11) command to check which packages have been installed.

Solaris 10:

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

Solaris 11:

```
# pkg info -l VRTSvlic package_name package_name
```

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

```
# /opt/VRTS/install/installsf<version>
```

Where `<version>` is the specific release version.

See “About the script-based installer” on page 46.

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 Symantec 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 packages, and the status (success or failure) of each package. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

Starting and stopping processes for the Symantec products

After the installation and configuration is complete, the Symantec product installer starts the processes that the installed products use. 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
  ```

  or

  ```
  # /opt/VRTS/install/installsf<version> -stop
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 46.

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
  ```

  or

  ```
  # /opt/VRTS/install/installsf<version> -start
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 46.

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

  For more details on hot relocation, see *Symantec Storage Foundation Administrator's Guide*..

Checking Veritas File System installation

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

Verifying Veritas File System kernel installation

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

```
# modinfo | grep vxfs
```

The `modinfo` command displays information about all modules loaded on the system. If the `vxfs` module is loaded, you will see an entry corresponding to `vxfs`. If not, follow the instructions load and then unload the file system module to complete the process.

See "Loading and unloading the file system module" on page 86.

Verifying command installation

Table 19-1 lists the directories with Veritas File System commands.

Table 19-1 VxFS command locations

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

Determine whether these subdirectories are present:

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

Make sure you have adjusted the environment variables accordingly.

See "Setting environment variables" on page 42.

Section 6

Uninstallation of SF

- Chapter 20. Uninstalling Storage Foundation
- Chapter 21. Uninstalling SF using response files

Chapter 20

Uninstalling Storage Foundation

This chapter includes the following topics:

- About removing Storage Foundation
- Preparing to uninstall
- Removing the Replicated Data Set
- Uninstalling SF packages using the script-based installer
- Uninstalling SF with the web-based installer
- Uninstalling Storage Foundation using the pkgrm or pkg uninstall command
- Manually uninstalling Storage Foundation packages on non-global zones on Solaris 11
- Removing the Storage Foundation for Databases (SFDB) repository after removing the product

About removing Storage Foundation

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

Only users with superuser privileges can uninstall Storage Foundation.

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

Preparing to uninstall

Review the following before removing the Veritas software.

Preparing to remove Veritas Volume Manager

This section describes the steps you need to take before you remove Veritas Volume Manager (VxVM) to preserve the contents of the volumes.

Warning: Failure to follow the preparations in this section might result in unexpected behavior.

On Solaris 11, the SMF service `vxvm-configure` must be online in order to uninstall VRTSvxvm successfully.

To verify that the vxvm-configure service is online

1. Check the state of the `vxvm-configure` service:

   ```
   # svcs -a | grep vxvm-configure
   ```

2. If the service is in disabled or maintenance state, use the following command to display information including the service log location:

   ```
   # svcs -xv vxvm-configure
   ```

3. If there are no issues, use the following ocmmand to bring the vxvm-configure service online:

   ```
   # svcadm enable vxvm-configure
   ```

Moving volumes from an encapsulated root disk

Use the following procedure to move volumes from an encapsulated root disk.

To uninstall VxVM if `root`, `swap`, `usr`, or `var` is a volume under Volume Manager control

1. Ensure that the `rootvol`, `swapvol`, `usr`, and `var` volumes have only one associated plex each.

   The plex must be contiguous, non-striped, non-spanned, and non-sparse. To obtain this information, enter the following:

   ```
   # vxprint -ht rootvol swapvol usr var
   ```

   If any of these volumes have more than one associated plex, remove the unnecessary plexes using the following command:

   ```
   # vxplex -g diskgroup -o rm dis plex_name
   ```

2. Run the `vxunroot` command:

   ```
   # /etc/vx/bin/vxunroot
   ```

   The `vxunroot` command changes the volume entries in `/etc/vfstab` to the underlying disk partitions for `rootvol`, `swapvol`, `usr`, and `var`. It also modifies `/etc/system` and prompts for a restart so that disk partitions are mounted instead of volumes for `root`, `swap`, `usr`, and `var`.

3. Once you have changed the `root`, `swap`, `usr`, and `var` volumes, move all remaining volumes to disk partitions.

   You can do this using one of the following procedures:

   - Back up the entire system to tape and then recover from tape.
   - Back up each file system individually and then recover them all after you create new file systems on disk partitions.
   - Move volumes incrementally to disk partitions.
     See "Moving volumes to disk partitions" on page 209.
     Otherwise, shut down VxVM.

Moving volumes to disk partitions

Use the following procedure to move volumes incrementally to disk partitions.

To move volumes incrementally to disk partitions

1. Evacuate disks using the `vxdiskadm` command, the VOM GUI, or the `vxevac` utility.

   Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.

2. Remove the evacuated disks from VxVM control by entering:

   ```
   # vxdg rmdisk diskname
   # vxdisk rm devname
   ```

3. Decide which volume to move first, and if the volume is mounted, unmount it.

4. If the volume is used as a raw partition for database applications, make sure that the application does not update the volume. Also make sure that you have applied the `sync` command to the data on the volume.

5. Create a partition on free disk space of the same size as the volume using the `format` command.

   If there is not enough free space for the partition, add a new disk to the system for the first volume removed. Subsequent volumes can use the free space that the removal of this first volume generates.

6. Copy the data on the volume onto the newly created disk partition using a command such as `dd`.

   ```
   # dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/dsk/c2t2d2s7
   ```

   where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly created partition.

7. Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.

8. Mount the disk partition if the corresponding volume was previously mounted.

9. Stop and remove the volume from VxVM using the commands.

   ```
   # vxvol -g diskgroup stop volume_name
   # vxedit -rf -g diskgroup rm volume_name
   ```

10 Remove any free disks (those disks that have no subdisks defined on them) by removing the volumes from VxVM control.

To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.

```
# vxprint -g diskgroup -F '%sdnum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VxVM control.

```
# vxdg rmdisk diskname
# vxdisk rm devname
```

Use the free space that is created for adding the data from the next volume you want to remove.

11 After you successfully convert all volumes into disk partitions, restart the system.

12 After the restart, make sure that none of the volumes are open by using the `vxprint` command.

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

13 If any volumes remain open, repeat the steps.

Example of moving volumes to disk partitions on Solaris

This example shows how to move the data on a volume to a disk partition. In the example, there are three disks: `disk1` and `disk2` are subdisks on volume `vol01` and `disk3` is a free disk. The data on `vol01` is copied to `disk3` using the `vxevac` command.

These are the contents of the disk group `voldg` before the data on `vol01` is copied to `disk3`.

```
# vxprint -g voldg -ht
DG NAME   NCONFIG    NLOG     MINORS    GROUP-ID
DM NAME   DEVICE     TYPE     PRIVLEN   PUBLEN    STATE
RV NAME   RLINK_CNT KSTATE   STATE     PRIMARY   DATAVOLS  SRL
RL NAME   RVG        KSTATE   STATE     REM_HOST REM_DG    REM_RLNK
V  NAME   RVG        KSTATE   STATE     LENGTH    READPOL   PREFPLEX UTYPE
PL NAME   VOLUME     KSTATE   STATE     LENGTH    LAYOUT    NCOL/WID MODE
SD NAME   PLEX       DISK     DISKOFFS LENGTH    [COL/]OFF DEVICE   MODE
SV NAME   PLEX       VOLNAME NVOLLAYR LENGTH    [COL/]OFF AM/NM    MODE
DC NAME   PARENTVOL LOGVOL
SP NAME   SNAPVOL    DCO
```

```
dg voldg default   default 115000
1017856044.1141.hostname.veritas.com

dm disk1 c1t12d0s2 sliced  2591     17900352 -
dm disk2 c1t14d0s2 sliced  2591     17899056 -
dm disk3 c1t3d0s2  sliced  2591     17899056 -

v  vol1  -         ENABLED ACTIVE   4196448  ROUND     -        fsgen
pl pl1   vol1      ENABLED ACTIVE   4196448  CONCAT    -        RW
sd sd1   pl1       disk1   0        2098224  0         c1t12d0  ENA
sd sd2   pl1       disk2   0        2098224  2098224   c1t14d0  ENA
```

Evacuate `disk1` to `disk3`.

```
# /etc/vx/bin/vxevac -g voldg disk1 disk3
# vxprint -g voldg -ht

DG NAME  NCONFIG    NLOG     MINORS   GROUP-ID
DM NAME  DEVICE     TYPE     PRIVLEN  PUBLEN   STATE
RV NAME  RLINK_CNT  KSTATE   STATE    PRIMARY  DATAVOLS  SRL
RL NAME  RVG        KSTATE   STATE    REM_HOST REM_DG    REM_RLNK
V  NAME  RVG        KSTATE   STATE    LENGTH   READPOL   PREFPLEX  UTYPE
PL NAME  VOLUME     KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID  MODE
SD NAME  PLEX       DISK     DISKOFFS LENGTH   [COL/]OFF DEVICE    MODE
SV NAME  PLEX       VOLNAME  NVOLLAYR LENGTH   [COL/]OFF AM/NM     MODE
DC NAME  PARENTVOL  LOGVOL
SP NAME  SNAPVOL    DCO

dg voldg default    default  115000
1017856044.1141.hostname.veritas.com

dm disk1 c1t12d0s2  sliced   2591     17900352 -
dm disk2 c1t14d0s2  sliced   2591     17899056 -
dm disk3 c1t3d0s2   sliced   2591     17899056 -

v  vol1  -          ENABLED  ACTIVE   4196448  ROUND     -         fsgen
pl pl1   vol1       ENABLED  ACTIVE   4196448  CONCAT    -         RW
sd disk3-01l1       disk3    0        2098224  0         c1t3d0    ENA
sd sd2   pl1        disk2    0        2098224  2098224   c1t14d0   ENA
```

Evacuate `disk2` to `disk3`.

```
# /etc/vx/bin/vxevac -g voldg disk2 disk3
# vxprint -g voldg -ht

DG NAME     NCONFIG   NLOG    MINORS   GROUP-ID
DM NAME     DEVICE    TYPE    PRIVLEN  PUBLEN   STATE
RV NAME     RLINK_CNT KSTATE  STATE    PRIMARY  DATAVOLS  SRL
RL NAME     RVG       KSTATE  STATE    REM_HOST REM_DG    REM_RLNK
V  NAME     RVG       KSTATE  STATE    LENGTH   READPOL   PREFPLEX UTYPE
PL NAME     VOLUME    KSTATE  STATE    LENGTH   LAYOUT    NCOL/WID MODE
SD NAME     PLEX      DISK    DISKOFFS LENGTH   [COL/]OFF DEVICE   MODE
SV NAME     PLEX      VOLNAME NVOLLAYR LENGTH   [COL/]OFF AM/NM    MODE
DC NAME     PARENTVOL LOGVOL
SP NAME     SNAPVOL   DCO

dg voldg    default   default 115000
1017856044.1141.hostname.veritas.com

dm disk1    c1t12d0s2 sliced  2591    17900352 -
dm disk2    c1t14d0s2 sliced  2591    17899056 -
dm disk3    c1t3d0s2  sliced  2591    17899056 -

v  vol1     -         ENABLED ACTIVE  4196448  ROUND   -      fsgen
pl pl1      vol1      ENABLED ACTIVE  4196448  CONCAT  -      RW
sd disk3-01 pl1       disk3   0       2098224  0       c1t3d0 ENA
sd disk3-02 pl1       disk3   2098224 2098224  2098224 c1t3d0 ENA
```

Remove the evacuated disks from VxVM control.

```
# vxdisk -g voldg list
DEVICE      TYPE      DISK         GROUP        STATUS
c1t3d0s2    sliced    disk3        voldg        online
c1t12d0s2   sliced    disk1        voldg        online
c1t14d0s2   sliced    disk2        voldg        online

# vxdg rmdisk disk1
# vxdg rmdisk disk2
# vxdisk rm c1t12d0
# vxdisk rm c1t14d0
```

Verify that the evacuated disks have been removed from VxVM control.

```
# vxdisk -g voldg list
DEVICE       TYPE      DISK         GROUP        STATUS
c1t3d0s2     sliced    disk3        voldg        online
```

Check to see whether the volume you want to move first is mounted.

```
# mount | grep vol1
/vol1 on /dev/vx/dsk/voldg/vol1
read/write/setuid/log/nolargefiles/dev=12dc138 on Wed Apr
3 10:13:11 2002
```

Create a partition on free disk space of the same size as the volume. In this example, a 2G partition is created on `disk1` (`c1t12d0s1`).

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
     0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
        /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
     1. c1t3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
        /sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
     2. c1t9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
        /sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
     3. c1t10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
        /sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
     4. c1t11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
        /sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
     5. c1t12d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
        /sbus@1f,0/SUNW,fas@2,8800000/sd@c,0
     6. c1t14d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
        /sbus@1f,0/SUNW,fas@2,8800000/sd@e,0
     7. c1t15d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
        /sbus@1f,0/SUNW,fas@2,8800000/sd@f,0
Specify disk (enter its number): 5
selecting c1t12d0
[disk formatted]

FORMAT MENU:
        disk       - select a disk
        type       - select (define) a disk type
        partition  - select (define) a partition table
        current    - describe the current disk
        format     - format and analyze the disk
        repair     - repair a defective sector
        label      - write label to the disk
        analyze    - surface analysis
        defect     - defect list management
```

```
        backup     - search for backup labels
        verify     - read and display labels
        save       - save new disk/partition definitions
        inquiry    - show vendor, product and revision
        volname    - set 8-character volume name
        !<cmd>     - execute <cmd>, then return
        quit
format> p

PARTITION MENU:
        0      - change '0' partition
        1      - change '1' partition
        2      - change '2' partition
        3      - change '3' partition
        4      - change '4' partition
        5      - change '5' partition
        6      - change '6' partition
        7      - change '7' partition
        select - select a predefined table
        modify - modify a predefined partition table
        name   - name the current table
        print  - display the current table
        label  - write partition map and label to the disk
        !<cmd> - execute <cmd>, then return
        quit
partition> 1
Part      Tag    Flag    Cylinders     Size          Blocks
  1 unassigned    wm      0            0       (0/0/0)            0
Enter partition id tag[unassigned]:
Enter partition permission flags[wm]:
Enter new starting cyl[0]:
Enter partition size[0b, 0c, 0.00mb, 0.00gb]: 2.00gb
partition> 1
Ready to label disk, continue? y

partition> p
Current partition table (unnamed):
Total disk cylinders available: 13814 + 2 (reserved cylinders)
Part      Tag    Flag    Cylinders     Size          Blocks
  0 unassigned    wm      0            0       (0/0/0)            0
  1 unassigned    wm      0 -  3236    2.00GB  (3237/0/0)   4195152
partition> q
```

Copy the data on `vol01` to the newly created disk partition.

```
# dd if=/dev/vx/dsk/voldg/vol01 of=/dev/dsk/c1t12d0s1
```

In the `/etc/vfstab` file, remove the following entry.

```
/dev/vx/dsk/voldg/vol1 /dev/vx/rdsk/voldg/vol1 /vol1 vxfs 4 yes rw
```

Replace it with an entry for the newly created partition.

```
/dev/dsk/c1t12d0s1 /dev/rdsk/c1t12d0s1 /vol01 vxfs 4  yes rw
```

Mount the disk partition.

```
# mount -F vxfs /dev/dsk/c1t12d0s1 /vol01
```

Remove `vol01` from VxVM.

```
# vxedit -rf -g voldg rm /dev/vx/dsk/voldg/vol01
```

To complete the procedure, follow the remaining steps.

Preparing to remove Veritas File System

The `VRTSvxfs` package cannot be removed if there are any mounted VxFS file systems or Storage Checkpoints. Unmount the VxFS file systems and Storage Checkpoints before uninstalling Symantec Storage Foundation. After you remove the `VRTSvxfs` package, VxFS file systems are not mountable or accessible until another `VRTSvxfs` package is installed.

To unmount a file system

1. Check if any VxFS file systems are mounted.

   ```
   # cat /etc/mnttab | grep vxfs
   ```

2. Unmount any file systems.

   ```
   # umount special | mount_point
   ```

   Specify the file system to be unmounted as a *mount_point* or *special* (the device on which the file system resides). See the `umount_vxfs`(1M) manual page for more information about this command and its available options.

   You can use the `-a` option to unmount all file systems except `/`, `/usr`, `/usr/kvm`, `/var`, `/proc`, `/dev/fd`, and `/tmp`.

To unmount a Storage Checkpoint

1 Check if any Storage Checkpoints are mounted.

```
# cat /etc/mnttab | grep vxfs
```

2 Unmount any Storage Checkpoints.

```
# umount /checkpoint_name
```

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 Volume Replicator, do not remove the Replicated Data Set.

To remove the Replicated Data Set

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

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

If the Secondary is not required to be up-to-date, proceed to 2 and stop replication using the `-f` option with the `vradmin stoprep` command.

2 Stop replication to the Secondary by issuing the following command on any host in the RDS:

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

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

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

3 Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

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

The argument local_rvgname is the name of the RVG on the local host and represents its RDS.

The argument sec_hostname is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

4 Remove the Primary from the RDS by issuing the following command on the Primary:

```
# vradmin -g diskgroup delpri local_rvgname
```

When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

The RDS is removed.

5 If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

```
# vxedit -r -g diskgroup rm srl_name
```

Uninstalling SF packages using the script-based installer

Use the following procedure to remove SF products.

Not all packages 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 6.1 with a previous version of SF.

Language packages are uninstalled when you uninstall the English language packages.

To shut down and remove the installed SF packages

1. Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.
2. Unmount all mount points for VxFS file systems.

   ```
   # umount /mount_point
   ```

3. If the VxVM package (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

   See "Preparing to remove Veritas Volume Manager" on page 208.
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
   ```

   ```
   # ./uninstallsf<version>
   ```

   Where `<version>` is the specific release version.

   Or, if you are using rsh, use the following:

   ```
   # ./uninstallsf<version> -rsh
   ```

   See "About the script-based installer" on page 46.
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, `sys1`:

   ```
   Enter the system names separated by spaces: [q?] sys1 sys2
   ```

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

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

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

9 To verify the removal of the packages, use the following commands:

Solaris 10:

```
# pkginfo | grep VRTS
```

Solaris 11:

```
# pkg list VRTS\*
```

Uninstalling SF with the web-based installer

This section describes how to uninstall using the 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 6.1 with a previous version of SF.

To uninstall SF

1 Perform the required steps to save any data that you want to preserve. For example, take backups of configuration files.

2 Start the web-based installer.

See "Starting the web-based installer" on page 52.

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**.

Most packages have kernel components. To ensure their complete removal, a system restart is recommended after all the packages have been removed.

Uninstalling Storage Foundation using the `pkgrm` or `pkg uninstall` command

Use the following procedure to uninstall Storage Foundation using the `pkgrm` command.

If you want to uninstall Storage Foundation using the `pkgrm` command, the packages must be removed in a specific order, or else the uninstallation fails. Removing the packages out of order results in some errors, including possible core dumps, although the packages are still removed.

To uninstall Storage Foundation

1 Unmount all mount points for file systems and Storage Checkpoints.

```
# umount /mount_point
```

Note: Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries can result in system boot problems later.

2 Stop all applications from accessing VxVM volumes, and close all volumes.

3 For Solaris 11.1 or later, if DMP native support is enabled, DMP controls the ZFS root pool. Turn off native support before removing Storage Foundation.

```
# vxdmpadm settune dmp_native_support=off
```

Note: If you do not disable native support, the system cannot be restarted after you remove DMP.

4 Stop any Veritas daemons that are running.

5 Remove the packages in the following order:

- For Storage Foundation (Solaris 10):

```
# pkgrm VRTSodm VRTSdbed VRTSfssdk \
VRTSfsadv VRTSvxfs VRTSsfmh VRTSob VRTSaslapm VRTSvxvm \
VRTSspt VRTSperl VRTSvlic VRTSsfcpi61
```

- For Storage Foundation (Solaris 11):

```
# pkg uninstall VRTSodm VRTSdbed VRTSfssdk VRTSfsadv\
VRTSvxfs VRTSsfmh VRTSob VRTSaslapm VRTSvxvm \
VRTSspt VRTSperl VRTSvlic VRTSsfcpi61
```

Uninstalling the language packages using the `pkgrm` command

If you want to remove only the language packages, you can do so with the `pkgrm` command.

If you use the product installer menu or the uninstallation script, you can remove the language packages along with the English packages.

To remove the language packages

- Use the `pkgrm` command to remove the appropriate packages.

  See "Chinese language packages" on page 259.

  See "Japanese language packages" on page 259.

  ```
  # pkgrm package_name package_name ...
  ```

  Because the packages do not contain any dependencies, you can remove them in any order.

Manually uninstalling Storage Foundation packages on non-global zones on Solaris 11

1. Log on to the non-global zone as a super user.
2. Uninstall SF packages from Solaris brand zones.

   ```
   # pkg uninstall VRTSperl VRTSvlic VRTSvcs VRTSvcsag VRTSvcsea

   VRTSvxfs
   ```

3. Uninstall SF packages from Solaris 10 brand zones.

   ```
   # pkgrm VRTSperl VRTSvlic VRTSvcs VRTSvcsag VRTSvcsea
   ```

Note: If you have SF packages installed inside non-global zones, perform the steps mentioned above to uninstall them from non-global zone before attempting to uninstall the packages from global zone.

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
```

Oracle:

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

DB2:

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

Oracle:

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

DB2:

```
# rm -rf /db2data/db2inst1/NODE0000/SQL00001/.sfae
```

3. Remove the repository location file.

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

This completes the removal of the SFDB repository.

Chapter 21

Uninstalling SF using response files

This chapter includes the following topics:

- Uninstalling SF using response files
- Response file variables to uninstall Storage Foundation
- Sample response file for SF uninstallation

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<version>
    -responsefile /tmp/response_file
   ```

   Where *<version>* is the specific release version, and /tmp/*response_file* is the response file's full path name.

   See "About the script-based installer" on page 46.

Response file variables to uninstall Storage Foundation

Table 21-1 lists the response file variables that you can define to configure SF.

Table 21-1 Response file variables for uninstalling SF

Variable	Description
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}{tmppath}	Defines the location where a working directory is created to store temporary files and the packages that are needed during the install. The default location is /var/tmp. List or scalar: scalar 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}{uninstall}	Uninstalls SF packages. List or scalar: scalar Optional or required: optional

Sample response file for SF uninstallation

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

```
our %CFG;

$CFG{opt}{redirect}=1;
$CFG{opt}{uninstall}=1;
$CFG{prod}="SF61";
$CFG{systems}=[ qw(thoropt89 thoropt90) ];

1;
```

Section 7

Installation reference

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

Appendix A

Installation scripts

This appendix includes the following topics:

- Installation script options

Installation script options

Table A-1 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 Symantec Storage Foundation product scripts, except where otherwise noted.

See "About the script-based installer" on page 46.

Table A-1 Available command line options

Command Line Option	Function
-allpkgs	Displays all packages required for the specified product. The packages are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The `-comcleanup` 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.
-comsetup	The `-comsetup` option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.
-configure	Configures the product after installation.

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

Command Line Option	Function
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases package upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-hotfix_path	Defines the path of a hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-hotfix2_path	Defines the path of a second hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-hotfix3_path	Defines the path of a third hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-hotfix4_path	Defines the path of a fourth hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-hotfix5_path	Defines the path of a fifth hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-installallpkgs	The `-installallpkgs` option is used to select all packages.
-installrecpkgs	The `-installrecpkgs`option is used to select the recommended packages set.
–installminpkgs	The `-installminpkgs`option is used to select the minimum packages set.

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

Command Line Option	Function
-ignorepatchreqs	The `-ignorepatchreqs` option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.
–jumpstart *dir_path*	Produces a sample finish file for Solaris JumpStart installation. The *dir_path* indicates the path to the directory in which to create the finish file.
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-I ssh_key_file` to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
-makeresponsefile	Use the `-makeresponsefile` option only to generate response files. No actual software installation occurs when you use this option.
-minpkgs	Displays the minimal packages required for the specified product. The packages are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-noipc	Disables the installer from making outbound networking calls to Symantec Operations Readiness Tool (SORT) in order to automatically obtain hot fix and release information updates.
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.

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

Command Line Option	Function
–pkginfo	Displays a list of packages and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS packages.
–pkgpath *package_path*	Designates the path of a directory that contains all packages to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
–pkgset	Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.
-pkgtable	Displays product's packages in correct installation order by group.
–postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
–recpkgs	Displays the recommended packages required for the specified product. The packages are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer hot fix file.
-requirements	The `-requirements` option displays required OS version, required packages and patches, file system space, and other system requirements in order to install the product.

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

Command Line Option	Function
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
–rootpath *root_path*	Specifies an alternative root directory on which to install packages. On Solaris operating systems, `-rootpath` passes `-R path` to `pkgadd` command.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See "About configuring secure shell or remote shell communication modes before installing products" on page 248.
–serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-settunables	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 `-tunablesfile` option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.

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

Command Line Option	Function
-timeout	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the `-timeout` option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 prevents the script from timing out. The `-timeout` option does not work with the `-serial` option
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file *tunables_file*	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 packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. Lists the installed patches, hot fixes, and available updates for the installed product if an Internet connection is available.

Appendix B

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 236.

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

```
# ./installer -tunablesfile tunables_file_name -settunables [
sys1 sys2 ...]
```

See “Setting tunables with no other installer-related operations” on page 237.

- 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 238.

See “About response files” on page 30.

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

See “Tunables value parameter definitions” on page 240.

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 240.

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 239.

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
 -settunables [sys1 sys2 ...]
```

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 240.

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 239.

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 `-settunables` option.

```
# ./installer -tunablesfile tunables_file_name -settunables [
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 240.

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 239.

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 `-responsefile` and `-tunablesfile` options.

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

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 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.

You can use the installer to create a tunables file template, or manually format tunables files you create.

To create a tunables file template

- Start the installer with the `-tunables` option. Enter the following:

  ```
  # ./installer -tunables
  ```

  You see a list of all supported tunables, and the location of the tunables file template.

To manually format tunables files

- 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 240.

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 *Symantec Storage Foundation Cluster File System High Availability Administrator's Guide* for detailed information on product tunable ranges and recommendations.

Table B-1 describes the supported tunable parameters that can be specified in a tunables file.

Table B-1 Supported tunable parameters

Tunable	Description
autoreminor	(Veritas Volume Manager) Enable reminoring in case of conflicts during disk group import.
autostartvolumes	(Veritas Volume Manager) Enable the automatic recovery of volumes.
dmp_cache_open	(Symantec Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached.
dmp_daemon_count	(Symantec Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks.
dmp_delayq_interval	(Symantec Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy.

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

Tunable	Description
dmp_fast_recovery	(Symantec Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Symantec Dynamic Multi-Pathing is started.
dmp_health_time	(Symantec Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy.
dmp_log_level	(Symantec Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed.
dmp_low_impact_probe	(Symantec Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled.
dmp_lun_retry_timeout	(Symantec Dynamic Multi-Pathing) The retry period for handling transient errors.
dmp_monitor_fabric	(Symantec Dynamic Multi-Pathing) Whether the Event Source daemon (`vxesd`) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Symantec Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Symantec Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events.
dmp_monitor_ownership	(Symantec Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored.
dmp_native_multipathing	(Symantec Dynamic Multi-Pathing) Whether DMP will intercept the I/Os directly on the raw OS paths or not.
dmp_native_support	(Symantec Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices.
dmp_path_age	(Symantec Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy.

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

Tunable	Description
dmp_pathswitch_blks_shift	(Symantec 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.
dmp_probe_idle_lun	(Symantec Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs.
dmp_probe_threshold	(Symantec Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon.
dmp_restore_cycles	(Symantec Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic.
dmp_restore_interval	(Symantec Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths.
dmp_restore_policy	(Symantec Dynamic Multi-Pathing) The policy used by DMP path restoration thread.
dmp_restore_state	(Symantec Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started.
dmp_retry_count	(Symantec Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed.
dmp_scsi_timeout	(Symantec Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP.
dmp_sfg_threshold	(Symantec Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature.
dmp_stat_interval	(Symantec Dynamic Multi-Pathing) The time interval between gathering DMP statistics.

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

Tunable	Description
fssmartmovethreshold	(Veritas Volume Manager) The file system usage threshold for SmartMove (percent). This tunable must be set after Veritas Volume Manager is started.
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer can only set the system default value of max_diskq. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
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 can only set the system default value of read_ahead. Refer to the tunefstab(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 can only set the system default value of read_nstream. Refer to the tunefstab(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 can only set the system default value of read_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
reclaim_on_delete_start_time	(Veritas Volume Manager) Time of day to start reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.

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

Tunable	Description
reclaim_on_delete_wait_period	(Veritas Volume Manager) Days to wait before starting reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.
same_key_for_alldgs	(Veritas Volume Manager) Use the same fencing key for all disk groups. This tunable must be set after Veritas Volume Manager is started.
sharedminorstart	(Veritas Volume Manager) Start of range to use for minor numbers for shared disk groups. This tunable must be set after Veritas Volume Manager is started.
storage_connectivity	(Veritas Volume Manager) The CVM storage connectivity type. This tunable must be set after Veritas Volume Manager is started.
usefssmartmove	(Veritas Volume Manager) Configure SmartMove feature (all, thinonly, none). This tunable must be set after Veritas Volume Manager is started.
vol_checkpt_default	(Veritas File System) Size of VxVM storage checkpoints (kBytes). This tunable requires a system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires a 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 a system reboot to take effect.

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

Tunable	Description
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires a 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).
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of Volume Replicator .
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (kBytes). This tunable requires a system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires a system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires a system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (kBytes). This tunable requires a 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) Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.

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

Tunable	Description
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires a system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions. This tunable requires a system reboot to take effect.
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires a system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (kBytes). This tunable requires a system reboot to take effect.
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 a system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires a system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires a 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 a system reboot to take effect.

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

Tunable	Description
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires a 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 a system reboot to take effect.
vx_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. This tunable requires a system reboot to take effect.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. This tunable requires a system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires a 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 can only set the system default value of write_nstream. Refer to the tunefstab(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 can only set the system default value of write_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.

Appendix C

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
- Restarting the ssh session
- Enabling and disabling rsh for Solaris

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

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

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

This section contains an example of how to set up ssh password free communication. The example sets up ssh between a source system (sys1) that contains the installation directories, and a target system (sys2). This procedure also applies to multiple target systems.

Note: The script- and web-based installers support establishing passwordless 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 C-1 illustrates this procedure.

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

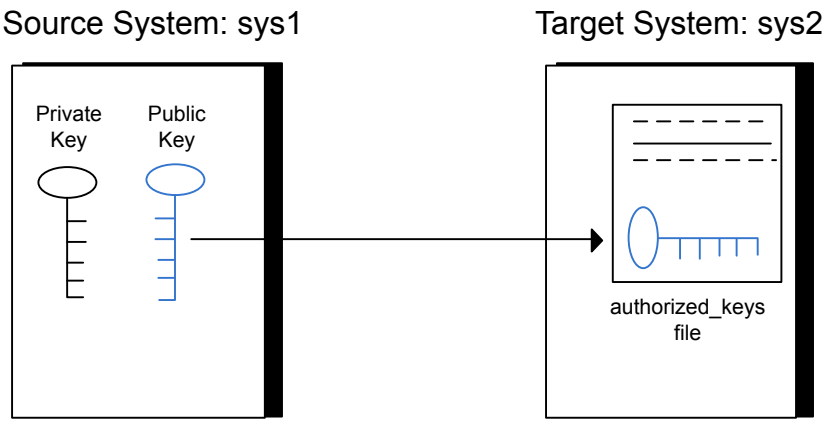

Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

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 (sys1), log in as root, and navigate to the root directory.

```
sys1 # cd /
```

2 Make sure the `/.ssh` directory is on all the target installation systems (sys2 in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

Solaris 10:

```
sys2 # mkdir /.ssh
```

Solaris 11:

```
sys2 # mkdir /root/.ssh
```

Change the permissions of this directory, to secure it.

Solaris 10:

```
sys2 # chmod go-w /.ssh
```

Solaris 11:

```
sys2 # chmod go-w /root/.ssh
```

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

```
sys1 # 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):
```

For Solaris 11:

```
Your identification has been saved in /root/.ssh/id_dsa.
Your public key has been saved in /root/.ssh/id_dsa.pub.
```

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

5 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.

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 (sys2 in this example).

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

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

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

To restart ssh on Solaris 10 and Solaris 11, type the following command:

```
sys1 # svcadm restart ssh
```

3 From the source system (sys1), move the public key to a temporary file on the target system (sys2).

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:

```
sys1 # sftp sys2
```

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 sys2 ...
The authenticity of host 'sys2 (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 'sys2,10.182.00.00'
(DSA) to the list of known hosts.
root@sys2 password:
```

5 Enter the root password of sys2.

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 (sys2 in this example), type the following command on sys1:

```
sys1 # ssh sys2
```

Enter the root password of sys2 at the prompt:

```
password:
```

9 After you log in to sys2, enter the following command to append the `id_dsa.pub` file to the `authorized_keys` file:

```
sys2 # cat /id_dsa.pub >> /.ssh/authorized_keys
```

10 After the `id_dsa.pub` public key file is copied to the target system (sys2), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, enter the following command on sys2:

```
sys2 # rm /id_dsa.pub
```

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

```
sys2 # exit
```

12 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`:

```
sys1 # exec /usr/bin/ssh-agent $SHELL
sys1 # 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 (sys1), enter the following command:

```
sys1 # ssh -l root sys2 uname -a
```

where sys2 is the name of the target system.

2 The command should execute from the source system (sys1) to the target system (sys2) without the system requesting a passphrase or password.

3 Repeat this procedure for each target system.

Restarting the ssh session

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

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

To restart ssh

1 On the source installation system (sys1), bring the private key into the shell environment.

```
sys1 # exec /usr/bin/ssh-agent $SHELL
```

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

```
sys1 # ssh-add
```

Enabling and disabling rsh for Solaris

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

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

See "Manually configuring and passwordless ssh" on page 249.

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

To enable rsh

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

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

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

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

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

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

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

```
# inetadm -e rlogin
```

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

```
# inetadm -d rlogin
```

4 Modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. This file must be modified for each user who remotely accesses the system using rsh. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system having access to the local system. For example, if the root user must remotely access `sys1` from `sys2`, you must add an entry for `sys2.`*`companyname`*`.com` in the `.rhosts` file on `sys1`.

```
# echo "sys2.companyname.com" >> $HOME/.rhosts
```

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

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

Appendix D

Storage Foundation components

This appendix includes the following topics:

- Storage Foundation installation packages
- Chinese language packages
- Japanese language packages
- Symantec Storage Foundation obsolete and reorganized installation packages

Storage Foundation installation packages

Table D-1 shows the package name and contents for each English language package for Storage Foundation. The table also gives you guidelines for which packages to install based whether you want the minimum, recommended, or advanced configuration.

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

Table D-1 Storage Foundation packages

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

Table D-1 Storage Foundation packages *(continued)*

packages	Contents	Configuration
VRTSperl	Perl 5.16.1 for Veritas	Minimum
VRTSvlic	Symantec 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	Storage Management Software for Databases	Recommended
VRTSob	Veritas Enterprise Administrator Service	Recommended
VRTSodm	Veritas Extension for Oracle Disk Manager Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle. Oracle Disk Manager enables Oracle to improve performance and manage system bandwidth.	Recommended

Table D-1 Storage Foundation packages *(continued)*

packages	Contents	Configuration
VRTSsfcpi61	Symantec Storage Foundation Installer The Storage Foundation Common Product installer package contains the installer libraries and product scripts that perform the following: ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ etc. You can use these script to simplify the native operating system installations, configurations, and upgrades.	Minimum
VRTSsfmh	Veritas Operations Manager Managed Host. Discovers configuration information on a Storage Foundation managed host. If you want a central server to manage and monitor this managed host, download and install the VRTSsfmcs package on a server, and add this managed host to the Central Server. The VRTSsfmcs package is not part of this release. You can download it separately from: http://www.symantec.com/veritas-operations-manager	Recommended
VRTSspt	Veritas Software Support Tools	Recommended
VRTSfsadv	Veritas File System Advanced	Minimum
VRTSfssdk	Veritas File System Software Developer Kit For VxFS APIs, the package contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.	All

Chinese language packages

The following table shows the package name and contents for each Chinese language package.

Table D-2 Chinese language packages

package	Contents
VRTSzhvm	Chinese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages

Japanese language packages

The following table show the package name and contents for each Japanese language package.

Table D-3 Japanese language packages

package	Contents
VRTSjacav	Japanese Symantec Cluster Server Agents for Storage Foundation Cluster File System – Manual Pages and Message Catalogs by Symantec
VRTSjacs	Symantec Cluster Server Japanese Message Catalogs by Symantec
VRTSjacse	Japanese Symantec High Availability Enterprise Agents by Symantec
VRTSjadba	Japanese Symantec Oracle Real Application Cluster Support package by Symantec
VRTSjadbe	Japanese Symantec Storage Foundation for Oracle from Symantec – Message Catalogs
VRTSjafs	Japanese Veritas File System – Message Catalog and Manual Pages
VRTSjaodm	Veritas Oracle Disk Manager Japanese Message Catalog and Manual Pages by Symantec
VRTSjavm	Japanese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages
VRTSmulic	Multi-language Symantec License Utilities

Symantec Storage Foundation obsolete and reorganized installation packages

Table D-4 lists the packages that are obsolete or reorganized for Storage Foundation.

Table D-4 Symantec Storage Foundation obsolete and reorganized packages

package	Description
Obsolete and reorganized for 6.1	
VRTSat	Obsolete
VRTSatZH	Obsolete
VRTSatJA	Obsolete
Obsolete and reorganized for 5.1	
Infrastructure	
SYMClma	Obsolete
VRTSaa	Included in VRTSsfmh
VRTSccg	Included in VRTSsfmh
VRTSdbms3	Obsolete
VRTSicsco	Obsolete
VRTSjre	Obsolete
VRTSjre15	Obsolete
VRTSmh	Included in VRTSsfmh
VRTSobc33	Obsolete
VRTSobweb	Obsolete
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete

Table D-4 Symantec Storage Foundation obsolete and reorganized packages *(continued)*

package	Description
Product packages	
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, the existing VRTSacclib is uninstalled and a new VRTSacclib is installed. ■ For uninstallation, VRTSacclib is not uninstalled.
VRTSalloc	Obsolete
VRTScmccc	Obsolete
VRTScmcm	Obsolete
VRTScmcs	Obsolete
VRTScscm	Obsolete
VRTScscw	Obsolete
VRTScsocw	Obsolete
VRTScssim	Obsolete
VRTScutil	Obsolete
VRTSd2gui	Included in VRTSdbed
VRTSdb2ed	Included in VRTSdbed
VRTSdbcom	Included in VRTSdbed
VRTSdbed	Included in VRTSdbed
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSdsa	Obsolete

Table D-4 Symantec Storage Foundation obsolete and reorganized packages *(continued)*

package	Description
VRTSfas	Obsolete
VRTSfasag	Obsolete
VRTSfsman	Included in the product's main package.
VRTSfsmnd	Included in the product's main package.
VRTSfspro	Included in VRTSsfmh
VRTSgapms	Obsolete
VRTSmapro	Included in VRTSsfmh
VRTSorgui	Obsolete
VRTSsybed	Included in VRTSdbed
VRTSvail	Obsolete
VRTSvcsdb	Included in VRTSvcsea
VRTSvcsmn	Included in VRTSvcs
VRTSvcsor	Included in VRTSvcsea
VRTSvcssy	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvcs
VRTSvdid	Obsolete
VRTSvmman	Included in the product's main package.
VRTSvmpro	Included in VRTSsfmh
VRTSvrpro	Included in VRTSob
VRTSvrw	Obsolete
VRTSvxmsa	Obsolete

Appendix E

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
- About the VRTSspt package troubleshooting tools
- Incorrect permissions for root on remote system
- Inaccessible system
- Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.1 (2184482)
- Troubleshooting the webinstaller

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 Symantec Storage
Foundation/Symantec 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' and validate using the command
'vxkeyless set NONE'.
```

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 Symantec product license keys" on page 36.
  After you install the license key, you must validate the license key using the following command:

  ```
  # /opt/VRTS/bin/vxlicrep
  ```

- 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

About the VRTSspt package troubleshooting tools

The VRTSspt package provides a group of tools for troubleshooting a system and collecting information on its configuration. If you install and use the VRTSspt package, it will be easier for Symantec Support to diagnose any issues you may have.

The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. Although the tools are not required for the operation of any Symantec product, Symantec recommends installing them should a support case be needed to be opened with Symantec Support. Use caution when you use the VRTSspt package, and always use it in concert with Symantec Support.

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.

```
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 248.

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

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
 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.

Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.1 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.1.

When upgrading from SF version 5.0 to SF 6.1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 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.

Troubleshooting the webinstaller

This section provides possible solutions to problems that may occur when using the `webinstaller` script:

- Issue: The `webinstaller` script may report an error.

You may receive a similar error message when using the webinstaller:

```
Error: could not get hostname and IP address
```

Solution: Check whether `/etc/hosts` and `/etc/resolv.conf` file are correctly configured.

- Issue: The hostname is not a fully qualified domain name.
  You must have a fully qualified domain name for the hostname in https://<*hostname*>:<*port*>/.
  Solution: Check whether the `domain` section is defined in `/etc/resolv.conf` file.
- Issue: FireFox 3 may report an error.
  You may receive a similar error message when using FireFox 3:

  ```
  Certificate contains the same serial number as another certificate.
  ```

  Solution: Visit FireFox knowledge base website:

http://support.mozilla.com/en-US/kb/Certificate+contains+the+same+serial+number+as+another+certificate

Appendix F

Compatability issues when installing Storage Foundation with other products

This appendix includes the following topics:

- Installing, uninstalling, or upgrading Storage Foundation products when other Symantec 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 Symantec products are present

Installing Storage Foundation when other Symantec 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 VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the VOM Central Server and Managed Host packages as is.
- When uninstalling Storage Foundation products where VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where 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 VRTSpbx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpbx, VRTSicsco, and VRTSat.

Index

A

about
- Deployment Server 90
- installation and configuration methods 28
- installation preparation 39
- installation using operating system-specific methods 64
- planning for installation 27
- response files 30
- SORT 17
- Symantec product licensing 33
- Veritas Operations Manager 17
- web-based installer 51

assesing system
- installation readiness 43

Automated installer
- about 65
- installing 65
- using 65

B

backup boot disk group 178
- rejoining 178

BE upgrade
- administering BE in Solaris 11 172
- completing Solaris 11 upgrade 170
- creating new Solaris 11 BE 168
- upgrading Solaris 11 using the installer 168
- verifying Solaris 11 upgrade 171

before using
- web-based installer 52

bootdg 85

C

changing root user 200

checking
- installation readiness 43

checking product versions 23

configuration daemon (vxconfigd)
- starting 83

configuring
- rsh 40
- ssh 40

creating
- /opt directory 42
- backups 120
- Flash archive 74
- post-deployment scripts 75

creating root user 41

D

default disk group 85

defaultdg 85

deploying
- specific Symantec release 108
- Symantec product updates to your environment 105

deployment management
- overview 92

deployment preferences
- setting 95

Deployment Script
- installing 91

Deployment Server
- about 90
- setting up 93

Deployment Server command line option
- for loading and downloading the most recent release information 97
- for specifying a non-default repository location 96

disabling
- external network connection attempts 25

disk groups
- bootdg 85
- default 85
- nodg 85
- root 85
- rootdg 83, 85

disk space requirements 22

downloading maintenance releases and hot fixes 23

E

EMC powerpath
converting a foreign disk to auto:simple 189
EMC PowerPath disk
converting a defined disk to auto:simple 191
converting a powervxvm disk to auto:simple 194

F

flarcreate 74
Flash archive 74
post-deployment scripts 75

I

I/O daemon (vxiod)
starting 84
Install Bundles
integration options 127
installer
about the script-based installer 46
installer hot fixes
obtaining either manually or automatically 24
Installing
SF with the web-based installer 54
web-based installer 54
installing
Automated Installer 65
JumpStart 70
language packages 50
SF using operating system-specific methods 64
Symantec product license keys 36
the Deployment Script 91
using Flash archive 74
using response files 57
using the system command 76

J

JumpStart
installing 70
Jumpstart
Generating the finish scripts 71
overview 70
Preparing installation resources 72

K

keyless licensing
setting or changing the product level 34

L

language packages
removal 222
licensing
installing Symantec product license keys 36
setting or changing the product level for keyless licensing 34
Live Upgrade
administering Solaris 10 boot environments 165
completing Solaris 10 upgrade 163
manual upgrade for Solaris 10 161
manual upgrade for Solaris 11 169
preparing 156
reverting to primary boot environment 165
Solaris 10 systems 155
supported upgrade paths 153
Switching boot environment for Solaris SPARC 165
upgrading Solaris 10 on alternate boot disk 157
upgrading Solaris 10 using the installer 161
verifying Solaris 10 upgrade 164
VVR environment 173
loading and downloading the most recent release information
Deployment Server command line option for 97
localized environment settings for using VVR
settings for uisng VVR in a localized environment 125

M

migrating
from Storage Foundation Basic to Storage Foundation Standard 174
mounting
software disc 42

N

nodg 85

O

obtaining
installer hot fixes either automatically or manually 24
security exception on Mozilla Firefox 53
overview
deployment management 92

P

planning to upgrade VVR 122
post-deployment scripts 75
post-upgrade
    adding JBOD support 187
    unsuppressing DMP for EMC PowerPath disks 188
    updating variables 186
    upgrading the array support library 187
    verifying 198
prechecking
    using the installer 43
preinstallation 122
preinstallation check
    web-based installer 54
preparing
    Live Upgrade 156
preparing to upgrade 117

R

rejoining
    backup boot disk group 178
release images
    viewing or downloading available 98
release information
    updating on systems without Internet access 109
release notes 20
releases
    finding out which releases you have, and which upgrades or updates you may need 106
removing
    the Replicated Data Set 217
Replicated Data Set
    removing the 217
repository images
    viewing and removing repository images stored in your repository 102
response files
    about 30
    installation 57
    syntax 31
    uninstalling 225
    upgrading 146
root disk group 83, 85
rootdg 85
rsh
    configuration 40

S

script-based installer
    about 46
setting
    deployment preferences 95
    environment variables 42
setting up
    Deployment Server 93
settings for uisng VVR in a localized environment
    localized environment settings for using VVR 125
SF installation
    preinstallation information 21
simulataneous install or upgrade 127
specifying
    non-default repository location 96
ssh
    configuration 40
starting
    web-based installer 52
starting vxconfigd configuration daemon 83
starting vxiod daemon 84
Storage Foundation Basic to Storage Foundation Standard
    migrating 174
supported upgrade paths
    Live Upgrade 153
Symantec product license keys
    installing 36
Symantec product updates
    deploying to your environment 105
Symantec products
    starting process 203
    stopping process 203
Symantec release
    deploying a specific release 108

T

tunables file
    about setting parameters 235
    parameter definitions 240
    preparing 239
    setting for configuration 236
    setting for installation 236
    setting for upgrade 236
    setting parameters 239
    setting with no other operations 237
    setting with un-integrated response file 238

U

uninstalling
about removing Storage Foundation 207
language packages 222
moving volumes from an encapsulated root disk 208
moving volumes to disk partitions 209
preparing to remove Veritas File System 216
preparing to remove Veritas Volume Manager 208
preparing to uninstall 208
using pkg uninstall command 221
using pkgrm command 221
using response files 225
using the web-based installer 220
unsuccessful upgrade 178
updating release information
on systems without Internet access 109
upgrade
array support 126
creating backups 120
getting ready 117
methods 112
supported upgrade paths 113
upgrades or updates
finding out which releases you have 106
upgrading
language packages 145
OS upgrade 136
Solaris operating system 141
using product installer 130
using response files 146
using the product installer 134
using the product installer or manual steps 133
using the web-based installer 139
using upgrade scripts 136
upgrading VVR
from 4.1 123
planning 122
using Live Upgrade 151

V

verifying
product installation 202
verifying installation
kernel component 205
viewing and removing repository images
stored in your repository 102
viewing or downloading
available release images 98
vradmin
delpri 218
stoprep 217
VVR 4.1
planning an upgrade from 123
vxconfigd configuration daemon
starting 83
vxdctl mode command 83
vxinstall program 84–85
vxinstall program, running 84
vxiod I/O daemon
starting 84
vxplex
used to remove mirrors of root disk volumes 131, 134

W

web-based installer 54
about 51
before using 52
installation 54
preinstallation check 54
starting 52
uninstalling 220
upgrading 139